A Statistical-Mechanical Approach to Data Assimilation for Nonlinear Dynamics

Gregory L. Eyink
Mathematical Sciences Department
The Johns Hopkins University, Baltimore, MD 21218, U.S.A.

Juan M. Restrepo
Department of Mathematics and Department of Physics
University of Arizona, Tucson, AZ 85721, U.S.A.

Francis J. Alexander
CCS-3, Los Alamos National Laboratory
Los Alamos, NM, 87545, U.S.A.

October 3, 2003

Abstract

The calculation of conditional statistics for the purpose of data assimilation in large-scale systems requires an approximation to the dynamical evolution of probability distributions. In this paper, two approximation schemes are considered: moment-closure methods and empirical-ensemble methods. Recently, a mean-field conditional analysis was proposed to incorporate the results of measurements [1]. Within that approximation, it is found to be advantageous to construct moment-closures to preserve an exact H-theorem for the "relative entropy." The closure and ensemble methods are compared here for a simple model system with bimodal statistics. The mean-field variational closure methods give results that are consistently more accurate than those from ensemble Kalman methods. The exact conditional analysis requires a closure of at least second-order, but the mean-field analysis can employ first-order closures as well as second-order. First-order closures are found to be nearly as accurate as second-order closures in the simple test problem.

Key Words: Filtering, Smoothing, Optimal Estimation, Ensemble Methods

1 Introduction

This paper presents a new probabilistic approach to data assimilation and inverse modeling for large-scale nonlinear dynamical systems, based upon the approximate calculation of conditional statistics. It builds on recent work in Eyink et al. [1], where we discussed the general probabilistic setting of optimal estimation for nonlinear stochastic dynamics. As described there, the exact calculation of conditional statistics requires the evolution of probability distributions, by means of the forward-backward Kolmogorov equations, and the combination or "analysis" of prior distributions with imprecise measurements, by means of Bayes formula. The main aim of Eyink et al. [1] was to develop a mean-field approximation to the "analysis" or conditioning step of the calculation. It was shown by consideration of a statistical mechanics of time-histories that a mean-field conditional analysis can be obtained by the minimization of a nonnegative, convex cost function, the *multi-time relative entropy.* The calculation of this cost function and its gradient requires also the solution of the forward-backward Kolmogorov equations. It is the purpose of this paper to consider approximations to the dynamical evolution given by these equations. Two approximations especially shall be considered here: Monte Carlo or empirical ensemble methods—as originally proposed by Leith [2]—and moment-closure methods such as used in theory of turbulent fluids or kinetic theory of gases.

It is important to emphasize that the Kolmogorov equations for the probability distributions are themselves completely impractical to solve for such large-scale systems. If the model dynamics of the system has D degrees-of-freedom, then the Kolmogorov equations are partial differential equations in a hyperspatial domain of dimension D. For all except the simplest, low-dimensional dynamical systems, it is impossible to solve such partial differential equations. Therefore, methods to calculate the exact conditional statistics by means of such equations are out of reach. This was already realized by one of the pioneers in the field, H. J. Kushner, who solved the problem of optimal filtering for nonlinear dynamics [see 3, 4]. Nevertheless Kushner realized that the exact mathematical solution by his (stochastic) partial differential equations could not be practically applied to spatially-extended or "distributed" systems of interest in many applications. Kushner [5] pointed out the analogy of this problem to the closure problem of turbulence theory and he developed moment-closure approximations to his exact equations for optimal filtering. Such approaches have also been advocated in the meteorological community to to study predictability of forecasts (see Ehrendorfer [6] for a review). In the present paper [see also 7] this closure method is extended to the problem of optimal smoothing, employing future measurements as well as past ones. For this purpose, we exploit an action principle for the forward-backward Kolmogorov equa-

tions, which are the corresponding Euler-Lagrange equations. Moment-closure methods can then be formulated as a Rayleigh-Ritz approximation, incorporating guesses of the statistics of the system by means of precise analytical ansätze. These methods can be applied both within the exact conditional analysis and the mean-field conditional analysis. However, the latter has the important advantage that for it first-order closures can be successfully employed, whereas closures using the exact conditional analysis must be at least second-order.

The concept of relative entropy, which already appeared fundamentally in formulating the mean-field analysis in [1], plays also an essential role in the moment-closure methods employed here. The important application of entropy and dissipation functions to state estimation has been appreciated in statistical physics since the work of Onsager and Machlup [8], a development of Onsager's earlier ground-breaking work on irreversible thermodynamics. For the statistical Markov models of greatest interest in the geosciences, the (single-time) relative entropy is a Lyapunov function, providing a generalized second law of thermodynamics. Recently, Kleeman [9] has emphasized this point for stochastic climate models and proposed to use relative entropy as a measure of the potential utility of dynamical predictions. The moment-closures considered in this work are entropy-based and ensure the preservation of the exact H-theorem for the relative entropy [see 10]. Such closure methods are modeled after those successful employed in kinetic theory [see 11, 12]. An important bonus of the entropy-based closure methods is that they provide a closed-form expression for the approximate cost function or *effective action* (see [7][1]. The resulting cost function has exactly the form of the Onsager-Machlup action in statistical physics. See section 2.1.2 below. This yields an intuitive interpretation of the newly-proposed cost function in terms of the enhanced dissipation and noise produced by small-scale chaotic or stochastic dynamics.

The other main evolution approximation considered in this work is the Monte Carlo scheme employing empirical ensembles of N sample realizations of the stochastic model systems, as proposed by Leith [2]. In recent years, such methods for data assimilation and inverse modeling in geophysical applications have been extensively developed by G. Evensen and his collaborators [see 13, 14, 15, 16]. See Tippett et al. [17] for a review of more recent developments. We regard such methods as among the most promising presently being considered and they shall be used as the primary basis of comparison for the moment-closure methods advocated in this work. We hope to convince the reader that the Rayleigh-Ritz or moment-closure approximations to the evolution are complementary to the empirical-ensemble methods and competitive, at least, in accuracy and numerical efficiency.

The detailed contents of this paper are as follows: In Section 2 we discuss the two main evolution approximations considered in this work,

moment-closure methods and empirical-ensemble methods. The former are shown there to correspond to a Rayleigh-Ritz approximation derived from a variational formulation of the forward-backward Kolmogorov equations. Two alternative Rayleigh-Ritz approaches are considered, which make different trial guesses for the solution of the backward equation, a linear ansatz and an exponential ansatz. The entropy-based closure scheme for the forward equation using also an exponential ansatz is briefly described. In Section 3 we develop concrete moment-closures for a simple model as a test case. Both first-order and second-order moment closures are considered. Finally, in Section 4 we give the results of assimilation experiments for the model problem. The approximation methods discussed in this paper, both ensemble methods and the array of closure methods developed in the earlier sections, are compared with each other and also with the methods using the exact dynamical evolution. Section 5 contains our summary and final discussion.

2 Evolution Approximations

2.1 Moment-Closure Methods (Rayleigh-Ritz)

We now introduce approximate methods of solving the evolution equations by means of moment-closure. The basis of these approximation schemes is the following variational principle. If $\hat{L}(t)$ is the generator of a Markov process, define an *action functional*

$$\Gamma[\mathcal{A},\mathcal{P}] := \int_{t_i}^{t_f} dt \int d\mathbf{x}\ \mathcal{A}(\mathbf{x},t)(\partial_t - \hat{L}(t))\mathcal{P}(\mathbf{x},t) \tag{1}$$

where $\mathcal{A}$ are bounded variables ($\mathcal{A} \in L^\infty$) and $\mathcal{P}$ are distributions with finite integral ($\mathcal{P} \in L^1$). Then, the forward and backward Kolmogorov equations, [see 18], are simultaneously obtained as Euler-Lagrange equations for stationarity of $\Gamma[\mathcal{A},\mathcal{P}]$. The forward equation

$$\partial_t \mathcal{P}_F(\mathbf{x},t) = \hat{L}(t)\mathcal{P}_F(\mathbf{x},t), \tag{2}$$

is obtained by variation over $\mathcal{A}$ with final condition $\mathcal{A}(t_f) \equiv 1$, while the backward equation

$$\partial_t \mathcal{A}_S(\mathbf{x},t) + \hat{L}^*(t)\mathcal{A}_S(\mathbf{x},t) = 0, \tag{3}$$

is obtained by variation over $\mathcal{P}$ with initial condition $\mathcal{P}(t_i) = \mathcal{P}_0$. As a particular example, consider a general stochastic diffusion process with "drift vector" $\mathbf{f}(\mathbf{x},t)$ and "diffusion matrix" $\mathbf{D}(\mathbf{x},t)$. In that case, $\hat{L}(t) = -\boldsymbol{\nabla}_\mathbf{x}\cdot[\mathbf{f}(\mathbf{x},t)(\cdot)] + \boldsymbol{\nabla}_\mathbf{x}\boldsymbol{\nabla}_\mathbf{x}^\top\!:[\mathbf{D}(\mathbf{x},t)(\cdot)]$ is the Fokker-Planck operator and

$\hat{L}^*(t) = \mathbf{f}(\mathbf{x},t)\cdot\boldsymbol{\nabla}_{\mathbf{x}} + \mathbf{D}(\mathbf{x},t):\boldsymbol{\nabla}_{\mathbf{x}}\boldsymbol{\nabla}_{\mathbf{x}}^{\top}$ is the adjoint Fokker-Planck operator. See Balian and Vénéroni [19] for the general approach. Such a variational formulation invites the use of Rayleigh-Ritz approximation methods, in which, rather than varying over all $\mathcal{A} \in L^{\infty}, \mathcal{P} \in L^1$, one varies only over finitely parametrized trial functions. Such methods turn out to be equivalent to standard moment-closure approximations. The right trial function $\mathcal{P}$ represents one's guess of the statistics of the system, while the left trial function $\mathcal{A}$ represents one's guess of the significant dynamical variables. Within this framework there is also a useful variational characterization of the effective action $\Gamma_X[\mathbf{x}]$, which was a principle tool in [1]. It is obtained by a *constrained variation* of the action $\Gamma[\mathcal{A},\mathcal{P}]$. In fact, the effective action is its stationary point (st.pt)

$$\Gamma_X[\mathbf{x}] = \text{st.pt.}_{\mathcal{A},\mathcal{P}}\Gamma[\mathcal{A},\mathcal{P}] \tag{4}$$

when varied over the same classes as above, but subject to constraints of fixed overlap

$$\int d\mathbf{x}\ \mathcal{A}(\mathbf{x},t)\mathcal{P}(\mathbf{x},t) = 1 \tag{5}$$

and fixed expectation

$$\int d\mathbf{x}\ \mathbf{x}\mathcal{A}(\mathbf{x},t)\mathcal{P}(\mathbf{x},t) = \mathbf{x}(t) \tag{6}$$

for all $t \in [t_i, t_f]$. See Eyink [7]. This theorem may be made the basis of an algorithm to calculate the effective action by employing the forward-backward Kolmogorov equations.

As discussed in detail in [1], the "mean-field approximation" to conditional statistics of the state variable given measurements may be obtained by minimizing the joint effective action $\Gamma_{X,Y}[\mathbf{x},\mathbf{y}]$ over the state variable $\mathbf{x}$ for fixed values of the measurements $\mathbf{y}$. The main quantity which needs to be calculated is $\Gamma_X[\mathbf{x}]$, since $\Gamma_{X,Y}[\mathbf{x},\mathbf{y}] = \Gamma_X[\mathbf{x}] + \Gamma_{Y|X}[\mathbf{y}|\mathbf{x}]$ and the "conditional action" of the measurements given the state variables is generally a simple known function. For example, for the measurement of a linear function $\mathbf{H}_m\mathbf{x}(t_m)$ with normally distributed errors of zero mean and covariance $\mathbf{R}_m$, at a sequence of times t_m, $m = 1, ..., M$, the conditional action is

$$\Gamma_{Y|X}[\mathbf{y}|\mathbf{x}] = \frac{1}{2}\sum_{m=1}^{M}[\mathbf{y}_m - \mathbf{H}_m\mathbf{x}(t_m)]^{\top}\mathbf{R}_m^{-1}[\mathbf{y}_m - \mathbf{H}_m\mathbf{x}(t_m)]. \tag{7}$$

Thus, the variational characterization of the effective action in (4) provides a means to determine the main unknown for variational data-assimilation within a moment-closure approximation of the time-evolution. In the context of discrete-time measurements, as in (7) above, it is sometimes useful

to consider the "multi-time (relative) entropy" which depends only upon state-variables at the measurement times. It is defined as

$$H_X(\mathbf{x}_1, ..., \mathbf{x}_m) = \min_{\{\mathbf{x}:\mathbf{x}(t_m)=\mathbf{x}_m,\, m=1,...,M\}} \Gamma_X[\mathbf{x}], \tag{8}$$

by "contraction" over all histories with the specified values at the measurement times. The variational principle (4) provides also the means to calculate this function.

We now discuss two somewhat different approximation schemes based upon these results.

2.1.1 Left-Linear Ansatz

In this approach, the trial functions $\mathcal{A}, \mathcal{P}$ employed are constructed from the usual elements of a moment-closure: a set of *moment functions* $M_i(\mathbf{x}, t),\ i = 1, ..., R$ and a probability density function or PDF ansatz $\mathcal{P}(\mathbf{x}, t; \boldsymbol{\mu})$, which is conveniently parametrized by the mean values which it attributes to the moment-functions, $\boldsymbol{\mu} := \int d\mathbf{x}\ \mathcal{P}(\mathbf{x}, t; \boldsymbol{\mu})\mathbf{M}(\mathbf{x}, t)$. $\mathcal{P}(\mathbf{x}, t; \boldsymbol{\mu})$ is itself the right trial function in this approach. The left trial function will be taken to be a linear combination of the (centered) moment functions

$$\mathcal{A}(\mathbf{x}, t; \boldsymbol{\alpha}) := 1 + \sum_{i=1}^{R} \alpha_i [M_i(\mathbf{x}, t) - \boldsymbol{\mu}(t)]. \tag{9}$$

The histories $\boldsymbol{\alpha}(t), \boldsymbol{\mu}(t)$ are the parameters to be varied over. Substituting the trial forms, one obtains the reduced action

$$\Gamma[\boldsymbol{\alpha}, \boldsymbol{\mu}] = \int_{t_i}^{t_f} dt\ \boldsymbol{\alpha}^\top(t)[\dot{\boldsymbol{\mu}}(t) - \mathbf{V}(\boldsymbol{\mu}(t), t)] \tag{10}$$

with

$$\mathbf{V}(\boldsymbol{\mu}, t) := \langle \dot{\mathbf{M}}(t) \rangle_{\boldsymbol{\mu}(t)} = \langle (\partial_t + \hat{L}^*)\mathbf{M}(t) \rangle_{\boldsymbol{\mu}(t)}. \tag{11}$$

$\langle \cdot \rangle_{\boldsymbol{\mu}(t)}$ denotes average with respect to the PDF ansatz. The Euler-Lagrange equations are

$$\dot{\boldsymbol{\mu}} = \mathbf{V}(\boldsymbol{\mu}, t) \tag{12}$$

and

$$\dot{\boldsymbol{\alpha}} + \left(\frac{\partial \mathbf{V}}{\partial \boldsymbol{\mu}}\right)^\top (\boldsymbol{\mu}, t)\boldsymbol{\alpha} = \mathbf{0}. \tag{13}$$

These are solved subject to initial condition $\boldsymbol{\mu}(t_i) = \boldsymbol{\mu}_0$ and final condition $\boldsymbol{\alpha}(t_f) = \mathbf{0}$. The forward equation (12) is just the standard moment-closure equation for the set of moments $M_i(\mathbf{x}, t),\ i = 1, ..., R$ and the PDF ansatz $\mathcal{P}(\mathbf{x}, t; \boldsymbol{\mu})$. The equation (13) is its adjoint. An advantage of this scheme

is that the Euler-Lagrange equations (12), (13) are partially decoupled: the forward equation does not depend upon the solution of the backward equation. Hence, the system may be solved by integrating (12) forward in time and storing its solution $\boldsymbol{\mu}(t)$ for use in the integration of (13) backward in time.

It was shown in Eyink et al. [1] that the "multi-time entropy" H_X and its Legendre dual function F_X may be calculated by solving the forward-backward Kolmogorov equations between measurement times, with certain "jump conditions" at the measurements: [1], equations (24),(28). To evaluate the functions F_X, H_X within a moment-closure approximation to the time-evolution, we must also formulate the jump conditions within the closure. For a general closure, it is not at all obvious how to do this. A practical way around this difficulty is to consider the measurements to be continuous in time, averaged with respect to a test function $\delta_\tau(t - t_m)$ localized within a time range τ of the measurement time t_m. This is a realistic assumption, where τ represents the duration of the measurement. However, it is useful to develop an approach in which proper jump conditions can be formulated as $\tau \to 0$. This can be done *(i)* if the quantity being measured, here the state $\mathbf{x}$, is included among the moment variables $M_i(\mathbf{x}, t),\ i = 1, ..., R$ in the closure and *(ii)* if a special form is used for the right trial function, an *exponential ansatz*

$$\mathcal{P}(\mathbf{x}, t; \boldsymbol{\lambda}) = \frac{\exp\left[\boldsymbol{\lambda}^\top \mathbf{M}(\mathbf{x}, t)\right]}{\mathcal{N}_M(\boldsymbol{\lambda}, t)} \mathcal{P}_*(\mathbf{x}, t), \tag{14}$$

where $\mathcal{P}_*(\mathbf{x}, t)$ is a suitable reference PDF, and

$$\mathcal{N}_M(\boldsymbol{\lambda}, t) = \int d\mathbf{x}\ \exp\left[\boldsymbol{\lambda}^\top \mathbf{M}(\mathbf{x}, t)\right] \mathcal{P}_*(\mathbf{x}, t). \tag{15}$$

The condition *(i)* is often satisfied, since the most natural closure equation is the average of the dynamical equation itself, e.g. for a diffusion process:

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t) + \mathbf{D}^{1/2}(\mathbf{x}, t)\mathbf{q}(t). \tag{16}$$

Condition *(ii)* of an exponential form (14) also does not seriously limit the formulation of suitable closures. All that is required is that exponential moments like those in (15) be calculable for the reference PDF $\mathcal{P}_*$ (either analytically or numerically).

We now state the jump conditions that apply when conditions *(i),(ii)* are satisfied. Since *(i)* holds, we might as well calculate the thermodynamic functions $F_M(\boldsymbol{\lambda}_1, ..., \boldsymbol{\lambda}_M), H_M(\mathbf{m}_1, ..., \mathbf{m}_M)$ appropriate to all of the closure variables $M_i(\mathbf{x}, t),\ i = 1, ..., R$ not just the state variables $x_i,\ i = 1, ..., d$.

Suitable jump conditions for the forward equation (12) are

$$\boldsymbol{\lambda}(\boldsymbol{\mu}_m^+, t_m) = \boldsymbol{\lambda}(\boldsymbol{\mu}_m^-, t_m) + \boldsymbol{\lambda}_m, \quad m = 1, ..., M \tag{17}$$

Here, $\boldsymbol{\lambda}(\boldsymbol{\mu}, t)$ is the inverse of the moment function

$$\boldsymbol{\mu}(\boldsymbol{\lambda}, t) = \int d\mathbf{x}\ \mathbf{M}(\mathbf{x}, t)\mathcal{P}(\mathbf{x}, t; \boldsymbol{\lambda}).$$

In fact, $\boldsymbol{\lambda}(\boldsymbol{\mu}, t)$ may be found from the Legendre transformation for the *(single-time) relative entropy*

$$H_M(\boldsymbol{\mu}, t) = \max_{\boldsymbol{\lambda}} \{\boldsymbol{\mu}^\top \boldsymbol{\lambda} - F_M(\boldsymbol{\lambda}, t)\} \tag{18}$$

where $F_M(\boldsymbol{\lambda}, t) = \log \mathcal{N}_M(\boldsymbol{\lambda}, t)$. That is, $\boldsymbol{\lambda}(\boldsymbol{\mu}, t)$ is the value of $\boldsymbol{\lambda}$ for which the maximum in (18) is achieved. The differences $(\Delta F)_m(\boldsymbol{\lambda}_1, ..., \boldsymbol{\lambda}_m) = F_M(\boldsymbol{\lambda}(\boldsymbol{\mu}_m^+, t_m), t_m) - F_M(\boldsymbol{\lambda}(\boldsymbol{\mu}_m^-, t_m), t_m)$ summed over $m = 1, ..., M$ give the function $F_M(\boldsymbol{\lambda}_1, ..., \boldsymbol{\lambda}_M)$. To obtain its Legendre transform $H_M(\mathbf{m}_1, ..., \mathbf{m}_M)$ one must also have the derivatives $\mathbf{m}_m = \partial F_M / \partial \boldsymbol{\lambda}_m$. These are obtained from a suitable adjoint algorithm. The expectation

$$\mathbf{m}(t) = \int d\mathbf{x}\, \mathbf{M}(\mathbf{x}, t)\mathcal{A}(\mathbf{x}, t)\mathcal{P}(\mathbf{x}, t)$$

becomes within the linear ansatz (9)

$$\mathbf{m}(t) = \boldsymbol{\mu}(t) + \mathbf{C}_M(\boldsymbol{\mu}, t)\boldsymbol{\alpha}(t) \tag{19}$$

where now $\mathbf{C}_M(\boldsymbol{\mu}, t) = \langle [\mathbf{M}(t) - \boldsymbol{\mu}(t)][\mathbf{M}(t) - \boldsymbol{\mu}(t)]^\top \rangle_{\boldsymbol{\mu}(t), t}$ is the single-time covariance of $\mathbf{M}(t)$ within the PDF ansatz (14). The variable $\boldsymbol{\alpha}(t)$ satisfies the backward equation (13) between measurements. Requiring that the smoother $\mathbf{m}(t)$ be *continuous* at the measurement times yields the jump conditions:

$$\boldsymbol{\mu}_m^+ + \mathbf{C}_M(\boldsymbol{\mu}_m^+, t_m)\boldsymbol{\alpha}_m^+ = \boldsymbol{\mu}_m^- + \mathbf{C}_M(\boldsymbol{\mu}_m^-, t_m)\boldsymbol{\alpha}_m^-, \tag{20}$$

$m = 1, ..., M$. Given the value $\boldsymbol{\alpha}_m^+ = \boldsymbol{\alpha}(t_m+)$ from the backward integration, this linear equation can be solved for $\boldsymbol{\alpha}_m^-$. The latter then becomes the final condition $\boldsymbol{\alpha}(t_m-)$ for integration backward to the next measurement time. Finally, $\mathbf{m}_m = \mathbf{m}(t_m)$ for all $m = 1, ..., M$. See Eyink [7].

In addition to the convenient jump conditions, (17),(20), such exponential-PDF closures have other good properties. Closures of this type were first introduced in Levermore [11, 12] and successfully used in kinetic theory, where they are shown to preserve an H-theorem. Likewise, if the reference PDF $P_*(x, t)$ is an exact solution of the forward Kolmogorov equation,

then there is an H-theorem for the closure dynamics (12) with the relative entropy (18) as the Lyapunov function [see 10]. Indeed, the *rate of entropy change*, $\eta_M(\boldsymbol{\mu},t) = \boldsymbol{\lambda}^\top(\boldsymbol{\mu},t)\mathbf{V}(\boldsymbol{\mu},t) + \frac{\partial H}{\partial t}(\boldsymbol{\mu},t)$, is nonpositive for all $\boldsymbol{\mu}$ and zero only for $\boldsymbol{\mu} = \boldsymbol{\mu}_*(t)$ which minimizes $H_*(\boldsymbol{\mu},t)$. This is a consequence of the stability of the Fokker-Planck equation, which implies a relaxation in time of the distribution $P(\mathbf{x},t)$ for any initial data to a unique stationary PDF $P_s(\mathbf{x})$. For the closure equations (12) by the exponential-PDF ansatz, $\boldsymbol{\mu}_*(\infty)$ is then the unique, global fixed point solution.

We have now formulated a complete set of prescriptions for calculating $F_M(\boldsymbol{\lambda}_1,...,\boldsymbol{\lambda}_M)$ and $H_M(\mathbf{m}_1,...,\mathbf{m}_M)$ with the right exponential ansatz (14) and the left linear ansatz (9). Although no closed-form expression is available for either of these functions, we can employ the above algorithms to calculate them numerically to any desired accuracy. This suffices in order to carry out the necessary optimizations to obtain the smoother estimate.

The same circle of ideas may be applied to constructing closures of the KSP equations. In fact, assume that the closure variables $\mathbf{M}$ consist of the state variables $\mathbf{x}$ and their tensor products $\mathbf{x}\mathbf{x}^\top$, $\mathbf{M} := (\mathbf{x}, \mathbf{x}\mathbf{x}^\top)$, with mean values given by $\boldsymbol{\mu} = (\boldsymbol{\xi}, \boldsymbol{\Xi})$ for an exponential PDF ansatz. Let the linear coefficients in the left trial state be denoted as $\boldsymbol{\alpha} = (\mathbf{a}, \mathbf{A})$ and the exponential parameters in the right trial state be $\boldsymbol{\lambda} = (\boldsymbol{\ell}, \boldsymbol{\Lambda})$. Because the states evolve by the forward and backward Kolmogorov equations between measurements, the closure KSP equations are of the same form as those in (12),(13). Simple jump conditions for the thermodynamic field variables $\boldsymbol{\lambda} = (\boldsymbol{\ell}, \boldsymbol{\Lambda})$ in the forward equation may be read off directly from Ref [1] Equation 8:

$$\boldsymbol{\ell}_m^+ = \boldsymbol{\ell}_m^- + \mathbf{R}_m^{-1}\mathbf{y}_m, \tag{21}$$

$$\boldsymbol{\Lambda}_m^+ = \boldsymbol{\Lambda}_m^- - \frac{1}{2}\mathbf{R}_m^{-1}, \tag{22}$$

for $m = 1,...,M$. The jump conditions for the variables $\boldsymbol{\alpha} = (\mathbf{a}, \mathbf{A})$ in the backward equation follow from the requirement of continuity of $\mathbf{m}(t) = \int d\mathbf{x}\, \mathbf{M}(\mathbf{x},t)\mathcal{A}(\mathbf{x},t)\mathcal{P}(\mathbf{x},t)$ and are precisely the same as (20) above. In this approach, the evolution between measurements is approximated by closure, but the analysis at measurement times is performed exactly in the forward equation. The analysis in the backward equation is still only approximate, because the linear ansatz for the left trial state does not allow an exact implementation of the jump conditions. We emphasize that an approximation of the KSP equations could be formulated in this manner only if a *second-order closure* is employed, which includes among the closure variables the squares of the state variables in addition to those variables themselves.

The linear ansatz (9) is a source of worry in both of the approaches outlined in this section, i.e. closure of the KSP equations directly or closure

of the variational equations. If the magnitudes of the linear coefficients $|\alpha_i(t)|$, $i = 1, ..., R$ grow large enough, then it is possible that the left trial function $\mathcal{A}(\mathbf{x}, t) < 0$ for some $\mathbf{x}, t$. However, from the statistical interpretation of $\mathcal{A}$, it should be true that that $\mathcal{A}(\mathbf{x}, t) \geq 0$ for *all* $\mathbf{x}, t$. Therefore, we may expect to encounter difficulties with this approach when measured values are far from the unconditioned ensemble average $\overline{\mathbf{m}}(t)$, obtained by solving the closure equation (12) with the initial datum $\overline{\mathbf{m}}_0 = \int d\mathbf{x}\ \mathbf{M}(\mathbf{x})\mathcal{P}_0(\mathbf{x})$. In fact, we shall see in later sections that such difficulties do indeed materialize.

2.1.2 Double-Exponential Ansatz

Within the context of exponential PDF closures a particularly symmetric and attractive choice is to make the *double-exponential ansatz*:

$$\mathcal{P}(\mathbf{x}, t) = \exp\left[\boldsymbol{\beta}^\top \mathbf{M}(\mathbf{x}, t) - F(\boldsymbol{\beta}, t)\right] \mathcal{P}_*(\mathbf{x}, t) \tag{23}$$

for the right trial state and

$$\mathcal{A}(\mathbf{x}, t) = \exp\left[\boldsymbol{\alpha}^\top \mathbf{M}(\mathbf{x}, t) - (\Delta_{\boldsymbol{\alpha}} F)(\boldsymbol{\beta}, t)\right] \tag{24}$$

for the left trial state. Here $(\Delta_{\boldsymbol{\alpha}} F)(\boldsymbol{\beta}, t) := F(\boldsymbol{\alpha} + \boldsymbol{\beta}, t) - F(\boldsymbol{\beta}, t)$ so that the normalization constraint $\int d\mathbf{x}\ \mathcal{A}(\mathbf{x}, t)\mathcal{P}(\mathbf{x}, t) = 1$ is automatically satisfied. For small $\boldsymbol{\alpha}$, (24) coincides with the linear ansatz (9). However, this new ansatz for the left trial function is globally nonnegative and symmetric in form to the exponential for the right trial state. An even more attractive feature of this double exponential ansatz is that, within it, the Rayleigh-Ritz effective action of the closure variables $\mathbf{M}$ themselves may be calculated analytically in closed form. The Rayleigh-Ritz approximation to the effective action is determined from the double exponential ansatz according to the theorem stated at the beginning of section 2.1. The result is:

$$\begin{aligned} \Gamma_M[\mathbf{m}] &= \frac{1}{4}\int_{t_i}^{t_f} dt\ [\dot{\mathbf{m}}(t) - \mathbf{V}(\mathbf{m}, t)]^\top \mathbf{Q}^{-1}(\mathbf{m}, t)[\dot{\mathbf{m}}(t) - \mathbf{V}(\mathbf{m}, t)] \\ &+ H_M(\mathbf{m}(t_i), t_i) \end{aligned} \tag{25}$$

where $\mathbf{V}(\mathbf{m}, t)$ is the same as in the previous subsection and

$$Q_{ij}(\mathbf{m}, t) := \langle (\boldsymbol{\nabla}_{\mathbf{x}} M_i)^\top \mathbf{D} (\boldsymbol{\nabla}_{\mathbf{x}} M_j) \rangle_{\boldsymbol{\lambda}(\mathbf{m}, t)}. \tag{26}$$

See Eyink [7] for some details of the derivation.

The effective action (25) has the same form as the Onsager-Machlup action in statistical mechanics (see Onsager and Machlup [8], also Eyink

[20]). This form is familiar to meteorologists, since it it is similar to the cost function already used in the standard least-squares estimation or "weak constraint" approach to data assimilation [see 21]:

$$\begin{aligned} I[\mathbf{x}] &= \frac{1}{4}\int_{t_i}^{t_f} dt\ [\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x},t)]^\top \mathbf{D}^{-1}(\mathbf{x},t)[\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x},t)] \\ &+ \frac{1}{2}[\mathbf{x}(t_i) - \mathbf{x}_0]^\top \mathbf{C}_0^{-1}[\mathbf{x}(t_i) - \mathbf{x}_0]. \end{aligned} \tag{27}$$

This is the same action functional which appears for the PDF on histories[1]. Therefore, minimizing the cost function $I[\mathbf{x}]$ gives a maximum-likelihood estimate. However, despite their superficial resemblance, the cost functions (25) and (27) have quite different properties and interpretations. Foremost, the cost function $I[\mathbf{x}]$ takes as the unbiased estimate the solution of the deterministic dynamical equation

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x},t), \quad \mathbf{x}(t_i) = \mathbf{x}_0 \tag{28}$$

with no model error. On the other hand, our cost function $\Gamma_M[\mathbf{m}]$ takes as unbiased estimate an approximation of the ensemble average $\overline{\mathbf{m}}(t)$, obtained by solving the closure equation (12) with the initial datum $\overline{\mathbf{m}}_0$. Whereas the deterministic dynamics (28) is likely to have multiple attractors—fixed points, limit cycles, or chaotic attractors—the closure dynamics (12) should have a single global stable fixed point $\overline{\mathbf{m}}_*$, which corresponds to the average $\int d\mathbf{x}\ \mathbf{M}(\mathbf{x})\mathcal{P}_s(\mathbf{x})$ in the unique invariant measure $\mathcal{P}_s$ for the process. The latter is the "climate state" of the model. Thus, in the absence of any observations, our cost function correctly produces as its estimate the climate-average state. We may note also that the matrix $\mathbf{Q}(\mathbf{m},t)$ in our cost function is a *chaos-generated noise covariance*, like that calculated by Miller et al., section 4.d(2) in a Monte Carlo experiment for the Lorenz model. However, in our approach, the chaos-generated noise is calculated analytically within the closure via the formula (26). This noise may be much larger than the microscopic noise covariance $\mathbf{D}(\mathbf{x},t)$ appearing in the stochastic equation (16) and the bare action (27).

Despite the differences from the standard least-squares cost function, the similarity in form of our closure cost function (25) allows many of the standard methods for (27) to be carried over. For example, the inverse matrix $\mathbf{Q}^{-1}$ may be eliminated by introducing a maximization over an auxillary variable $\boldsymbol{\alpha}$. In fact, the latter just correspond to the parameters in the left trial state (24). Including the cost function for the observations, the total action to be extremized may then be written as

$$\Gamma_*[\boldsymbol{\alpha},\mathbf{m}] = \int_{t_i}^{t_f} dt\ \left\{\boldsymbol{\alpha}^\top[\dot{\mathbf{m}} - \mathbf{V}(\mathbf{m},t)] - \boldsymbol{\alpha}^\top\mathbf{Q}(\mathbf{m},t)\boldsymbol{\alpha}\right\} + H_M(\mathbf{m}(t_i),t_i)$$

$$+\frac{1}{2}\sum_{m=1}^{M}[\mathbf{m}(t_m)-\mathbf{y}_m]^\top\mathbf{R}_m^{-1}[\mathbf{m}(t_m)-\mathbf{y}_m]. \quad (29)$$

The Euler-Lagrange equations for this functional are

$$\dot{\mathbf{m}} = \mathbf{V}(\mathbf{m},t) + 2\mathbf{Q}(\mathbf{m},t)\boldsymbol{\alpha}, \quad (30)$$

$$\dot{\boldsymbol{\alpha}} + \left(\frac{\partial \mathbf{V}}{\partial \mathbf{m}}\right)^\top \boldsymbol{\alpha} + \frac{\partial}{\partial \mathbf{m}}\left(\boldsymbol{\alpha}^\top \mathbf{Q}\boldsymbol{\alpha}\right) = \sum_{m=1}^{M}\mathbf{R}_m^{-1}[\mathbf{m}(t_m)-\mathbf{y}_m]\delta(t-t_m). \quad (31)$$

Solving these equations with boundary values $\boldsymbol{\alpha}(t_i) = \boldsymbol{\lambda}(\mathbf{m}(t_i),t_i)$ and $\boldsymbol{\alpha}(t_f) = \mathbf{0}$ can give directly the optimal history, without the need of applying any explicit minimization algorithm. In contrast to the equations (12), (13) which arise from the left-linear ansatz, the equations (30), (31) above are fully coupled and must be solved as a two-time boundary-value problem. Furthermore, there are only jumps at measurements in the solution of the backward equation and none in solving the forward equation, which provides directly the smoother estimate $\mathbf{m}(t)$.

As in the previous subsection on the left-linear ansatz, we may also use the double-exponential ansatz to construct closures of the KSP equations directly. As before, we must assume that the closure variables $\mathbf{M}$ consist of the state variables $\mathbf{x}$ and their tensor products $\mathbf{x}\mathbf{x}^\top$, $\mathbf{M} := (\mathbf{x}, \mathbf{x}\mathbf{x}^\top)$, with mean values given by $\mathbf{m} = (\mathsf{x}, \mathbf{X})$ for a double exponential ansatz. Let the exponential parameters in the left trial state be denoted as $\boldsymbol{\alpha} = (\mathbf{a}, \mathbf{A})$ and those in the right trial state as $\boldsymbol{\beta} = (\mathbf{b}, \mathbf{B})$. The closure KSP equations are now of the same form as those in (30),(31). As in (30),(31), there are no jumps at measurement times in the equations for $\mathbf{m} = (\mathsf{x}, \mathbf{X})$. On the other hand, as follows from [1], equation (10), there are simple jump conditions for the adjoint variables $\boldsymbol{\alpha} = (\mathbf{a}, \mathbf{A})$,

$$\mathbf{a}_m^- = \mathbf{a}_m^+ + \mathbf{R}_m^{-1}\mathbf{y}_m, \quad (32)$$

$$\mathbf{A}_m^- = \mathbf{A}_m^+ - \frac{1}{2}\mathbf{R}_m^{-1}, \quad (33)$$

for $m = 1, ..., M$. In this approach, the evolution between measurements is approximated by closure, but the analysis at measurement times is performed exactly. As before, the price paid for this is that a *second-order closure* must be employed.

2.2 Empirical Ensemble Methods

An alternative method of approximating the dynamical evolution is by the use of empirical ensembles of solution histories, $\{\mathbf{x}^{(n)}(t) : t \in [t_i, t_f]\}]$ for $n = 1, ..., N$. In particular, such empirical ensemble or Monte Carlo methods

have been advocated and developed by Evensen. In the case of zero model error, when $\mathbf{D} \equiv \mathbf{0}$, the forward Kolmogorov or Liouville equation

$$\partial_t \mathcal{P}(\mathbf{x}, t) = -\boldsymbol{\nabla}_{\mathbf{x}} \cdot [\mathbf{f}(\mathbf{x}, t) \mathcal{P}(\mathbf{x}, t)] \tag{34}$$

is hyperbolic and the ensemble methods correspond to solution by the method of characteristics. In fact, the solution trajectories $S^t \mathbf{x}_0$ of (16) are just the characteristics of (34) and

$$\mathcal{P}(\mathbf{x}, t) = \mathcal{P}_0(S^{-t}\mathbf{x}) \left| \frac{\partial (S^{-t}\mathbf{x})}{\partial \mathbf{x}} \right| \tag{35}$$

is the solution of (34) written in terms of the initial distribution and backward characteristics. For averages of moment functions $\mathbf{M}(\mathbf{x}, t)$ with respect to $\mathcal{P}(\mathbf{x}, t)$ this yields the formula

$$\langle \mathbf{M}(\mathbf{x}, t) \rangle_t = \int d\mathbf{x}_0 \ \mathcal{P}_0(\mathbf{x}_0) \mathbf{M}(S^t \mathbf{x}_0, t). \tag{36}$$

The latter average over initial data may be approximated by an N-sample empirical ensemble average

$$\langle \mathbf{M}(\mathbf{x}, t) \rangle_t \approx \frac{1}{N} \sum_{n=1}^{N} \mathbf{M}(\mathbf{x}^{(n)}(t), t). \tag{37}$$

where $\mathbf{x}^{(n)}(t) = S^t \mathbf{x}_0^{(n)}$ and $\mathbf{x}_0^{(n)}$, $n = 1, ..., N$ are independently sampled from the distribution $\mathcal{P}_0$. For low-order moment functions $\mathbf{M}(\mathbf{x}, t)$ one may expect to achieve convergence of (37) for sample sizes N of reasonable magnitude. In the case where $\mathbf{D} \neq \mathbf{0}$ and the forward Kolmogorov equation is parabolic, then one must include in (36) also an average over realizations of the random noise. This is the analogue of the method of characteristics for the parabolic Fokker-Planck equation. In this case the N sample histories in (37) must also be constructed with temporal sequences of model errors drawn independently from the white-noise distribution. Such methods may be combined with approximate analysis methods—notably the linear analysis—to generate approximations to N-sample ensembles conditioned upon measurements. See Evensen [13], van Leeuwen and Evensen [14], Burgers et al. [15], Evensen and van Leeuwen [16] and references in [17]. Such approximations shall be the primary object of comparison for the moment-closure closure schemes advocated in this work.

3 Closures for the Model Problem

We consider as our sample problem the following stochastically forced double-well system

$$\dot{x}(t) = f(x(t)) + \kappa \eta(t) \tag{38}$$

where

$$f(x) = 4x(1 - x^2), \tag{39}$$

$\eta(t)$ is white-noise with zero mean and covariance $\langle \eta(t)\eta(t')\rangle = \delta(t-t')$, and $\kappa = 0.5$. As for any nonlinear dynamical system, the statistical moments $\mu_n = \langle x^n \rangle,\ n = 1, 2, 3, ...$ do not satisfy closed equations at any finite order. Instead, there is an infinite hierarchy of moment equations. For our simple model, the equations of the hierarchy are

$$\dot{\mu}_1 = 4\mu_1 - 4\mu_3, \tag{40}$$

$$\dot{\mu}_2 = 8\mu_2 - 8\mu_4 + \kappa^2,\ \ldots \tag{41}$$

and so forth for higher-order moments. Each equation for a moment of a given order involves moments of higher order. Moment-closure methods attempt to solve the hierarchy truncated at a given order by making a model or *closure* of the unknown higher-order moments in terms of the retained lower-order ones. Most commonly, a *first-order closure* is made, in which only the average of the dynamical equation is closed. This corresponds to keeping only (40) in our hierarchy. Occasionally, one considers *second-order closures* in which an equation also for 2nd-order moments of the state variables are considered, like our (41). For spatially-extended systems with many degrees of freedom, such second-order closures are already at the limit of practicality, because the number of variables to be integrated is the square of that for a first-order closure. Closures of third and higher order are therefore rarely considered.

We shall apply such closure methods to the double-well model with either $\mathbf{M}(x) = (x)$ (first-order closure) or $\mathbf{M}(x) = (x, x^2)$ (second-order closure). We discuss each of these in turn.

3.1 First-Order Closure

In the first-order closure, we choose our reference PDF $\mathcal{P}_*(x)$, approximating $\mathcal{P}_s(x)$, as the distribution of a *surrogate variable* X_* constructed in the following way. First, let $X_\pm$ be normal random variables centered at ± 1 with the PDF's

$$\mathcal{P}_\pm(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x \mp 1)^2}{2\sigma^2}\right) \tag{42}$$

where $\sigma^2 = \kappa^2/16$. Because $U''(\pm 1) = 8$, these random variables are good surrogates for the fluctuations of the solution $x(t)$ of (38) within the wells at $x = \pm 1$. We finally define X_* to be the random variable which takes the same value as X_+ with probability $\frac{1}{2}$ and the same value as X_- also with probability $\frac{1}{2}$. Essentially, we toss a fair coin in order to determine

in which well of the potential U the surrogate X_* lies. This method allows us to construct a highly non-Gaussian, bimodal reference PDF $\mathcal{P}_*(x)$ by randomly selecting from normal random variables, themselves having Gaussian distributions. The same method should be readily applicable to more realistic, spatially extended systems.

In the present simple example, using this choice of $\mathcal{P}_*$, it is straightforward to calculate the cumulant-generating function or "free-energy" $F_M(\lambda_1) = \log\left(\int dx\ e^{\lambda_1 x}\mathcal{P}_*(x)\right)$ in closed form:

$$F_M(\lambda_1) = \frac{1}{2}\sigma^2\lambda_1^2 + \log\cosh(\lambda_1). \tag{43}$$

A closed form expression for F_M is not required for our method to be applicable: it is also possible to use a numerical approximation of the normalization integral $\mathcal{N}_M(\lambda_1)$ in order to determine $F_M(\lambda_1)$. However, it is convenient here to use the analytical expression (43). All of the moments μ_n required may be obtained by expansion into cumulants C_n, and the latter are obtained in turn as n-fold derivatives of $F_M(\lambda_1)$. Thus, $\mu_1 = C_1, \mu_2 = C_2 + C_1^2$ and $\mu_3 = C_3 + 3C_1C_2 + C_1^3$ with

$$\begin{aligned} C_1 &= F_M'(\lambda_1) = \sigma^2\lambda_1 + \tanh\lambda_1 \\ C_2 &= F_M''(\lambda_1) = \sigma^2 + \mathrm{sech}^2\lambda_1 \\ C_3 &= F_M'''(\lambda_1) = -2\mathrm{sech}^2\lambda_1\tanh\lambda_1. \end{aligned} \tag{44}$$

In particular, the moments μ_1, μ_3 appearing in $V_1 = 4\mu_1 - 4\mu_3$ in (40) are both obtained in terms of λ_1, yielding a closure of that equation. Integrating the resulting closure dynamics yields results that are virtually indistinguishable from those using the first-order closure discussed in Eyink and Restrepo [22], which was constructed with a somewhat different PDF ansatz.

For our simple model, the closure equation in terms of the thermodynamic field λ_1 may be written out explicitly. It is

$$(\sigma^2 + \mathrm{sech}^2\lambda_1)\frac{d\lambda_1}{dt} = -4\sigma^2(2+3\sigma^2)\lambda_1 - 4\sigma^6\lambda_1^3 - 12\sigma^2(1+\sigma^2\lambda_1^2)\tanh\lambda_1 \tag{45}$$

The relative entropy likewise can be given as

$$H_M(\lambda_1) = \frac{1}{2}\sigma^2\lambda_1 + \lambda_1\tanh\lambda_1 - \log\cosh\lambda_1. \tag{46}$$

In the present case one can verify that $\eta_M(\lambda_1) \leq 0$ for all λ_1 and $= 0$ only for $\lambda_1 = 0$, so that H_M is a Lyapunov function. In Eyink and Levermore [10] it is shown that this is true in general for exponential PDF closures, not only when the reference PDF $\mathcal{P}_*(x)$ is taken to be the exact stationary PDF $\mathcal{P}_s(x)$, but also if it is a sufficiently good approximation. Because of

this H-theorem for the dynamics (45), the origin 0, which is an unstable fixed point for the deterministic dynamics $\dot{x} = f(x)$, is the global stable fixed point of the closure dynamics.

For the Rayleigh-Ritz calculation of effective actions, we need also one other closure quantity, the "effective noise covariance" $\mathbf{Q}$. According to the general formula (26), in our model $\mathbf{Q}(\boldsymbol{\mu}) = 8\sigma^2\langle \mathbf{M}'(x)\mathbf{M}'(x)^\top\rangle_{\boldsymbol{\mu}}$ with $\mathbf{M}'(x) = d\mathbf{M}(x)/dx$. Thus, for the first-order closure,

$$Q(\mu_1) = 8\sigma^2. \tag{47}$$

Note that Q is the same as the bare diffusion constant $\kappa^2/2$. This will be true for the first-order closure of any stochastic dynamical system with purely additive noise.

3.2 Second-Order Closure

We take here the same reference PDF $\mathcal{P}_*(x)$ as for the first-order closure. We must then find $F_M(\lambda_1, \lambda_2) = \log\left(\int dx\; e^{\lambda_1 x + \lambda_2 x^2}\mathcal{P}_*(x)\right)$. Because the reference PDF is a linear combination of Gaussians, it is easy to calculate this also in closed form:

$$F_M(\lambda_1, \lambda_2) = \frac{\sigma^2\lambda_1^2 + 2\lambda_2}{2(1 - 2\lambda_2\sigma^2)} - \frac{1}{2}\log(1 - 2\lambda_2\sigma^2) + \log\cosh\left[\frac{\lambda_1}{1 - 2\lambda_2\sigma^2}\right] \tag{48}$$

Note that F_M is only finite when $\lambda_2 < \frac{1}{2\sigma^2}$. This constraint is obviously necessary because the tails of the reference PDF for $|x| \gg 1$ have the asymptotic form $e^{-x^2/2\sigma^2}$. Thus, the normalization integral will only be finite if the inequality is satisfied.

As before, we obtain all of the moments in the closure vector

$$\mathbf{V}(\boldsymbol{\mu}) = \begin{bmatrix} 4\mu_1 - 4\mu_3 \\ 8\mu_2 - 8\mu_4 + \kappa^2 \end{bmatrix} \tag{49}$$

in terms of λ_1, λ_2 (and thence μ_1, μ_2) by expanding them in terms of cumulants $C_n = \partial^n F_M/\partial\lambda_1^n$. This algorithm generates the values of the components V_1, V_2 easily for numerical computation, but analytical expressions are a little complicated to write out in general and will not be given here. However, for $\lambda_1 = 0$ it is easy to show that $V_1 = 0$ for all λ_2 and

$$V_2 = 16\sigma^2 + 8\sigma^2\xi - (24\sigma^4 - 8)\xi^2 - 48\sigma^2\xi^3 - 8\xi^4 \tag{50}$$

with $\xi = 1/(1 - 2\lambda_2\sigma^2)$. The equation $V_2 = 0$ has precisely one positive root ξ^s. For $\kappa = 0.5$, it is $\xi^s \approx 0.978$, which corresponds to $\lambda_2^s \approx -0.724$. Thus, the closure dynamics has a fixed point at $(\lambda_1^s, \lambda_2^s) \approx (0, -0.724)$

and numerical investigation shows that it is unique and stable. The reference PDF for our closure should, in fact, really be taken to be $\widetilde{\mathcal{P}}_*(x) = e^{\boldsymbol{\lambda}^{s\top}\mathbf{M}(x)}\mathcal{P}_*(x)/\mathcal{N}(\boldsymbol{\lambda}^s)$. The latter is the best approximation to the stationary PDF $\mathcal{P}_s(x)$ within our two-parameter ansatz because it satisfies the stationarity condition

$$\int dx\ M_i(x)\,\hat{L}\widetilde{\mathcal{P}}_*(x) = 0, \quad i = 1,2 \tag{51}$$

Here $\hat{L} = -\frac{\partial}{\partial x}\,(f(x)\cdot) + \frac{\kappa^2}{2}\frac{\partial^2}{\partial x^2}$ is the Fokker-Planck operator for the stochastic equation (38). However, taking $\widetilde{\mathcal{P}}_*(x)$ as the reference PDF amounts just to making some trivial shifts in variables, e.g. $\widetilde{F}_M(\boldsymbol{\lambda}) = F_M(\boldsymbol{\lambda}+\boldsymbol{\lambda}^s) - F_M(\boldsymbol{\lambda}^s)$. The only essential point is that it is the corresponding entropy $\widetilde{H}_M(\boldsymbol{\mu}) = H_M(\boldsymbol{\mu}) - H_M(\boldsymbol{\mu}^s) - \boldsymbol{\lambda}^{s\top}(\boldsymbol{\mu}-\boldsymbol{\mu}^s)$ with $\boldsymbol{\mu}^s \approx (0, 0.972)$ which is the Lyapunov function for the closure dynamics, and not $H_M(\boldsymbol{\mu})$.

Finally, we need to give the effective noise covariance in the second-order closure:

$$\mathbf{Q}(\mu_1, \mu_2) = 8\sigma^2 \begin{bmatrix} 1 & 2\mu_1 \\ 2\mu_1 & 4\mu_2 \end{bmatrix} \tag{52}$$

This is easily verified from the formula in the previous subsection.

4 Results for the Model Problem

We shall now consider results of data assimilation experiments for the model (38). We use the same sample history that appeared in Figure 1 of Miller et al. 1994 and in Figure 1 of [1]. For ease of comparison, the five sets of "measurements" we consider shall also be the same as in [1], and designated in the same way, as datasets A-E. We defer details of the numerical implementation of all the estimation methods to [24].

4.1 Comparison of Filtering Schemes

To begin, we compare three filtering schemes. The main new method we consider is a second-order closure of the Kushner-Stratonovich (KS) filter equations (KS-2E), using the exponential PDF ansatz of section 3.2. For comparison, we consider the Ensemble Kalman Filter (EnKF) of Evensen [13], Burgers et al. [15] discussed in [24] section 3.4.1, and finally the exact KS filter, discussed in [1]. Since filters are suboptimal estimators and not the main subject of this paper, we shall compare the above three methods for the single dataset A.

In Figure 1 we plot the mean history and also the mean history plus or minus the standard deviation for each of the three filtering methods. In

Figure 1(a) are plotted the exact conditional statistics from the KS filter. In Figure 1(b) we plot the results of KS-2E closure, and in Figure 1(c) the results for EnKF. Obviously, KS and KS-2E agree quite well, except that the variance predicted by the closure approximation tends to be somewhat larger than the true conditional variance from KS. Apart from this minor discrepancy, the second-order exponential PDF closure of the KS filter seems to be remarkably good. The EnKF method, however, has a serious discrepancy when compared with the previous two methods. Whereas the transition in the optimal KS filter—and also KS-2E—begins at time $t = 3$ and is completed at time $t = 5$, for EnKF there is a lag in the transition. For EnKF the transition is just barely indicated at time $t = 4$, really begins at $t = 5$ and is only completed at time $t = 6$. This defect of EnKF has already been observed in Miller et al. [25]. The reason for this failure to track the transition properly is the linear analysis employed in EnKF. For a discussion of that, see [1] section 3.2 and Evensen [13], Burgers et al. [15]. Therefore, more recent ensemble Kalman filtering schemes [17] employing the same linear analysis are likely to show the same deficiency. On the other hand, novel ensemble filters introduced in [26], which employ a proper conditional analysis rather than a Kalman approximation, perform much better in this situation.

4.2 Comparison of Smoothing Schemes

Now we shall compare six smoothing algorithms, which we denote by acronyms KSP, MFV, MFV-1LL, MFV-1DE, MFV-2DE, and EnKS.

The first is the KSP algorithm[1]. It implements the exact analysis and the exact dynamics (except for negligible discretization errors in solving the PDE's). Hence, this algorithm gives the primary results for comparison with all the suboptimal estimators. We also consider estimators that employ the mean-field variational analysis. When implemented with the exact dynamics via the Kolmogorov equations, as already discussed in [1], we denote the algorithm just as MFV. Our main object here is to discuss the mean-field estimators within moment-closure approximations to the time-evolution. We discuss several such approximations: MFV-1LL, using the first-order closure of section III.A and the left-linear ansatz of section II.A.1; MFV-1DE, using the first-order closure and the double-exponential ansatz of section II.A.2; and, finally, MFV-2DE, using the second-order closure of section III.B and the double-exponential ansatz. These represent a series of approximations of increasing sophistication and, hopefully, accuracy. We do not present any results for a closure of the MFV equations using the second-order closure and left-linear ansatz, because such an approximation fails to give usable results.

Finally, we consider a suboptimal estimator using a linear analysis and

also an ensemble approximation to the dynamical evolution as discussed in section 2.2 of this paper. The method we employ is the Ensemble Kalman Smoother (EnKS) of Evensen and van Leeuwen [16].

4.2.1 Dataset A

We first compare the methods for the dataset A, with results shown in Figure 2. This is a case with measurement error variance 0.04 or 20% rms error, where the MFV and exact KSP values of the conditional mean and variance are here very close [1]. The most important fact learned from Figure 2 is that *all* of the approximate MFV methods are successful in tracking the transition in dataset A occuring at times $t = 3 - 5$ but that the EnKS method lags by one measurement in making the transition. This is the same failure that was observed in the EnKF method in Section 4.1 above and it is not cured by the backward-in-time adjustments made in EnKS. This is again a failure of the linear analysis performed in EnKF and EnKS, and it is cured by the mean-field conditional analysis made in the MFV methods. While all of the MFV methods are successful at the most basic level of tracking the transition, we see a quantitative improvement, as expected, in going from MFV-1LL, to MFV-1DE, to MFV-2DE. The worst of these approximations, MFV-1LL, greatly overpredicts the conditional variance $\sigma_*^2(t)$. Thus, it would lead to a lower confidence in the conditional mean $x_*(t)$ as an accurate estimate than, in fact, is warranted. The MFV-1DE method, although based upon the same first-order closure as MFV-1LL, nevertheless gives a much more accurate value of the variance. The mean histories in the two methods MFV-1LL and MFV-1DE are quite similar, and are good approximations to the exact $x_*(t)$ except before the first measurement and after the last one, where they relax toward zero a little too quickly. This is a deficiency of the first-order closure in the double-well problem, that it overpredicts the relaxation back to equilibrium (i.e. to the climate state). This also accounts for the somewhat too large variances in the MFV-1DE method, before the first measurement and after the last one. Indeed, in equilibrium with equal weights of the peaks in the PDF at the two values $x = \pm 1$, the variance is $O(1)$. The MFV-2DE method cures this deficiency, because the second-order closure relaxes to equilibrium at a slower rate, in better agreement with the exact statistical evolution by the forward-backward Kolmogorov equations. In fact, the MFV-2DE closure leads to results for dataset A almost indistinguishable from those of the full MFV method, which solves the Kolmogorov equations.

4.2.2 Dataset B

In Figure 3 we compare the six methods for the dataset B. This dataset has the same measured values as dataset A, but now it is assumed that the observation error variance was 0.16 or 40% rms error. The MFV method gives now a less accurate approximation to the optimal KSP results (for reasons discussed in [1]). The general trends for the other approximations are similar to those for dataset A in several respects. Again, all of the approximate MFV methods continue to track the transition, even with the larger assumed observation error. However, the EnKS method gives an estimated history which lags even worse than for dataset A, the transition now being completed only at $t = 7$. The variance $\sigma_*^2(t)$ calculated from MFV-1LL is again too large, but those for MFV-1DE and MFV-2DE closure are about the same as for MFV itself. The results for MFV-1DE are now much closer to those for MFV-2DE than to MFV-1LL, except that before the first measurement and after the last one the MFV-1DE history relaxes a little fast and the variance grows a little too fast. These problems are corrected by the MFV-2DE closure, which gives results very similar to MFV. All of the MFV methods show the important feature of *statistical stability*: the results for dataset B are quite similar to those for dataset A, except, of course, the calculated variances are larger. The EnKS method is considerably less stable.

4.2.3 Dataset C

Figure 4 shows the results of the six methods for the dataset C. This is another case with measurement error variance 0.04 or 20%, but with a different set of "measured" values which happen to have smaller magnitudes than for dataset A. The suboptimal MFV method now gives a slightly less accurate approximation to KSP, but still quite good. All of the MFV closure methods also track the transition well. In this case, EnKS tracks the transition better than in any of the previous cases, but it still lags a bit behind the measurements. As before, MFV-1LL gives variances too large. In this case, MFV-1DE and MFV-2DE give quite similar results, although the MFV-2DE closure is little closer to MFV. We see that all of the MFV methods possess statistical stability also in the sense that with the same history but different "measured" values (i.e. a different realization of measurement errors) the calculated mean history and variance do not change greatly. The EnKS method is less stable.

4.2.4 Dataset D

The case of dataset D presented in Figure 5 is interesting because here the optimal estimate for the "measured" data, with error variance 0.09

or 30% rms error, does *not* exactly track the transition in the "actual" time-history. This means that with such poor measurements, a correct estimation method should also not track the transition. Instead, a correct method should show a quick crossover at time $t = 5$, but should then reverse back to the other well and give thereafter large variances, indicating an inability to predict with accuracy the state of the system. The MFV method catches the crossover at time $t = 5$ but gives afterward variances which are too small, projecting too much confidence in the estimate at those times. All of the MFV closure methods also show the crossover at time $t = 5$, but the EnKS method misses it entirely. In fact, the EnKS method gives essentially no hint at all of the transition which "actually" occurred. The MFV closure methods all show larger variances than MFV, especially MFV-1LL. The MFV-1DE and MFV-2DE are quite similar to each other and, by happenstance, give somewhat better agreement with the exact KSP results than does MVF itself. This is a result of "cancellation of errors" between the approximate analysis and the approximate evolution in those methods.

4.2.5 Dataset E

The final case is dataset E, the results for which are compared in Figure 6. This case combines features of datasets C and D: the magnitudes of "measured" values are small and the measurement error variance is large, now 0.36, or 60% rms error. In this case the exact conditional statistics given by KSP show almost no evidence of the transition in the "actual" history except for a slight bend downward in $x_*(t)$ at time $t = 4$, but the variance $\sigma_*^2(t)$ thereafter becomes large, indicating great uncertainty in the actual state of the system. With these poor measurements, there is an intrinsic loss of predictability. MFV performs the worst for this case. Its estimated history $x_*(t)$ is too small in magnitude, although it shows the slight bend at time $t = 4$. The variance $\sigma_*^2(t)$ given by MFV is too large before the bend and too small afterward. The MFV closure methods and the EnKS method give quite similar results for this case. The estimated history $x_*(t)$ for all these methods is too small in magnitude and, while it shows the bend at time $t = 4$, this is actually too pronounced, especially for MFV-1LL. All of the approximation methods tend to overpredict the variance before the bend and underpredict it afterward. Hence, none of the approximations is particularly successful for this case of highly inaccurate measurements. On the other hand, none of them is unacceptably bad either.

4.2.6 General Conclusions

We may draw some general conclusions based upon the above comparisons. We see that of all the approximation methods, EnKS gives the poorest results for the conditional history and variance. The MFV-1LL method, while it gives good results for the history, tends to greatly overestimate the variance. The best methods are consistently MFV-1DE and MFV-2DE. Reassuringly, MFV-2DE is always closer to MVF, so that adding more moments to the closure produces a better result, as one would hope. However, the MFV-1DE method tends to give nearly as good results as MFV-2DE. Therefore, there may be a true advantage in using the mean-field variational methods, since they allow the use of *first-order* closure methods with, possibly, little sacrifice in the accuracy of the results.

5 Discussion and Summary

In this work we have considered two methods of approximating the evolution of probability distributions, for the purpose of calculating conditional statistics: moment-closure methods and empirical-ensemble methods. Several versions of the former were considered. It was found to be particularly advantageous to construct the moment-closure to preserve an exact H-theorem for the Markov evolution, in terms of the relative entropy. We have compared these methods for a simple model system with a non-Gaussian, bimodal steady-state distribution. Applied within the mean-field variational method [1], first-order closures were nearly as accurate as second-order. Approximation of the KSP equations for the exact conditional statistics instead requires a closure of at least second-order. In the simple test problem, the mean-field variational closure methods consistently out-performed the empirical-ensemble or Monte Carlo methods. Approximate results for both the conditional mean and variance were consistently more accurate for the closure methods. This is more likely due to the over-simplified linear analysis employed in current ensemble estimation schemes, rather than to an intrinsic deficiency in approximating the dynamical evolution of system statistics by N-sample ensembles. At least for the filtering problem, newer ensemble methods that employ a correct conditional analysis perform much better [26]. It is therefore a promising avenue for future work to develop corresponding ensemble smoothers which go beyond the linear, Kalman approximation and calculate proper conditional statistics.

Let us make a few final remarks comparing moment-closure and empirical-ensemble methods in general.

For ease of application, ensemble methods are rather superior to closure methods. They are, in fact, cookbook methods that can be applied in

a rote fashion. Closure methods, however, require both insight into the statistics and dynamics of the system and also some skill in constructing closures that are simultaneously physically accurate and computationally tractable. On the other hand, this apparent disadvantage may in fact be an advantage in disguise. Since successful closures provide theoretical understanding, they may assist in describing other distinct but related problems. The ensemble method, in contrast, must be repeated *de novo* for each application. On the other hand, careful study of the results of an ensemble calculation may suggest theoretical ideas useful in constructing accurate closure approximations.

Closure methods probably have the potential to achieve the greatest computational economy in calculating conditional statistics. This is especially true for the large-scale, spatially-extended systems that appear in geophysical applications. In Alexander et al. [27] the methods presented in this paper are applied to a stochastic advection-diffusion equation, for the problem of tracking the past history of a passive scalar contaminant in a random flow. While a standard first-order closure yields there forward-backward evolution equations which are PDE's (advection of the scalar by the mean-flow field plus eddy-diffusion), more ingenious closures can be devised which build in more information about the realizations of the problem. In that case, the estimation problem can be reduced to solving ODE's in just a few number of variables. The Rayleigh-Ritz method is generally one which rewards physical understanding of the problem to be solved. If one can make a good, insightful guess of the solution, then very significant savings in computing effort may be obtained.

We view the different methods of approximating the evolution —moment-closure and empirical ensembles —not as competitive but instead as co-operative. Neither method is exact. Furthermore, methods to calculate the exact conditional statistics are generally intractable to apply to the spatially-extended, strongly nonlinear systems of real interest. Therefore, the only way to assess the accuracy of approximation schemes is by inter-comparison. Because moment closure and empirical ensemble methods make rather distinct and complementary approximations, the agreement of their results would lend considerable credibility to both.

6 Acknowledgments

We wish to thank C. D. Levermore and M. Ghil in particular for much useful advice on the subject of this paper and, also, the anonymous reviewers of a previous version of the paper. This work, LAUR 02-3057, was carried out in part at Los Alamos National Laboratory under the auspices of the Department of Energy and supported by LDRD-ER 2000047. We also

received support from NSF/ITR, Grant DMS-0113649 (GLE,JMR), as well as from NASA, Goddard Space Flight Center, Grant NAG5-11163 (JMR).

References

[1] G. L. Eyink, J. M. Restrepo, and F. J. Alexander. A mean field approximation in data assimilation for nonlinear dynamics, *submitted*, 2003.

[2] C. E. Leith. Theoretical skill of monte carlo forecasts. *Mon Wea. Rev.*, 102:409–418, 1974.

[3] H. J. Kushner. On the differential equations satisfied by conditional probability densities of markov processes, with applications. *J. SIAM Control, Ser.A*, 2:106–119, 1962.

[4] H. J. Kushner. Dynamical equations for optimal nonlinear filtering. *J. Diff. Eq.*, 3:179–190, 1967.

[5] H. J. Kushner. Approximation to optimal nonlinear filters. *IEEE Trans. Auto. Contr.*, 12:546–556, 1967.

[6] M. Ehrendorfer. The liouville equation and its potential usefulness for the prediction of forecast skill. part i. theory. *Mon Wea. Rev.*, 122: 703–713, 1994.

[7] G. L. Eyink. A variational formulation of optimal nonlinear estimation, *submitted*, 2002.

[8] L. Onsager and S. Machlup. Fluctuations and irreversible processes. *Physical Review*, 91:1505–1512, 1953.

[9] R. Kleeman. Measuring dynamical prediction utility using relative entropy (to appear). *J. Atmos. Sci.*, 2001.

[10] G. L. Eyink and C. D. Levermore. Entropy-based closure of nonlinear stochastic dynamics, *in preparation*, 2002.

[11] C. D. Levermore. Moment closure hierarchies for kinetic theories. *J. Stat. Phys.*, 83:1021–1065, 1996.

[12] C. D. Levermore. Entropy-based moment closures for kinetic equations. *Transp. Theor. Stat. Phys.*, 26:591–606, 1997.

[13] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *J. Geophys. Res.*, 99 (C5):10 143–10 162, 1994.

[14] P. J. van Leeuwen and G. Evensen. Data assimilation and inverse methods in terms of a probabilistic formulation. *Mon. Wea. Rev.*, 124:2898–2913, 1996.

[15] G. Burgers, P. J. van Leeuwen, and G. Evensen. Analysis scheme in the ensemble kalman filter. *Mon. Wea. Rev.*, 126:1719–1724, 1998.

[16] G. Evensen and P. J. van Leeuwen. An ensemble kalman smoother for nonlinear dynamics. *Mon. Weather Rev.*, 128:1852–1867, 2000.

[17] M. K. Tippett, J. L. Anderson, C. H. Bishop, T. H. Hamill, and J. S. Whitaker. Ensemble square-root filters. *Monthly Weather Review, to appear*, 2003.

[18] H. Risken. *The Fokker-Planck Equation.* Springer-Verlag, New York, 1984.

[19] R. Balian and M. Vénéroni. Time-dependent variational principle for the expectation value of an observable: mean-field applications. *Ann. Phys.*, 164:334–410, 1985.

[20] G. L. Eyink. Dissipation and large thermodynamic fluctuations. *J. Stat. Phys.*, 61:533–572, 1990.

[21] Y. Sasaki. An objective analysis based on the variational method. *J. Meteor. Soc. Japan*, 36:77–88, 1958.

[22] G. L. Eyink and J. R. Restrepo. Most probable histories for nonlinear dynamics: tracking climate transitions. *J. Stat. Phys.*, 101:459–472, 2000.

[23] R. N. Miller, M. Ghil, and P. Gauthiez. Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmos. Sci.*, 51:1037–1056, 1994.

[24] G. L. Eyink, J. M. Restrepo, and F. J. Alexander. Algorithms for the practical implementation of a statistical-mechanical approach to data assimilation, *submitted*, 2003.

[25] R. N. Miller, Jr. E. F. Carter, and S. T. Blue. Data assimilation into nonlinear stochastic models. *Tellus*, 51A:167–194, 1999.

[26] S. Kim, G. L. Eyink, J. M. Restrepo, F. J. Alexander, and G. Johnson. Ensemble filtering for nonlinear dynamics. *Mon Wea. Rev., to appear*, 2003.

[27] F. J. Alexander, G. L. Eyink, K. Pao, and D. Tartakovsky. Calculating contaminant histories in a model of random advection: a statistical approach to inverse modeling, *in preparation*, 2002.

Figure Captions

1. Dataset A (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KS filter, (b) KS-2E closure, (c) EnKF.

2. Dataset A (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

3. Dataset B (40 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

4. Dataset C (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

5. Dataset D (30 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

6. Dataset E (60 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

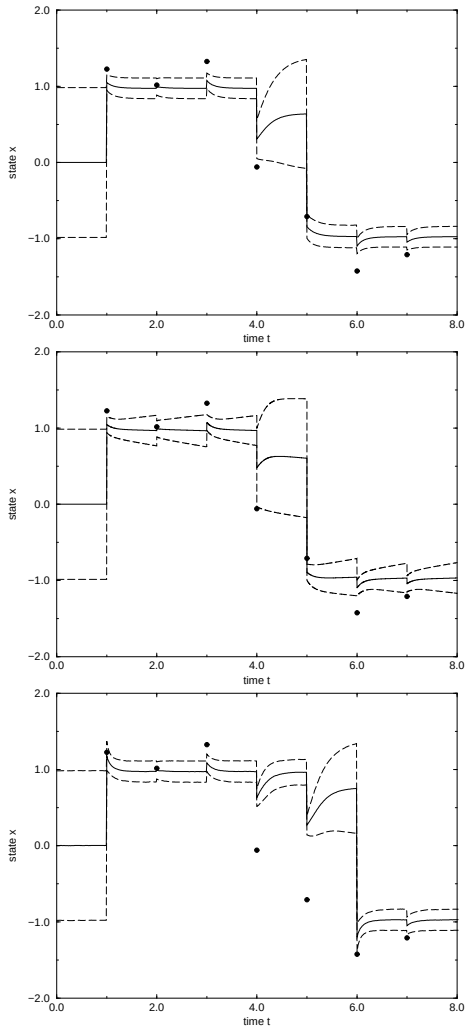

Figure 1: Dataset A (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. Clockwise, (a) KS filter, (b) KS-2E closure, (c) EnKF.

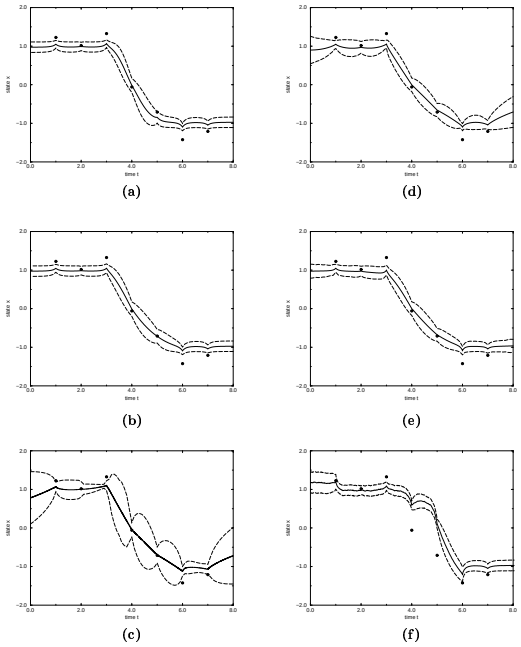

Figure 2: Dataset A (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

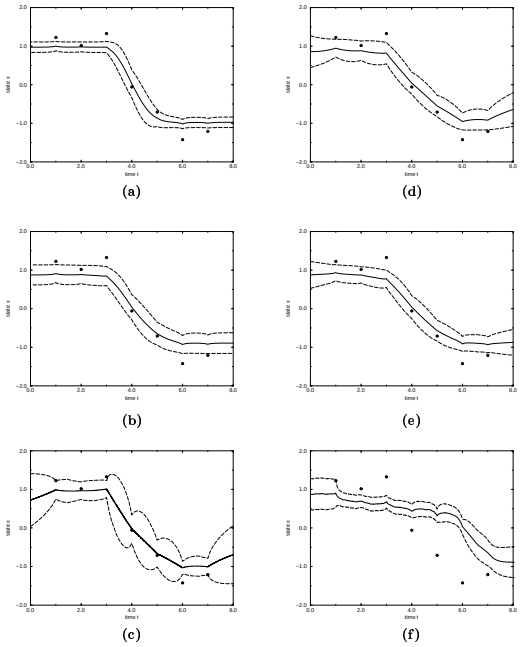

Figure 3: Dataset B (40 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

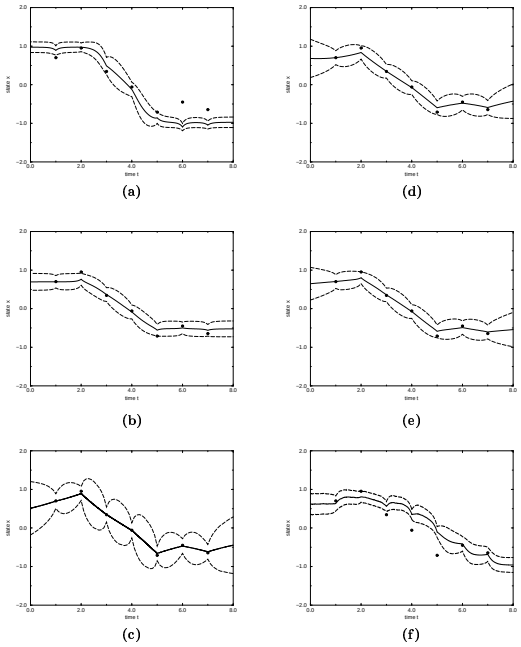

Figure 4: Dataset C (20 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

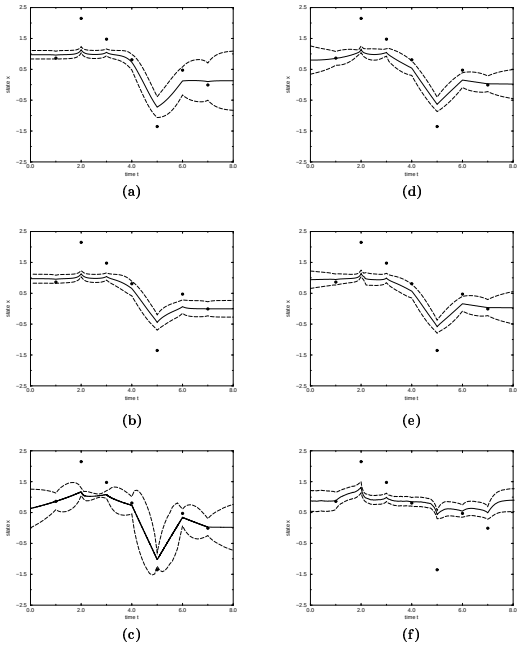

Figure 5: Dataset D (30 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.

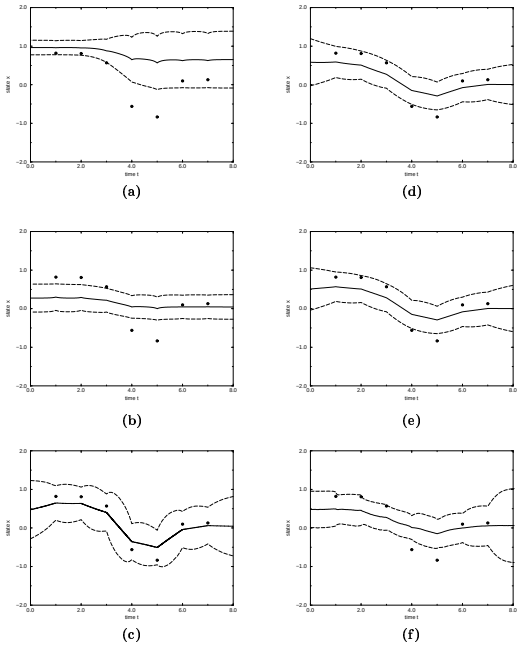

Figure 6: Dataset E (60 % rms error) represented by filled circles. Solid line: mean, dashed line: mean ± standard deviation. (a) KSP, (b) MFV, (c) MFV-1LL, (d) MFV-1DE, (e) MFV-2DE, (f) EnKS.