

Ensemble Generation for Models of Multimodal Systems

ROBERT N. MILLER

Laboratoire de Météorologie Dynamique, Ecole Normale Supérieure, Paris, France, and College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, Oregon

LAURA L. EHRET

College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, Oregon

(Manuscript received 5 December 2000, in final form 10 January 2002)

ABSTRACT

In this work the performance of ensembles generated by commonly used methods in a nonlinear system with multiple attractors is examined. The model used here is a spectral truncation of a barotropic quasigeostrophic channel model. The system studied here has 44 state variables, great enough to exhibit the problems associated with high state dimension, but small enough so that experiments with very large ensembles are practical, and relevant probability density functions (PDFs) can be evaluated explicitly. The attracting sets include two stable limit cycles.

To begin, the basins of attraction of two known stable limit cycles are characterized. Large ensembles are then used to calculate the evolution of initially Gaussian PDFs with a range of initial covariances. If the initial covariances are small, the PDF remains essentially unimodal, and the probability that a point drawn from the initial PDF lies in a different basin of attraction from the mean of that PDF is small. If the initial covariances are so large that there is significant probability that a given point in the initial ensemble does not lie in the same basin of attraction as the mean, the initial Gaussian PDF will evolve into a bimodal PDF. In this case, graphical representation of the PDF appears to split into two distinct regions of relatively high probability.

The ability of smaller ensembles drawn from spaces spanned by singular vectors and by bred vectors to capture this splitting behavior is then investigated, with the objective here being to see how well they capture multimodality in a highly nonlinear system. The performance of similarly small random ensembles drawn without dynamical constraints is also evaluated.

In this application, small ensembles chosen from subspaces of singular vectors performed well, their weakest performance being for an ensemble with relatively large initial variance for which the Gaussian character of the initial PDF remained intact. This was the best case for the bred vectors because of their tendency to align tangent to the attractor, but the bred vectors were at a disadvantage in detection of the tendency of an initially Gaussian PDF to evolve into a bimodal one, as were the unconstrained ensembles.

1. Introduction

In many applications in ocean modeling and numerical weather prediction, ensembles of model-generated forecasts are used to evaluate the sensitivity of the model to variations in initial and boundary conditions and in forcing. Along with a forecast based on one's best estimate of forcing and initial conditions, some number of forecasts is generated in parallel, with forcing, initial conditions and possibly other parameters differing slightly from the original forecast. The perturbed initial, boundary, and forcing data used to generate this collection, or ensemble, of forecasts must be chosen judiciously, since it is impractical to represent all possible

types of errors within the ensemble. Perturbations are therefore chosen either at random according to an explicit estimate of analysis error statistics or according to some dynamically based importance criteria. The question of which, if any, importance criterion is the best one is a matter of some controversy in the numerical weather prediction community. In this work, we investigate the performance of ensembles in the context of a strongly nonlinear system.

The initial conditions for a forecast are considered here as a random vector, with some probability density function (PDF). The evolution of the PDF as governed by the dynamics of the forecast model and the noise contains all available information about the consequences of errors in the initial conditions, boundary conditions, and forcing. If the underlying dynamics are nonlinear, the evolution of the PDF can be quite complex, and simple initial PDFs can evolve under the influence

Corresponding author address: Dr. Robert N. Miller, College of Oceanic and Atmospheric Sciences, Oregon State University, Ocean Admin. Bldg. 104, Corvallis, OR 97331-5503.
E-mail: miller@coas.oregonstate.edu

of the dynamics into very different forms over the course of a model run.

In order for a small ensemble to contain quantitative information about the consequences of the initial, boundary, and forcing errors, it must contain some of the essential features of the PDF. We seek the most efficient methods of ensemble generation in strongly nonlinear systems, that is, those methods that allow generation of small ensembles with the best performance. The general task of evaluating the performance of small ensembles in highly nonlinear systems is an extremely broad one, due to the vast repertoire of nonlinear behaviors. In this work we restrict our attention to the consequences of dealing with a system with multiple attractors.

The small ensembles used in our computations are generated by the methods used by major weather centers (see, e.g., Houtekamer and Derome 1995; Molteni et al. 1996; Toth and Kalnay 1993, 1997). We compare the results of these calculations among themselves and to the results of calculations with large ensembles generated at random from given PDFs.

In the ensemble prediction system in use at the European Centre for Medium-Range Weather Forecasts (ECMWF), ensemble members are generated according to the singular vectors of a forecast evolution operator linearized about a forecast based on the current initial condition. The singular vectors specify the directions of greatest growth of the linearized system over a predetermined time interval. Since singular vectors are specified relative to a fixed norm, the norm must also be specified; this introduces an additional element of choice into the generation of the ensemble. The ideal choice for the norm would reflect the distribution of errors in the initial analysis (see, e.g., Ehrendorfer and Tribbia 1997). Since reliable estimates of initial error covariance are rarely available, proxies are used; see, for example, Molteni et al. (1996). Application of singular vectors to study predictability of the atmosphere date back to Lorenz (1965). The literature on application of singular vectors is extensive; see Ehrendorfer and Tribbia (1997) or Molteni et al. (1996) and references therein.

The ensemble system in operational use at the National Centers for Environmental Prediction (NCEP) is based on a scheme in which ensemble members are determined iteratively in a process referred to as “breeding”; see, e.g., Toth and Kalnay (1993, 1996, 1997), Kalnay and Toth (1996), or Szunyogh et al. (1997). New perturbations are calculated each forecast cycle and normalized to have the same amplitudes as the perturbations from the previous cycle. The vectors derived from the breeding process are known as “bred vectors,” hereafter BVs. The initial amplitude of the perturbations is a parameter in the breeding algorithm, as are the parameters that define the norm used in the normalization process. Since the breeding process bears some similarity to the analysis cycle, it is plausible that the results of breeding can reflect the errors in initial conditions. Two points

of view on this topic were recently expressed by Errico and Langland (1999) and Toth et al. (1999).

Much work on the consequences of different methods of ensemble generation has been performed with the Lorenz (1963) model; see, for example, Trevisan and Pancotti (1998) and Anderson (1996, 1997), and references therein. This is a useful model, which captures much of the indeterminacy of numerical weather prediction, but does not deal with the problems of higher-dimensional models. Other studies, for examples, Houtekamer and Derome (1995), Ehrendorfer and Tribbia (1997), or Szunyogh et al. (1997), have described experiments with numerical weather prediction models of varying resolution. These have the obvious advantage of being similar to working forecast models, but they are highly complex and computationally demanding, due in large part to their high-dimensional state vectors. Interpretation of the results is a daunting task, and demands on computer resources place limits on the size of ensembles that can be generated, and the number of ensemble experiments that can be performed.

Hamill et al. (2000) performed a series of ensemble experiments with a baroclinic quasigeostrophic periodic channel model. Their calculations included ensembles of BVs, ensembles of singular vectors, and ensembles chosen according to statistical criteria but not subject to explicit dynamical constraints, as are BV and singular-vector ensembles. This latter method is similar to the “perturbed observation” (PO) method used operationally by the Canadian Meteorological Centre. In their comparisons of different ensemble generation methods, Hamill et al. found that PO ensembles had desirable statistical characteristics, as revealed by a series of specialized statistical tests.

Little has been done on ensemble forecasting in ocean models. There are fewer observations of the ocean available, and global first guesses with which to initialize models are often limited to poorly resolved climatology. The best choice of ensemble generation method for any given ocean modeling system may be influenced by different considerations from those encountered in numerical weather prediction. Moore (1999) performed a series of experiments in which he applied different methods of ensemble generation to a quasigeostrophic model of the Gulf Stream, with promising results: he found statistically significant relationships between forecast skill and ensemble spread in a number of cases.

Multiple attractors in the models of the atmosphere have been the subject of intense study since the late 1970s by a number of investigators in ongoing attempts to understand the dynamics of planetary flow regimes. Early studies such as that of Charney and DeVore (1979) demonstrated multiple steady states in a highly truncated model of a barotropic channel with topography. Legras and Ghil (1985) found detailed attracting structures in a spectral truncation of the barotropic potential vorticity equation on a sphere. Results from these highly trun-

cated models and discussion of their motivation appear in Ghil and Childress (1987).

As more detail is added to these simple models, a level of complexity is reached that precludes explicit calculation of steady states or more complex attracting structures. Attractors in such models are often studied by computing marginal PDFs (i.e., PDFs calculated from projections of model states upon low-dimensional subspaces) based on long calculations. Hansen and Sutra (1986) examined a dataset from the National Meteorological Center (now known as NCEP) consisting of daily 500-mb height fields from 16 winters. They found bimodal structure in one-dimensional marginal PDFs of a wavenumber index for zonal wavenumbers 2–4. Molteni (1996) calculated marginal PDFs from long runs of a three-layer quasigeostrophic hemispheric model that showed distinct evidence of bimodality. Berner and Branstator (2000) presented marginal PDFs calculated from projections of a long run of a global circulation model (GCM) onto planes defined by empirical orthogonal functions (EOFs) of the GCM's 500-hPa height field. These marginal PDFs differed significantly from multivariate Gaussian PDFs, and while they did not exhibit obvious multimodality, regimelike behavior was found in the mean phase space trajectories. A nonlinear perspective on climate prediction, including a discussion of multiple quasi-stationary regimes and climate trends, was presented by Palmer (1999).

Multimodality is also found in models of wind-driven ocean circulation; see, for example, Schmeits and Dijkstra (2000), Speich et al. (1995), and Jiang et al. (1995), and discussion and references therein. These are simplified models, in that none admits baroclinic instability or effects of topography, but the model of Schmeits and Dijkstra (2000) incorporates realistic basin geometry. Qualitative comparisons to long GCM runs and to observations are promising, though not unambiguous.

An evident example of multimodality in ocean circulation is the well-known path variation of the Kuroshio south of Japan. Multiple regimes in the path of the current along the coast are well documented in observations and in a variety of numerical models; see, for example, the discussions in Sekine (1990) and Masuda et al. (1999), and references therein. The current along the coast either follows a straight path, or forms a large meander that separates from the coastline and reattaches downstream, trapping a cyclonic eddy between the main current axis and the coast. Kumaran and Miller (1995) found three steady solutions of a barotropic quasigeostrophic model of the Kuroshio off the coast of Japan, which corresponded to the straight and meandering paths, along with an unstable intermediate state. This is an example of a regional model, with explicit representation of local topography and specified inflow and outflow.

No commonly used method of ensemble generation was designed to address the problem of probabilistic evaluation of forecast skill in a system with multiple

attractors, but it is likely that multiple attractors appear in systems of interest. We wish to find out what common methods of ensemble generation will do when they encounter this situation.

Given an estimate of the PDF, it is possible to assimilate data by applying Bayes's theorem to calculate an approximate conditional PDF for analysis based on a combination of observations and the model forecast. This is the nonlinear filtering approach; see, for example, Rozovskii (1990). Miller et al. (1999, hereafter MCB) and Anderson and Anderson (1999) applied this Bayesian approach to several simple highly nonlinear systems.

In a system with linear dynamics, Gaussian initial error, and Gaussian white model noise and measurement noise, the Bayesian approach reduces to the Kalman filter (see Jazwinski 1970; Evensen and van Leeuwen 1996, and references therein). This is the basis of Evensen's (1994) ensemble Kalman filter, in which the prediction is the ensemble mean, and the Kalman gain is constructed from the ensemble covariance about the mean. Other approaches have been proposed in similar spirit. Lermusiaux and Robinson (1999a,b) and Pham et al. [1998; see also Ballabrera-Poy et al. (2001)] have devised and tested methods for capturing model and analysis errors in restricted subspaces of the model space. Both of these approaches account for the effects of nonlinearity by using the model dynamics to calculate the evolution of these subspaces. For these applications also, it will be useful to find the most efficient method of generating ensembles.

We use the spectral truncation of a nonlinear barotropic quasigeostrophic β -plane channel model used in MCB. This is a strongly nonlinear model whose solution structure is reasonably well understood. Our version has 44 state variables, enough to introduce the difficulties of high-dimensional systems and likely to exhibit behavior to be found in natural geophysical fluid dynamical systems, but small enough to be economical for the large ensemble calculations necessary to determine the explicit evolution of the PDF. This provides a reference dataset that can be used to evaluate the performance of small ensembles. Understanding the behavior of ensemble calculations in this highly nonlinear setting depends on detailed understanding of the structure of solutions to the model system. Much of this paper is therefore devoted to the examination of details of the solutions to the spectral β -plane model and the evolution of the PDFs from random initial conditions.

We begin with the simple case of "strong constraint" or "perfect model" ensembles; that is, no stochastic component is introduced into the ensemble beyond initial conditions, as a first step toward examination of "weak constraint" or "imperfect model" ensembles, in which the model equations are not assumed to be satisfied exactly. Moore (1999) achieved encouraging results with ensemble experiments with a model that was not assumed to be perfect.

The model is described in section 2. The calculation of BVs and singular vectors is described in section 3. Results of the large ensemble experiments are described in section 4, along with comparisons to experiments with small ensembles of singular vectors, bred vectors, and subsamples of the large ensemble. In this section, we document the evolution of PDFs from Gaussian initial conditions to bimodal structures. Section 5 contains discussion and summary.

2. The truncated spectral barotropic model

In this section we examine a truncated spectral model of a barotropic system in a periodic channel with sinusoidal topography on a β plane. While still a simplified system, this example is drawn from a model of the midlatitude atmosphere. This model is of intermediate dimensionality, so it forces us to deal with the problems of high dimension, while still being sufficiently economical that computations with very large ensembles are practical.

Similar systems have been investigated since the work of Charney and DeVore (1979) and Pedlosky (1981) for the purpose of elucidating the occurrence of multiple equilibria in oceanic and atmospheric models. Jin and Ghil (1990) worked with a similar model, including baroclinic extensions. Hamill et al. (2000) worked with a baroclinic quasigeostrophic channel model with eight vertical levels. Their model incorporated a finite-difference method in the horizontal direction, and it had $O(10^4)$ state variables. Here we work with a system described by Gravel and Derome (1993).

The model describes the evolution of the deviation from uniform flow with speed u^* . The total streamfunction is given by $\Psi = -u^*y + \phi(x, y, t)$. Dissipation is by Rayleigh friction with time constant τ , H is the mean depth of the fluid, and the variable topography is given by h . The perturbation streamfunction ϕ then evolves according to the equation of conservation of potential vorticity, which we write in the form

$$\begin{aligned} \frac{\partial}{\partial t} \nabla^2 \phi + J \left(\phi, \nabla^2 \phi + \beta y + \frac{f_0}{H} h \right) \\ + u^* \frac{\partial}{\partial x} \left(\nabla^2 \phi + \frac{f_0}{H} h \right) = -\frac{1}{\tau} \nabla^2 \phi. \end{aligned} \quad (1)$$

As in MCB, we chose the parameters of our channel to mimic the midlatitude atmosphere at 45°N . Our channel has length $L = 28\,000$ km and width $D = 4000$ km. At this latitude, we have the Coriolis parameter $f_0 = 1.03 \times 10^{-4} \text{ s}^{-1}$ and β , the meridional rate of change of the Coriolis parameter given by $\beta = 1.6 \times 10^{-8} \text{ s}^{-1} \text{ km}^{-1}$. The mean depth $H = 10$ km and h is chosen to have peak amplitude of 500 m with functional form given by $2 \sin(2\pi x/L) \sin(\pi y/D)$. We chose $\tau = 22$ days based on Gravel and Derome's (1993) results of a series of experiments with different values of τ and h .

We adopt the following scaling:

$$U = \text{scale speed}, \quad (2)$$

$$x \sim L, \quad y \sim D, \quad \phi \sim UD, \quad (3)$$

$$t \sim L/U, \quad \text{the advective timescale}, \quad (4)$$

$$u^* \sim U, \quad \text{and} \quad (5)$$

$$h \sim H. \quad (6)$$

In this scaling (1) becomes

$$\begin{aligned} \frac{\partial}{\partial t} \nabla^2 \phi + J \left[\phi, \nabla^2 \phi + \left(\frac{\beta^* \alpha}{R_0} \right) y + \left(\frac{\alpha}{R_0} \right) h \right] \\ + u^* \frac{\partial}{\partial x} \left[\nabla^2 \phi + \left(\frac{\alpha}{R_0} \right) h \right] = -\frac{1}{\tau^*} \nabla^2 \phi, \end{aligned} \quad (7)$$

where

$$\alpha = D/L \text{ is the aspect ratio}, \quad (8)$$

$$\beta^* = \beta D/f_0 \text{ is the dimensionless form of } \beta, \quad (9)$$

$$R_0 = U/f_0 L = \text{the Rossby number, and} \quad (10)$$

$$\tau^* = U\tau/L \text{ is the dimensionless decay time}. \quad (11)$$

The Laplacian operator becomes $\nabla^2 = \alpha^2 \partial^2 / \partial x^2 + \partial^2 / \partial y^2$.

The perturbation streamfunction ϕ is expanded spectrally:

$$\phi(x, y, t) = \sum_i x_i(t) G_i(x, y), \quad (12)$$

where the basis functions G are from the set

$$\begin{aligned} G_i \in \left\{ \sqrt{2} \cos \pi n y, 2 \sin 2\pi m x \sin \pi n y, \right. \\ \left. 2 \cos 2\pi m x \sin \pi n y \right\} \\ m = 1, 2, \dots, M, \quad n = 1, 2, \dots, N. \end{aligned} \quad (13)$$

The resulting equations for the spectral components are

$$\begin{aligned} \frac{dx_i}{dt} = \frac{1}{a_i^2 R_0} \left\{ u^* \sum_j b_{ji} (\alpha h_j - R_0 a_j^2 x_j) \right. \\ + \sum_{j,k} c_{ijk} x_j (\alpha h_k - R_0 a_k^2 x_k) \\ \left. + \beta^* \alpha \sum_j b_{ji} x_j - \left(\frac{R_0}{\tau^*} \right) a_i^2 x_i \right\}. \end{aligned} \quad (14)$$

The coefficients a_i , b_{ji} and c_{ijk} are given by

$$a_i^2 = 4\alpha^2 m^2 \pi^2 + n^2 \pi^2 \quad (15)$$

$$b_{ji} = \int_x \int_y G_j \frac{\partial G_i}{\partial x} dy dx, \quad \text{and} \quad (16)$$

$$c_{kij} = \int_x \int_y G_k J(G_i, G_j) dy dx, \quad (17)$$

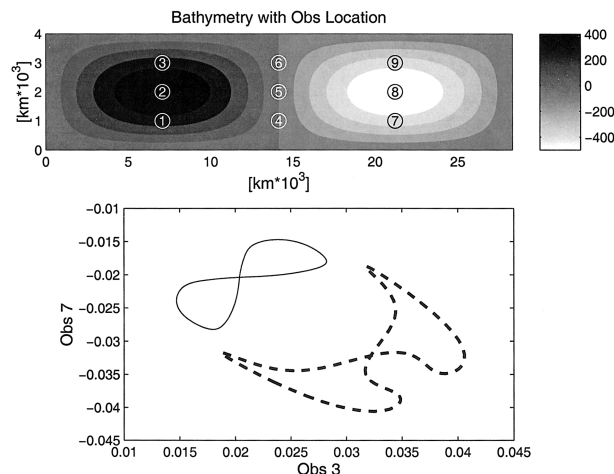


FIG. 1. Schematic diagram of model geometry and locations of simulated observations, along with projections of the two stable limit cycles into a plane defined by values of the perturbation streamfunction at two observation locations. (top) The nine observation points, superimposed on a shaded map of the topography. The shading bar shows contours in m. (bottom) Two stable limit cycles for this system, projected onto a two-dimensional subspace defined by the values of the perturbation streamfunction ϕ at two points in physical space. The horizontal axis is the value of ϕ at location 3. The vertical axis is ϕ at observation 7. Another projection of the limit cycles into a space defined by observations is shown in the top panel of Fig. 20 in MCB.

where m and n in (15) are the zonal and meridional wavenumbers corresponding to the i th spectral component. Following Gravel and Derome (1993), we use a spectral expansion with five zonal wavenumbers and four meridional wavenumbers ($M = 5$ and $N = 4$). In this model, the stability of the equilibrium solutions is governed by Hopf bifurcations in the parameter u^* . The equilibria of the system lose stability through a Hopf bifurcation near 18 m s^{-1} and regain stability through another Hopf bifurcation near 28 m s^{-1} . For this choice of parameters there is no choice of u^* for which multiple stable equilibria were found (MCB; Gravel and Derome 1993). The bifurcation diagram for equilibrium solutions is shown in Fig. 11 of MCB.

There are stable limit cycles in much of the range in which the equilibria are unstable, and at least one case in which multiple stable limit cycles coexist at the same parameter values. The qualitative behavior of the system over a broad parameter range is described in MCB. At $u^* = 19.998 \text{ m s}^{-1}$ there are at least four limit cycles, of which at least two are stable. Examples of four limit cycles projected into three dimensions are shown in MCB. The bifurcation diagram for limit cycles shown in Fig. 12 of MCB is by no means complete. We believe that there are still other limit cycles at this value of u^* , but did not pursue the calculation of other branches. There may be still more complex attracting structures.

In the present study, we work with simulated observations of our model system in physical space. Figure 1 shows a schematic diagram of the model geometry,

and the nine points at which simulated observations of the perturbation streamfunction ϕ are taken. The bottom panel of Fig. 1 shows the projections of the two stable limit cycles into the two-dimensional subspace defined by two widely separated observations.

A model solution with an arbitrary initial condition might eventually approach one of the two limit cycles shown in Fig. 1. There may be other possibilities. These limit cycles are examples of attractors; that is, there is some neighborhood known as the *domain* or *basin of attraction* of each limit cycle (see, e.g., Guckenheimer and Holmes 1983) in which the long-time limit of any trajectory beginning in that neighborhood is the limit cycle itself.

In a model such as this one with multiple attractors, if the errors in the initial condition are sufficiently great that one cannot determine with confidence the basin of attraction in which the initial condition lies, an initially Gaussian PDF might be expected to evolve into one with distinctly bimodal structure. This possibility was illustrated schematically by Anderson (1997) in Fig. 11 of that article. The structure of such a PDF can be captured with a very large ensemble, but this is inefficient and cumbersome. Here we compare small ensembles generated by different methods to see which ones contain the most information about the consequences of initial errors.

In models with large state dimension, PDFs with complex structure may be difficult to discern. Here we address the difficulty of representing the PDF of a complex system directly. In order to do so, we must first determine the basins of attraction of the two limit cycles shown in Fig. 1.

In order to find the boundary between basins of attraction of the two known stable limit cycles, we arbitrarily chose 50 points equally spaced in time on each limit cycle, and constructed the line segments in state space between corresponding points, which were determined by matching extrema. We then chose 20 equally spaced points on each of these line segments, and performed 1000 model runs with initial points chosen at each of the 20 points on each of the 50 line segments joining the two limit cycles. Each of these model runs was of seven periods duration. We assigned each of the 1000 trajectories from the above-described initial points in phase space to one basin of attraction or the other, depending on an average 44-dimensional Euclidean distance from the last period of the trajectory to each limit cycle; a period here refers to the period of the orbit shown as a dashed curve in Fig. 1. There was no essential change in the results from the sixth to the seventh periods. The raw results of this series of experiments are shown in Fig. 2, in which the two orbits and the cylindrical surface containing the 1000 initial points are shown projected into a space defined by three observations in physical space. In Fig. 2 the dashed curve and darker area are associated with the orbit shown as a dashed curve in Fig. 1, while the solid curve and




FIG. 2. Results of a series of runs with initial points chosen on lines between the two stable limit cycles. Axes are values of observations 3, 7, and 8; refer to the top panel of Fig. 1. The gridded surface is a projection of the cylinder bounded by the two limit cycles into this three-dimensional space. The solid and dashed limit cycles correspond to the limit cycles in the bottom panel of Fig. 1. Points on the lighter part of the surface lie in the basin of attraction of the solid orbit, and points on the darker part of the surface lie in the basin of attraction of the dashed orbit, according to the criterion described in the text.

lighter area are associated with the solid curve in Fig. 1. We were thus able to estimate the position of the boundary between the two basins of attraction. We refer to this boundary as a separatrix.

Our conceptual picture of the invariant sets of a system such as this one with two stable limit cycles separated by a separatrix surface is shown in the top panel of Fig. 3. We expect the projections of the separatrix in our system from the full 44-dimensional state space to a 2- or 3-dimensional space for the purpose of visualization to be tortuous. Our visualization of the separatrix with the two attractors, projected into a three-dimensional space, is shown in the bottom panel of Fig. 3. (The reader may refer to the appendix for details of how this figure was constructed.) From Fig. 3 we see that the projection of the separatrix into three dimensions is indeed tortuous, and it seems to fold around the lower (red) limit cycle. Qualitatively, we might expect that initial PDFs centered on this lower limit cycle would be more likely to develop into bimodal structures than initial PDFs with similar covariance structures centered on the upper-limit cycle.

3. Methods for generation of limited ensembles

a. Singular values and singular vectors

The singular vectors, or optimal vectors corresponding to a given solution of a dynamical system can be

seen as those perturbations that result in the greatest divergence from the given solution over a fixed time interval. These vectors may not coincide with the eigenvectors of the linearized system. In relevant cases of solutions that are asymptotically stable, that is, all small perturbations decay in the limit of infinite time, there may be perturbations that grow for a finite time. Methods for calculation of these singular vectors and the implications of their existence have been widely discussed in the literature; see, for example, Ehrendorfer and Tribbia (1997), Ehrendorfer et al. (1996), Legras and Vautard (1996) and references therein.

The treatment here follows Ehrendorfer and Tribbia (1997). Given a solution \mathbf{u} to the nonlinear ordinary differential equation,

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(\mathbf{u}), \quad (18)$$

with initial conditions $\mathbf{u}(0) = \mathbf{u}_0$, the difference $\delta\mathbf{u}(t)$ between $\mathbf{u}(t)$ and a nearby trajectory $\mathbf{v}(t)$ with initial conditions $\mathbf{v}(0) = \mathbf{u}_0 + \delta\mathbf{u}(0)$ can be approximated by the linear equation:

$$\frac{d\delta\mathbf{u}}{dt} = \mathbf{f}'(\mathbf{u})\delta\mathbf{u}. \quad (19)$$

Under fairly general conditions, the nonlinear solution operator Φ_T , which maps an initial condition for (18) into the solution at time T , will be differentiable

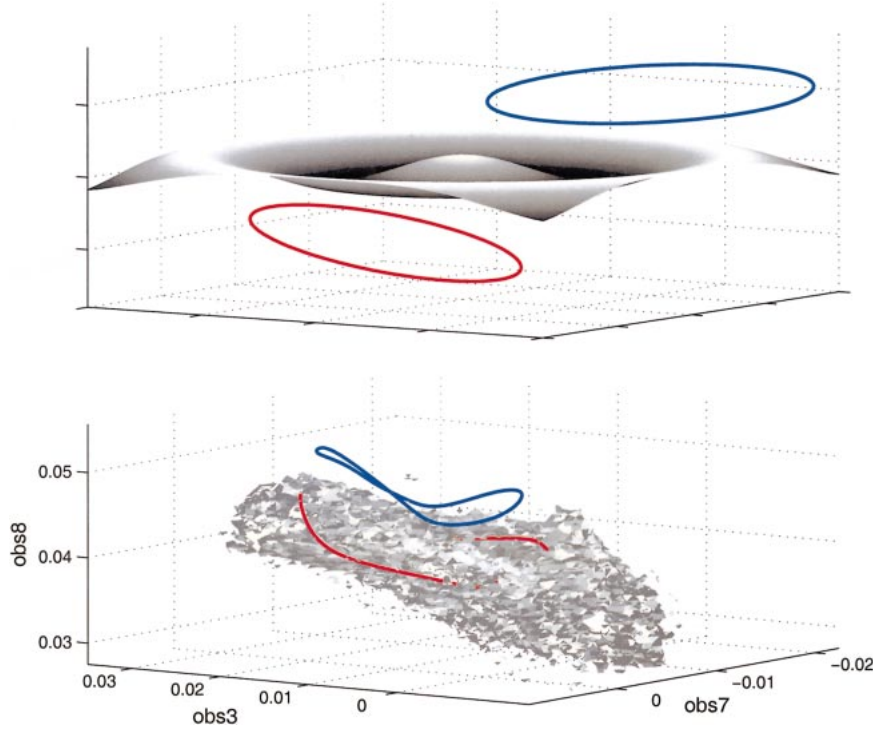


FIG. 3. (top) Schematic diagram of the attracting set of a three-dimensional system with two limit cycles, separated by a two-dimensional separatrix surface, i.e., the boundary between their basins of attraction. (bottom) Projection into a three-dimensional subspace of two stable limit cycles of our model, and a portion of the separatrix. The upper (blue) limit cycle corresponds to the solid limit cycle in Fig. 2 or the solid limit cycle in the bottom panel of Fig. 1 and the lower (red) limit cycle corresponds to the dashed limit cycle. The axes are observations 3, 7, and 8, as in Fig. 2, but the viewpoint has been changed for clarity.

with respect to the initial condition, and the result of perturbing the initial condition to (18) can be expressed by the equation

$$\mathbf{v}(T) = \Phi_T \mathbf{v}(0) \approx \mathbf{u}(T) + \mathbf{M}(0, T) \delta \mathbf{u}(0), \quad (20)$$

where $\mathbf{M}(0, T)$ is the matrix of partials of Φ_T with respect to the components of the initial conditions, evaluated at \mathbf{u}_0 . The vector $\mathbf{u}(T) + \mathbf{M}(0, T) \delta \mathbf{u}(0)$ is called the tangent linear approximation, and the matrix $\mathbf{M}(0, T)$ is called the resolvent. In this case it represents an approximation of the evolution from time $t = 0$ to $t = T$ of deviations from the solution $\mathbf{u}(t)$. Most relevant systems of the form (18) will be sufficiently smooth that the matrices $\mathbf{M}(t_1, t_2)$, that is, the approximate evolution from time $t = t_1$ to $t = t_2$ of deviations from a given solution, will exist and be well behaved.

The singular values of \mathbf{M} (we drop the explicit time dependence in this section), that is, the square roots of the eigenvalues of $\mathbf{M}^T \mathbf{M}$, determine whether the L^2 norm of a given perturbation will grow or decay locally in time, but the impact of initial errors on the accuracy of a given forecast depends on the distribution of those errors.

We consider the strong constraint formulation; that is, we examine only the effect of errors in initial con-

dition. Suppose the initial errors are normally distributed with zero mean and covariance \mathbf{V} . Then we should find the deviation vector $\delta \mathbf{u}$, which maximizes the inner product $(\mathbf{M} \delta \mathbf{u}, \mathbf{M} \delta \mathbf{u})$ subject to $(\delta \mathbf{u}, \mathbf{V}^{-1} \delta \mathbf{u}) = 1$. The inner product $(\cdot, \mathbf{V}^{-1} \cdot)$ defines a norm in terms of the positive definite matrix \mathbf{V}^{-1} . The set of all points \mathbf{x} in state space with $(\mathbf{x}, \mathbf{V}^{-1} \mathbf{x}) = 1$ is an ellipsoid, with its longest axis aligned along the leading EOF of the initial error distribution. We therefore seek the maximum of the functional:

$$J(\mathbf{v}) = (\mathbf{M} \mathbf{v}, \mathbf{M} \mathbf{v}) + \lambda(1 - (\mathbf{v}, \mathbf{V}^{-1} \mathbf{v})). \quad (21)$$

Setting the first variation δJ equal to zero leads to the generalized eigenvalue problem:

$$\mathbf{M}^T \mathbf{M} \mathbf{v} = \lambda \mathbf{V}^{-1} \mathbf{v}. \quad (22)$$

The covariance matrix \mathbf{V} is positive definite, so we can define a new variable \mathbf{y} according to $\mathbf{v} = \mathbf{V}^{1/2} \mathbf{y}$. Substituting this expression for \mathbf{v} into (22) and multiplying on the left by $\mathbf{V}^{1/2}$ leads to the symmetric eigenvalue problem:

$$\mathbf{V}^{1/2} \mathbf{M}^T \mathbf{M} \mathbf{V}^{1/2} \mathbf{y} = \lambda \mathbf{y}, \quad (23)$$

which can be written as

$$(\mathbf{M}\mathbf{V}^{1/2})^T(\mathbf{M}\mathbf{V}^{1/2})\mathbf{y} = \lambda\mathbf{y}. \quad (24)$$

The matrix on the left-hand side of (24) can be written

$$\mathbf{G}^T\mathbf{G} = \mathbf{Y}\mathbf{\Lambda}\mathbf{Y}^T, \quad (25)$$

where $\mathbf{G} = \mathbf{M}\mathbf{V}^{1/2}$, $\mathbf{\Lambda}$ is a diagonal matrix whose nonzero elements are the eigenvalues of $\mathbf{G}^T\mathbf{G}$, and \mathbf{Y} is an orthogonal matrix whose columns are the eigenvectors of $\mathbf{G}^T\mathbf{G}$. These eigenvectors are the singular vectors of \mathbf{M} with respect to the norm defined by \mathbf{V}^{-1} .

Now write $\mathbf{\Pi} = \mathbf{G}\mathbf{Y}\mathbf{\Lambda}^{-1/2}$. We then have

$$\mathbf{\Pi}^T\mathbf{\Pi} = \mathbf{\Lambda}^{-1/2}\mathbf{Y}^T\mathbf{G}^T\mathbf{G}\mathbf{Y}\mathbf{\Lambda}^{-1/2} \quad (26)$$

$$= \mathbf{I}, \quad (27)$$

so $\mathbf{\Pi}$ is an orthogonal matrix. Now,

$$\mathbf{\Pi}\mathbf{\Lambda}\mathbf{\Pi}^T = \mathbf{G}\mathbf{Y}\mathbf{\Lambda}^{-1/2}\mathbf{\Lambda}\mathbf{\Lambda}^{-1/2}\mathbf{Y}^T\mathbf{G}^T \quad (29)$$

$$= \mathbf{G}\mathbf{G}^T \quad (29)$$

$$= \mathbf{M}\mathbf{V}^{1/2}\mathbf{V}^{1/2}\mathbf{M}^T \quad (30)$$

$$= \mathbf{M}\mathbf{M}^T \quad (31)$$

the standard expression for the evolution of the initial error covariance, an expression seen as part of the Kalman filter; see, for example, Bennett (1992).

From (31) we can derive a relatively inexpensive way to calculate an efficient approximation to the covariance evolution of a complex model. Put another way, it allows us to make quantitative estimates of the consequences of errors in the initial conditions. A related method for approximating the evolution of the initial error covariance has been used at ECMWF (cf. Fisher 1996).

We can do this by calculating a few leading eigenvalues and eigenvectors of $\mathbf{G}^T\mathbf{G}$, and writing

$$\mathbf{M}\mathbf{M}^T \approx \hat{\mathbf{\Pi}}\hat{\mathbf{\Lambda}}\hat{\mathbf{\Pi}}^T, \quad (32)$$

where the quantities with the hats ($\hat{\cdot}$) are truncated in rank by deleting all but a restricted number of columns of \mathbf{Y} in the definition of $\mathbf{\Pi}$ and retaining only the submatrix of $\mathbf{\Lambda}^{-1/2}$ corresponding to the retained columns.

Actually (cf. Ehrendorfer and Tribbia, 1997), the definition of $\mathbf{\Pi}$ implies that

$$\mathbf{G} = \mathbf{\Pi}\mathbf{\Lambda}^{1/2}\mathbf{Y}^T, \quad (33)$$

which can be identified as the singular-value decomposition (SVD) of \mathbf{G} , so the desired vectors could be obtained by performing an SVD directly on \mathbf{G} (cf. Ehrendorfer and Tribbia 1997). Here, $\mathbf{G}\mathbf{G}^T$ will be an approximation of the new covariance, with the columns of $\mathbf{\Pi}$ being the EOFs. Computations involve the resolvent of the forward model and its adjoint. There is no need to write down the enormous matrices involved in the calculation implied by (31).

The singular values calculated in this fashion can be related to the ordinary L^2 singular values, that is, the singular values with the matrix \mathbf{V} in (21) chosen to be the identity. Here, \mathbf{M} admits the SVD:

$$\mathbf{M} = \mathbf{Q}\mathbf{D}^{1/2}\mathbf{U}^T, \quad (34)$$

where \mathbf{Q} and \mathbf{U} are orthogonal and $\mathbf{D}^{1/2}$ is diagonal. We may now write an alternative form $\tilde{\mathbf{G}}$ of \mathbf{G} in the form $\tilde{\mathbf{G}} = \mathbf{D}^{1/2}\mathbf{U}^T\mathbf{V}^{1/2}$, so

$$\tilde{\mathbf{G}}^T\tilde{\mathbf{G}} = \mathbf{V}^{1/2}\mathbf{U}\mathbf{D}\mathbf{U}^T\mathbf{V}^{1/2} \quad (35)$$

$$= \mathbf{V}^{1/2}\mathbf{M}^T\mathbf{M}\mathbf{V}^{1/2} \quad (36)$$

$$= \mathbf{Y}\mathbf{\Lambda}\mathbf{Y}^T = \mathbf{G}^T\mathbf{G}, \quad (37)$$

where \mathbf{Y} , \mathbf{G} , and $\mathbf{\Lambda}$ are as in (25). The SVD $\tilde{\mathbf{G}} = \tilde{\mathbf{\Pi}}\mathbf{\Lambda}^{1/2}\mathbf{Y}^T$ leads to

$$\tilde{\mathbf{\Pi}}\mathbf{\Lambda}\tilde{\mathbf{\Pi}}^T = \mathbf{D}^{1/2}\mathbf{U}^T\mathbf{V}\mathbf{U}\mathbf{D}^{1/2} \quad (38)$$

$$= \mathbf{Q}^T\mathbf{M}\mathbf{V}\mathbf{M}^T\mathbf{Q} \quad (39)$$

in analogy to (28)–(31). This tells us that the effect of the resolvent operator \mathbf{M} on the initial covariance \mathbf{V} amounts to a change of coordinates followed by a modification of the variances referred to the new coordinate system, followed by yet another change of coordinates. The variances in the new coordinate system increase or decrease depending on whether the L^2 singular values of \mathbf{M} are greater or less than one; if any L^2 singular value of \mathbf{M} exceeds one, then one can find a direction in phase space in which errors can be expected to increase.

Figure 4 shows the eigenvalues and singular values of the resolvent of the solution operator for (14) over one period of each of the stable limit cycles. Singular values greater than 1.0 correspond to locally growing directions in state space. Examination of the eigenvalues (left panels) shows that both limit cycles are stable, but it is clear from the singular values shown in the right panels that there is a fairly large number of directions in phase space in which perturbations grow over the course of a single period for both limit cycles.

b. Breeding

Ensembles at NCEP are generated by the breeding method, described by Toth and Kalnay (1993, 1996, 1997), Kalnay and Toth (1996), and Szunyogh et al. (1997). The bred vectors (BVs) are the normalized differences between two nearby trajectories at the end of an integration of a given length, known as a breeding cycle. As noted by Toth and Kalnay (1997), “this method ‘breeds’ the nonlinear perturbations that grow fastest on the trajectory taken by the evolving atmosphere in phase space.” As pointed out by numerous authors (e.g., Hamill et al. 2000), the BVs can be thought of as representing the result of perturbations that have grown rapidly in the recent past. These may not be the perturbations that will grow most rapidly in the future.

The BVs are often related to the local Liapunov vectors, and the growth rates derived from the normalization factors are similarly related to the Liapunov exponents; see, for example, Lichtenberg and Lieberman (1983) or Legras and Vautard (1996). Errico and Lang-

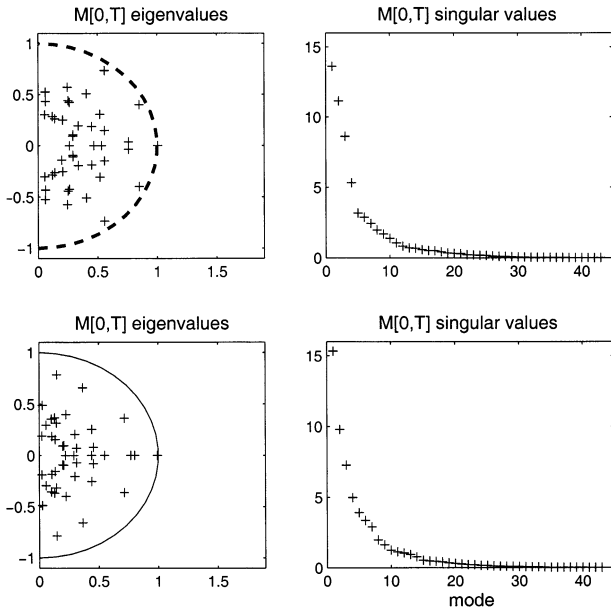


FIG. 4. Stability of the limit cycles for (14). The four panels show eigenvalues and singular values of the resolvent of the solution operator Φ_T for T given by one period. (top panels) Eigenvalues and singular values of the resolvent evaluated at a point on the dashed limit cycle shown in Figs. 1 and 2. Axes on the left panels are real and imaginary components. The fact that the eigenvalues lie inside the unit circle in the complex plane indicates that the orbits are stable. (right panels) Singular values of the resolvent. Values exceeding unity signify local growth in time. (bottom panels) Similar to top panels but for the solid orbit.

land (1999), however, note specific differences between the breeding process and common methods for estimating Liapunov exponents.

The Liapunov exponents are defined in terms of the resolvent $\mathbf{M}(t_1, t_2)$ of the system about a given trajectory. Following Legras and Vautard (1996), we examine the limit

$$\lambda(\mathbf{e}) = \lim_{t_2 \rightarrow \infty} \frac{1}{t_2 - t_1} \ln \left(\frac{\|\mathbf{M}(t_1, t_2)\mathbf{e}\|}{\|\mathbf{e}\|} \right) \quad (40)$$

for \mathbf{e} in state space. This limit exists, is independent of t_1 , and takes on at most n values $\lambda_1 > \lambda_2 > \dots > \lambda_n$. The λ_i are known as the Liapunov exponents. We can also define the limit operator:

$$S_\infty = \lim_{t_2 \rightarrow \infty} \frac{1}{t_2 - t_1} [(\mathbf{M}(t_1, t_2))^* \mathbf{M}(t_1, t_2)]^{1/2(t_2 - t_1)}, \quad (41)$$

where the superscript $*$ denotes conjugate transpose. The eigenvectors of S_∞ are known as the local Liapunov vectors. They depend on t_1 and on the norm in which the limit in (41) is chosen. We distinguish the case $t_2 \rightarrow \infty$ from $t_1 \rightarrow -\infty$ and define $S_{-\infty}$ for this latter limit in a manner similar to (41). Legras and Vautard (1996) refer to the eigenvectors of S_∞ and $S_{-\infty}$ as the forward and backward Liapunov vectors, and identify the forward Liapunov vectors with the limit of the singular

vectors as time increases, and the backward Liapunov vectors as the limit of the BVs as the length of the breeding cycle increases. Reynolds and Errico (1999) investigated the convergence of the singular vectors to the forward and backward Liapunov vectors in a three-layer quasigeostrophic model of the atmosphere. Lichtenberg and Lieberman (1983, chapter 5) point out that the Liapunov exponents for a limit cycle correspond to the eigenvalues of $\mathbf{M}(0, T)$, the resolvent evaluated over a period. Since both of the limit cycles considered here are stable, the largest eigenvalue of the $\mathbf{M}(0, T)$ is unity, and the local Liapunov vector is tangent to the orbit.

We performed four breeding experiments on our model, two beginning on or near each of the two known stable periodic orbits. In these pairs of experiments, we attempted to mimic the breeding strategy of Kalnay and Toth (1996), Szunyogh et al. (1997), and Toth and Kalnay (1993, 1996, 1997). Our approach in the first pair of breeding experiments was to simulate a process in which the reference state was a point on the limit cycle, and the only information available to our forecast and analysis system consisted of noisy measurements of the perturbation streamfunction taken at the nine points shown in Fig. 1.

In our analysis system for the first pair of experiments, data were assimilated by a simple optimal interpolation scheme, in which the model error covariance matrix was assumed to be diagonal in the model state space. The model error was assumed to be isotropic; that is, the error variance was a function of the square of the wavenumber. The variance was red in wavenumber space, falling off exponentially as the square of the wavenumber. We obtained the first initial condition by starting with the reference point, generating a vector of nine noisy observations by adding independent random noise to the reference solution at the nine points shown in Fig. 1, and using our optimal interpolation scheme to assimilate the noisy data. This, of course, had the effect of moving the analysis away from the reference solution. The result was used as an initial condition for an integration that was designated the control. We then added to and subtracted from the control an arbitrary perturbation with amplitude equal to $0.05 \times (\sum_i \sigma_i^2)^{1/2}$, where σ_i^2 is the mean square deviation of the i th component of the state vector from its mean value, averaged around a single period. This gave us two additional initial conditions, which we designated as the perturbations. Integrations were then performed with the control and perturbation initial conditions for a time interval of one period. To begin the next breeding cycle, initial conditions for the new control were obtained by using our optimal interpolation scheme to assimilate new noisy observations. Since the breeding interval is one period, successive measurements differ only by observation noise. The initial conditions for the perturbation forecasts for the next breeding cycle were determined as follows. First, the differences \mathbf{d}_+ and \mathbf{d}_- between the results of the perturbations and the result of the original




FIG. 5. Two cycles of the breeding process for the (top) dashed and (bottom) solid limit cycles referred to an analysis based on a small set of noisy observations. Horizontal axis is scaled time. Breeding cycle is one period. Vertical axis is the difference between the perturbation streamfunction at location 3 (see top panel in Fig. 1) and that value for the limit cycle itself. The three curves shown in each panel are the control and the two perturbed trajectories. Initial points for each breeding cycle are depicted as circled symbols. This plot is designed to have the same form as plots in the literature (Kalnay and Toth 1996; Szunyogh et al. 1997; Toth and Kalnay 1993, 1996, 1997) illustrating the breeding process.

control forecast were calculated. In general, we will have $\mathbf{d}_+ + \mathbf{d}_- \approx 0$. In order to generate the new perturbations, we formed the quantity $(\mathbf{d}_+ - \mathbf{d}_-)/2$. This difference was normalized to have the same amplitude as the initial perturbations, and added to and subtracted from the initial conditions for the new control. This process is illustrated schematically in Fig. 4 of Toth and Kalnay (1993). Results of these experiments for the dashed and solid limit cycles are shown in Fig. 5. These plots are designed to have the same form as plots in the literature (Kalnay and Toth 1996; Szunyogh et al. 1997; Toth and Kalnay 1993, 1996, 1997) illustrating the breeding process.

Use of noisy observations in the breeding procedure introduces a stochastic component into the process whose effect is not clear. The presence of this stochastic component constitutes one distinct difference between the breeding method and methods commonly used to calculate Liapunov exponents of dynamical systems (e.g., Lichtenberg and Lieberman 1983, chapter 5). One can imagine that observation noise could force transition of the trajectories calculated from the breeding process from one basin of attraction into another; one could imagine a situation in which the control and perturbation trajectories were in different basins, but the noise-free

breeding process could also produce this outcome and we have not observed such behavior.

Another difference between the breeding method and methods used for calculation of Liapunov exponents lies in the generation of the control state at the beginning of each breeding cycle. In our breeding experiments, this control state is derived from assimilating noisy observations into the reference state on the limit cycle. The perturbations added to this new control state do not depend explicitly on the data. In our case there is no error of representativeness, and the observations, though sparse, are distributed uniformly in space (see Fig. 1). After 18 breeding cycles of the first experiments, neither the BVs from the solid limit cycle nor those from the dashed appeared to converge to the tangent vector to the limit cycle. However, the normalization factors for the solid orbit cluster about unity with fairly small spread, as expected. A histogram of normalization factors from the breeding process for the dashed limit cycle (i.e., the one more likely to give rise to a split PDF) shows two peaks, both less than one. This result is consistent with the hypothesis that an occasional perturbation vector is generated in the basin of the other (solid) orbit; recall the proximity of the separatrix to the dashed (red) limit cycle in Fig. 3.




FIG. 6. Similar to Fig. 5 but for the breeding process referred to the limit cycle rather than to a small set of noisy observations.

The two other breeding experiments were performed without introducing synthetic observations. Since we have the luxury of knowing the “true” solution in this case, we can use the true solution in place of the control. In this series of experiments, we constructed the perturbed initial conditions by adding perturbations of equal magnitude and opposite directions to the limit cycle as defined above, integrating for a period, normalizing the difference between the solution based on the perturbed initial condition and the true solution, and starting over. We observed convergence of the BVs to the tangent vector, and the convergence of the normalization factors to one, as expected. The breeding process referred to the exact solution converges very slowly, since several conjugate pairs of eigenvalues have moduli close to 1.0 (see Fig. 4). Even so, the approximation applied by Samelson (2001) of the BVs being tangent to the limit cycles is reasonable. The results of this series of experiments are shown in Fig. 6.

4. Ensemble calculations

In this section, results of calculations performed with small ensembles generated according to three different strategies are compared among themselves and with results of calculations with large ensembles drawn from given distributions without regard to dynamical constraints. The small ensembles were chosen from linear combinations of singular vectors, from linear combi-

nations of BVs, and by subsampling the large ensembles.

A total of six large ensemble calculations with 1000 samples were performed, three for each of the two stable limit cycles. We observed little difference in the statistics between ensembles of 500, 750, and 1000 samples and therefore concluded that the results of our experiments with 1000 samples were reliable. In many cases, the statistics based on 250 or fewer samples were discernibly different from the results of the ensemble experiment with 1000. The initial samples were drawn from a Gaussian population with diagonal covariance. The variance of each initial component was chosen to be $\Sigma \times \sigma_i^2$, where σ_i^2 is the mean square deviation of the i th component of the state vector from its mean value, averaged around a single period and Σ is a scalar used to normalize the total variance. For a given limit cycle, the three ensemble experiments differed in their total initial variance, that is, the factor Σ . Covariance statistics were calculated every quarter period.

If the initial ensemble were centered on the limit cycle and lay entirely in its basin of attraction, we would expect the ensemble to spread out in a few directions, and decay in the others at first. The perturbations that grow most rapidly at first will be those in the directions of the leading singular vectors of the linearized differential operator, that is, the matrix $\mathbf{f}'(\mathbf{u})$ in (19). This growth will not, in general, be modal, and the perturbations that grow most rapidly over a finite time interval




FIG. 7. Evolution of PDFs in a neighborhood of a stable periodic orbit. Axes are aligned along the backward singular vectors corresponding to the largest singular values of the resolvent for the solution operator for one period. The initial PDF was Gaussian, with diagonal covariance matrix. The standard deviations of each component were set to 50% component-by-component rms amplitudes, evaluated over a period of the limit cycle, as indicated by the label above the upper left-hand panel. These are two-dimensional marginal PDFs, calculated by a kernel method from ensembles projected into this space. Times relative to the beginning of the experiment are shown in the upper right-hand corner of each panel, e.g., the panel labeled $[1/4T]$ depicts the PDF at one-quarter period following the initial time.

of length T will be those that lie in the direction of the singular vectors of the resolvent $\mathbf{M}(0, T)$ [see (20)]. Eventually all perturbations in the basin of attraction will approach the limit cycle itself. In this fashion, after an initial increase in variance in some directions, the total variance in directions transverse to the stable orbit would decrease asymptotically to zero after a long time. Perturbations in directions tangent to the orbit are neutrally stable, and the result of such a perturbation in the long-time limit is to produce a point displaced along the periodic orbit from the original point. For those perturbations of the initial ensemble outside the basin of attraction of the limit cycle, the points would approach the other limit cycle, or possibly other attracting sets, if they exist. A significant number of the latter would result in a final distribution that would be far from Gaussian.

Since each of the ensemble members evolves according to the exact model dynamics, the different trajectories differ only in their initial conditions, and the sequence of PDFs one might infer from the ensemble would be an approximate solution to Liouville's equation:

$$\frac{\partial P}{\partial t} = -\nabla \cdot (\mathbf{f}(\mathbf{x}, t)P), \quad (42)$$

where P is the PDF and \mathbf{f} represents the right-hand side of (14). A study of the importance of Liouville's equation in the study of predictability appears in the article by Ehrendorfer (1994). Results of our ensemble calculations are shown in Figs. 7–13.

Figures 7 and 8 depict the evolution of an initially Gaussian PDF in a neighborhood of the solid limit cycle. These figures show a marginal PDF, that is, a two-dimensional PDF calculated by a kernel method (see, e.g., Silverman 1986) from the projection of the ensemble onto the plane defined by the two backward singular vectors corresponding to the two largest singular values. The “backward singular vectors” (Legras and Vautard 1996), also referred to as “left singular vectors,” “final time singular vectors,” or “evolved singular vectors” are the normalized results of applying the resolvent to the singular vectors. The initial PDF was chosen with $\Sigma = 0.5$; that is, the variance was chosen to be half that of the limit cycle itself about its mean. Subspaces were chosen for these figures and for the figures relating to our experiments with the dashed orbit for clarity of pre-




FIG. 8. Similar to Fig. 7 but showing detail of the sixth period of integration of the ensemble.

sensation. Nearly all of the points of the initial ensemble lie in the basin of attraction of the solid limit cycle; that is, nearly all remain near that limit cycle for the duration of our six-period ensemble experiment. A small region containing most of the support of the PDF travels around

the limit cycle. The progress of this region around the limit cycle in quarter-period increments for the first period of integration is easily observed in Fig. 7. We see that in the first three quarter period of the model run the PDF seems to spread over both orbits, but this is

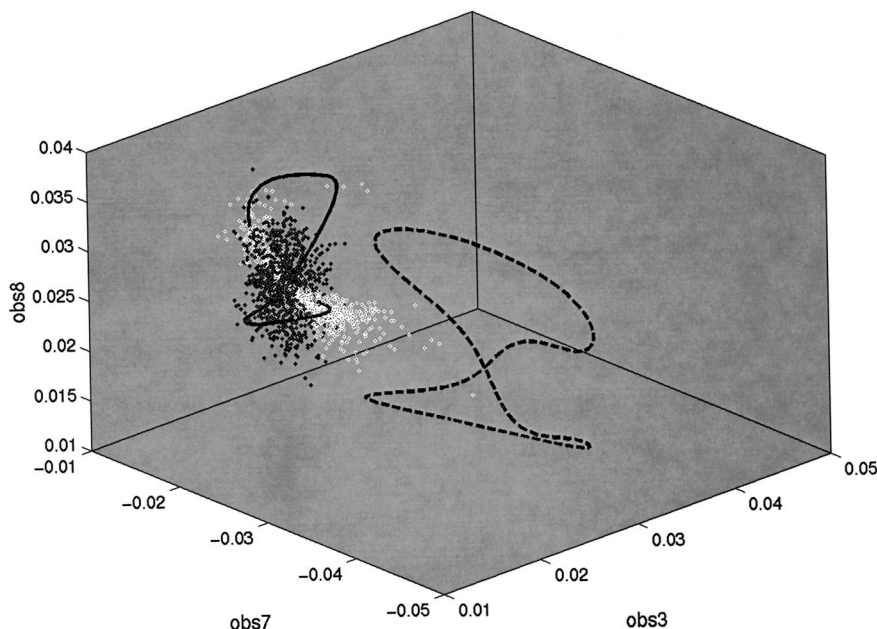


FIG. 9. The ensemble at the beginning and end of a six-period run. Periodic orbits are shown as solid and dashed. Initial ensemble points are shown as black dots. Ensemble points after six periods are shown as white dots. Axes are measurements shown in Fig. 1.

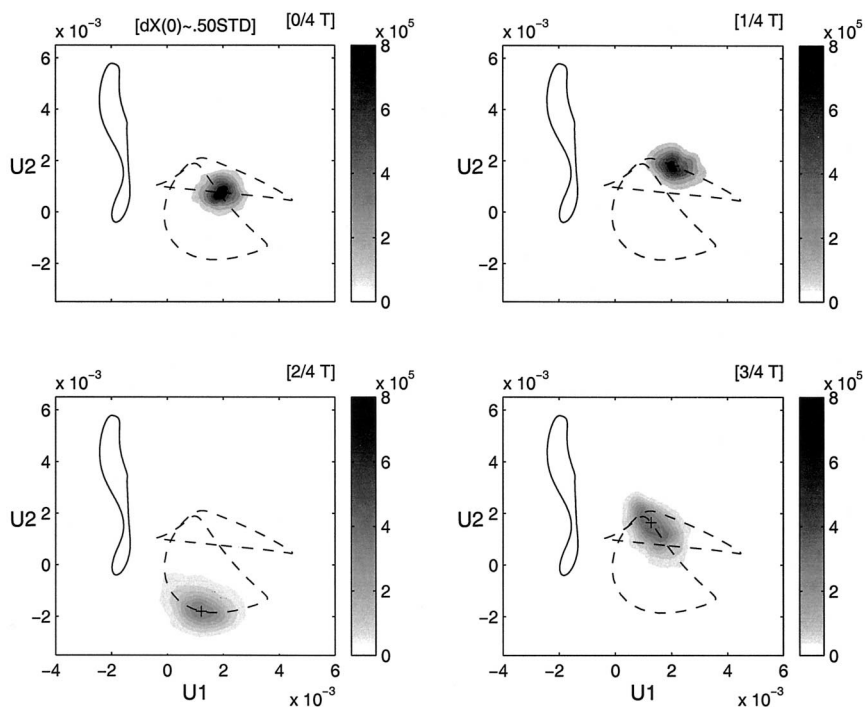


FIG. 10. Evolution of PDFs in a neighborhood of a stable periodic orbit. Panel labels are similar to those in Fig. 7.




FIG. 11. Continuation of Fig. 10, showing alternating marginal PDFs for a 10-period run. Shading values for times beyond the initial have been reduced from Fig. 10 since the maximum of the PDF by this time has been reduced by an order of magnitude, reflecting the spread of the PDF. No such changes were necessary in Figs. 7 and 8 because of the concentration of the PDF in that experiment.




FIG. 12. Detail of last two periods of the 10-period run shown in Figs. 10 and 11. Shading is the same as in Fig. 11.

an artifact of the projection of the ensembles into this two-dimensional space. Figure 8 shows the PDF at quarter-period intervals during the sixth period of the ensemble simulation. By the end of the six-period run the PDF has concentrated along the solid periodic orbit. The probability of a random initial point drawn from the initial PDF shown in the upper-left panel of Fig. 7 resulting in an orbit whose long-time limit approaches the dashed orbit is small. We do not expect the PDF to evolve into a form with significant concentration of probability outside the basin of attraction of the solid orbit in a longer simulation, since we observe little change in the progress of the PDF around the orbit from period 5 to period 6.

To address the possibility that a significant number of ensemble points lies in the basin of attraction of the dashed orbit after six periods of integration, we plotted, in Fig. 9, the projection of the two orbits into a three-dimensional space, along with a black dot for the initial condition of each ensemble member, and a white dot

for the state space location of each ensemble member after six periods of integration. The subspace chosen here is defined by measurements of the streamfunction at three of the points in physical space shown in Fig. 1. From this view, only 1 out of 1000 initial ensemble points can be said with confidence to have ended up in the basin of attraction of the dashed orbit.

Figures 10–12 depict the evolution of an initially Gaussian PDF in a neighborhood of the dashed limit cycle projected onto the plane defined by the backward singular vectors calculated for that limit cycle. In this case, we see a clear evolution from an essentially Gaussian PDF centered on the limit cycle to a bimodal PDF, with significant concentrations of probability density near both orbits. From this we conclude that significant proportions of the initial ensemble lay on both sides of the separatrix shown in Fig. 3. The splitting process is shown clearly in Fig. 11. After six periods of integration, the greatest concentration of ensemble points remains near the maximum of the initial PDF, but the distribution

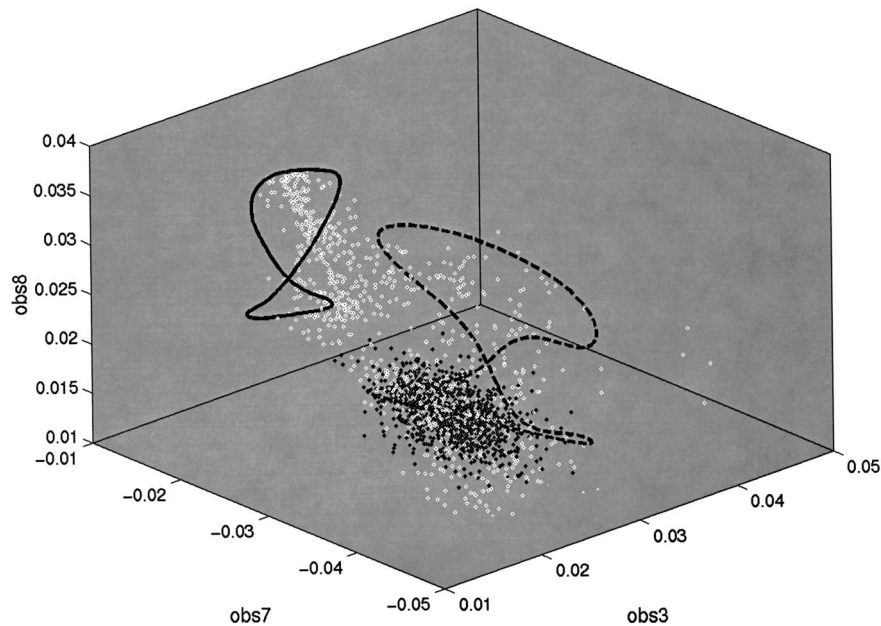


FIG. 13. Similar to Fig. 9 but for the dashed orbit.

has broadened considerably to include a second local maximum, which appears to have formed on the dashed orbit. The reader should note that a change in scale was necessary after the first period and rapid decrease in the maximum of the PDF in order to resolve detailed structure of the PDF in subsequent panels.

Details of the 9th and 10th periods of integration are shown in Fig. 12. In this view, in all panels, the PDF appears to be split, in obvious contrast with Fig. 8. In the panels depicting the PDF at full and half intervals, the secondary PDF maxima appear on the solid orbit. Comparison of panels from the eighth and ninth orbits shows that the secondary maxima of the PDF at the quarter- and three-quarter-period intervals is similarly progressing toward the solid orbit. The reader is again reminded to take care in interpreting 2-dimensional marginal PDFs derived from 44-dimensional ones. Figure 13 is similar to Fig. 9 and clearly shows the evolution of the PDF from a single Gaussian at the initial time to a bimodal structure after six periods of integration.

We now turn to the performance of small ensembles chosen by breeding and by generating singular vectors. Here our main goal is to see which method allows us to best capture the evolution of the PDF with a small ensemble. So far, we have evaluated the evolution of the covariance matrix and examined the extent to which the eigenvalues of the covariance matrix generated from the small ensembles are able to reproduce the eigenvalues of the covariance matrix from the large ensemble. We expect the evolution of the covariance to reflect the splitting process.

Evolution of the small ensembles was calculated for six periods of each of the known stable limit cycles. As in the large ensemble experiments, calculations were

performed for three different values of the initial total variance. Results of our experiments with small ensembles are shown in Figs. 14 and 15.

We performed our singular-vector experiments with L^2 singular vectors. If singular vectors were to be calculated as part of a system with an ongoing assimilation scheme, we would calculate those singular vectors relative to the norm defined by the error covariance as opposed to the L^2 norm. The ECMWF system uses the total energy norm to generate singular vectors (Molteni et al. 1996). This was chosen on the basis of intercomparisons between several candidates. Barkmeijer et al. (1998) compared singular vectors of a T21L5 primitive equation model of the atmosphere generated from the total energy norm to singular vectors generated from the Hessian of the cost function of a three-dimensional variational data assimilation scheme, and found significant differences between results of the two approaches. Moore (1999), in his study of ensemble methods for a model of the Gulf Stream, found that ensembles of singular vectors defined in terms of the total enstrophy norm were more effective for prediction of forecast skill than ensembles of singular vectors defined in terms of the total energy norm.

Our three small ensembles in each experiment were generated by subsampling the 1000-member unconstrained ensemble and by choosing random vectors from a subspace spanned by a small number of singular vectors or BVs. Suppose our best estimate of the initial condition is unbiased, and the errors are Gaussian with covariance \mathbf{V} . In our experiments we chose \mathbf{V} to be the covariance matrix used to generate the corresponding large ensemble. Now let \mathbf{a} be a random vector drawn from the distribution of the initial errors. We can write

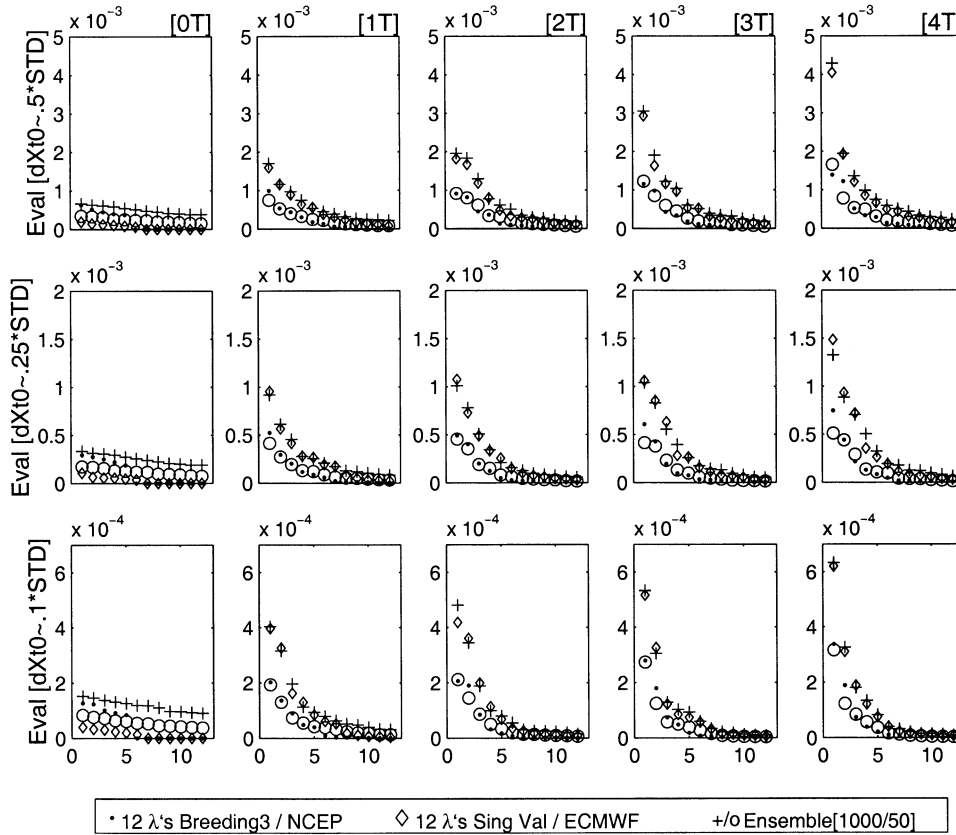


FIG. 14. Results of experiments with small ensembles for the dashed limit cycle. Eigenvalues of the covariance matrices at different times in each experiment are shown for the different methods of ensemble generation. Different symbols depict different methods of choosing the ensemble. Plus signs depict the original ensemble of 1000 members, diamonds depict an ensemble of 50 members chosen from a space spanned by six singular vectors, and dots depict an ensemble drawn from a three-dimensional space derived from BVs. The leading 12 of the 44 eigenvalues are plotted. The sequence of panels from left to right shows increasing time, from initial time to four periods, in intervals of one period. Sequence of panels from bottom to top shows increasing initial standard deviation. The bottom row shows 10% of the component-by-component rms amplitude of the limit cycle; middle row, 25%; and top row, 50%.

$$\mathbf{a} = \hat{\mathbf{Y}}\mathbf{b} + \mathbf{Z}, \quad (43)$$

where $\hat{\mathbf{Y}}$ is the matrix whose columns are the leading N singular vectors of \mathbf{M} , \mathbf{b} is a vector of coefficients, and \mathbf{Z} is a vector orthogonal to the columns of $\hat{\mathbf{Y}}$. The vector of coefficients \mathbf{b} is a random variable with covariance $\mathbf{V}_s = \hat{\mathbf{Y}}^T \mathbf{V} \hat{\mathbf{Y}}$. We therefore choose a relatively small collection of random deviates $\{\mathbf{b}_j, j = 1, 2, \dots\}$ with covariance \mathbf{V}_s and examine the extent to which the behavior of the ensemble $\{\hat{\mathbf{Y}}\mathbf{b}_j\}$ contains the behavior of large ensembles.

We chose the number N of singular vectors to span our restricted space by examining the diagonal matrix \mathbf{D} from the ordinary L^2 SVD of the resolvent of the solution operator for one period of the periodic orbit under consideration. The number N of singular values was chosen to be the smallest integer N with the property

$$\frac{\sum_{i=1}^N \mathbf{D}_{ii}}{\text{tr}(\mathbf{D})} \geq 1 - \epsilon, \quad (44)$$

where we fixed $\epsilon = 0.05$ arbitrarily. Application of this criterion resulted in $N = 6$ for both the solid and dashed orbits.

We chose a restricted space of BVs in a similar fashion. We performed an 18-period breeding run, discarded the results of the first period, and formed a 44×34 matrix \mathbf{B} whose columns consisted of the BVs; recall that there are two BVs for each experiment, resulting from adding and subtracting the normalized perturbation from the analysis. In this breeding run, noisy data were assimilated at the end of each breeding cycle in order to form initial conditions for the next breeding cycle. We performed an SVD on this matrix,

$$\mathbf{B} = \mathbf{U}\mathbf{D}\mathbf{V}^T, \quad (45)$$

and applied (44) to choose the dimension of the restricted subspace. The result was $N = 4$ for the solid orbit and $N = 3$ for the dashed orbit. The corresponding columns of \mathbf{U} form an orthonormal basis of that subspace. Toth et al. (1996) suggested a similar procedure,

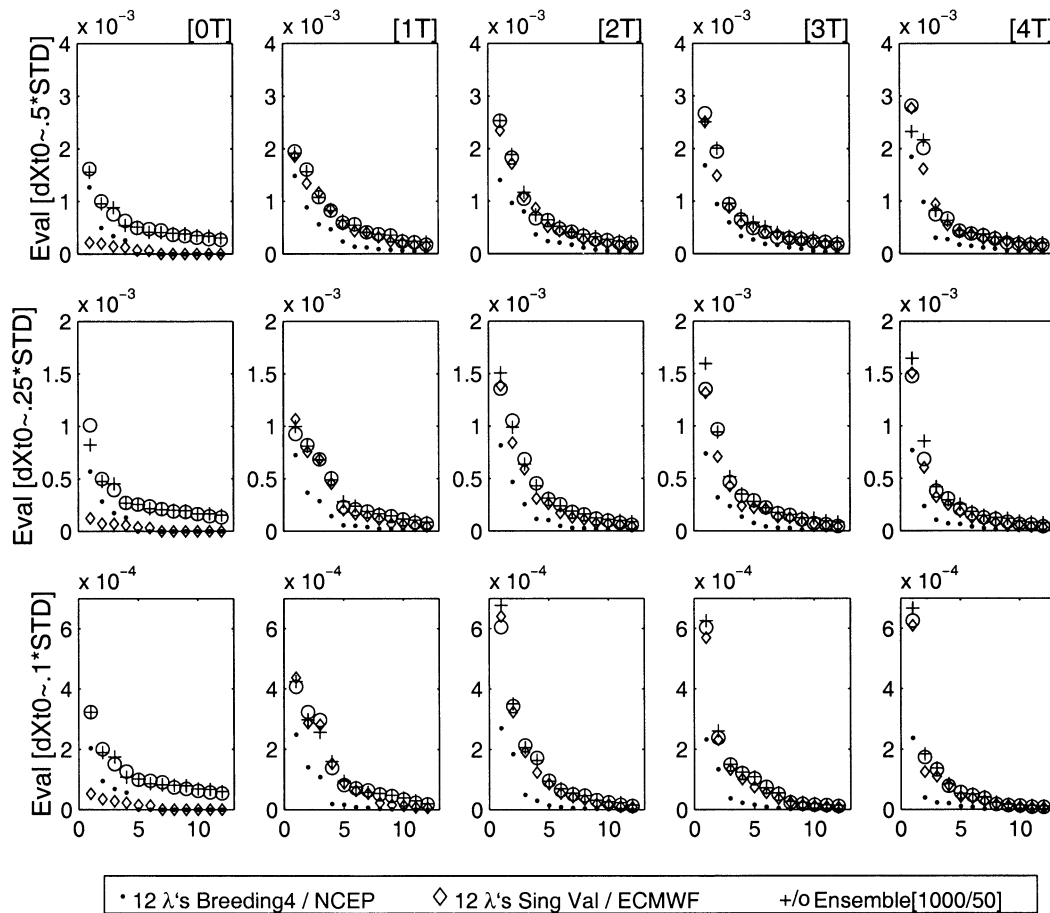


FIG. 15. Similar to Fig. 14 but for solid limit cycle for which the space derived from breeding vectors is four-dimensional.

in which singular vectors might be chosen from a subspace of the full-state space defined by BVs. To the extent that the PDFs remain Gaussian, we can evaluate the ability of the small ensemble to reproduce the properties of the large ensemble by comparing the eigenvalues and eigenvectors of the calculated covariance matrices from the two experiments. Obviously the covariance matrix calculated from the small ensemble will not be of full rank, so comparisons will be restricted to the first few eigenvalues, but we can see the proportion of the variance contained in these first few components.

For each of the two stable limit cycles, four ensemble experiments were performed for three choices of initial ensemble variance. Each row of Figs. 14 and 15 shows the evolution of the eigenvalues of sample covariance matrices calculated from our four different ensemble experiments, one with 1000 members chosen from a Gaussian distribution in state space and three with 50 members: one, the “unrestricted ensemble,” subsampled from the 1000-member ensemble; one with 50 members chosen from a subspace defined by six singular vectors; and one chosen from a subspace defined by three or four BVs, depending on the limit cycle. Only

four periods of our six-period ensemble experiments are shown here. The results after six periods do not differ appreciably from the four-period results shown in the rightmost column of these figures.

Figure 14, which refers to the dashed orbit, the one evidently closer to the separatrix and hence more likely to exhibit splitting of an initially Gaussian PDF, shows the greatest disparity among the results from the different ensembles. While the ensembles of singular vectors seem to capture the least of the initial covariances, at the end of four periods, they produce results similar to those obtained from the large ensemble. Results from the small unrestricted ensemble and the ensemble of BVs are similar to one another, but quite different from the results obtained from the large ensemble or the ensemble of singular vectors. This is probably due to the splitting phenomenon itself. The small unrestricted ensemble is less likely to contain members that lie in the basin of the solid orbit, and we suspect that members of the ensemble of BVs are not oriented in directions favorable to splitting. We remind the reader yet again of the local nature of the breeding and singular vector techniques.

Results of ensemble experiments with the solid orbit are shown in Fig. 15. In this case, as in Fig. 14, for one period into the simulation and later, the ensemble of singular vectors and the unrestricted ensemble produce results similar to those obtained from the large ensemble. As in Fig. 14, the ensemble of BVs underestimates the variances.

5. Discussion and summary

We set out to evaluate the performance of initial ensembles in a system with multiple attractors. We found a specific example in which a Gaussian PDF evolved into a multimodal one, and investigated the performance of small ensembles in that case. We expect that a small ensemble will only exhibit useful statistical behavior in this extremely non-Gaussian setting if some members of the ensemble will reliably be selected on either side of the separatrix.

The question of whether an initially unimodal PDF will split into a multimodal one is easy to quantify in theory. Assume the separatrix (see Fig. 3) of the two basins of attraction examined here divides state space into two regions, call them D_1 and D_2 . There may be other attractors we do not know about, but we have not seen evidence of them. We can say we expect the PDF to split if the integrals of the initial PDF over both D_1 and D_2 exceed some predetermined threshold. We could phrase this question in terms of regime transitions, but it would not be meaningful in this case, in which randomness enters only through initial conditions.

The system under study has 44 state variables, a large state dimension, but small enough that detailed structure of the attractors can be determined and PDFs derived as solutions of Liouville's equation (42) can be conveniently computed by Monte Carlo techniques. This we did by using large ensemble calculations to evaluate the evolution of a series of initially Gaussian PDFs, with different covariances. We then generated smaller ensembles restricted to subspaces spanned by singular vectors and BVs, as well as small subsamples of the large ensemble used to approximate the true PDF, and evaluated the ability of these smaller ensembles to determine whether the initial PDF was likely to split into a bimodal structure. We then compared the evolution of the eigenvectors of the covariance matrices calculated from our different ensembles. We found that small ensembles chosen from a subspace of singular vectors produced eigenvalues that were similar to those resulting from a very large unconstrained ensemble, while small ensembles chosen from BVs consistently underestimated the variances. Moore (1999) found a similar result in his study of predictability of a model of the Gulf Stream.

Eigenvalues of covariance matrices calculated from unrestricted small samples, that is, those chosen by subsampling the large ensemble, were comparable to those calculated from the large ensemble in the case of the solid orbit, that is, the one for which the PDF was un-

likely to split. The performance of the small unrestricted ensemble in the case of the dashed orbit, that is, the one for which the PDF was more likely to split, was comparable to the ensemble of BVs. Hamill et al. (2000) performed a series of statistical tests on unrestricted ensembles and on ensembles of singular vectors and BVs. They found that relatively small unrestricted ensembles actually performed better than ensembles of singular vectors and BVs. Toth and Kalnay (1997) found that ensembles of BVs performed better than unrestricted ensembles. These results are not directly comparable. Hamill et al. (2000) used a model that was somewhat similar to ours in that it was quasigeostrophic model in a periodic channel with relaxation to zonal flow, but it included variability in the vertical with eight vertical layers, and it did not have topography or any other surface forcing other than Ekman pumping. Toth and Kalnay (1997) used a T62L18 version of the NCEP medium-range forecast model.

A number of authors, for example, Toth et al. (1996), Anderson (1997), and Trevisan and Pancotti (1998), have suggested that BVs eventually come to lie tangent to the attractor, while singular vectors point off the attractor. The attractors in our case are limit cycles. There is only one tangent direction, and the Liapunov exponent is related to the largest eigenvalue of the return map (see, e.g., Lichtenberg and Lieberman 1983), in this case 1.0. We therefore expect the breeding process, when conducted with no added noise in the analysis process, to converge to a vector tangent to the limit cycle, with a unit growth rate, corresponding to the unit eigenvalue in Fig. 4. We did not find this convergence in our experiments in which noisy observations were used in the breeding process, but, as noted in section 3b, convergence was observed when the exact limit cycle was used instead of an analysis based on noisy data.

It appears from Fig. 3 that the shortest distance from any point on one of the limit cycles to the separatrix is probably in a direction transverse to the limit cycle itself, rather than tangent to it, but it is not obvious that the singular vectors define the nearest approach to the separatrix. These results are consistent with those of Toth et al. (1996) and Trevisan and Pancotti (1998).

Anderson (1996) and Smith (2000) pointed out that in the special case of systems in which the state point can be safely assumed to lie on a low-dimensional attracting set, choosing ensembles at random from distributions given in state space without explicit dynamical constraints can result in qualitative errors in the estimates of the evolution of the PDF and statistics derived from it. Most practical models of the ocean or atmosphere are sufficiently complex that quantitative descriptors of the attracting sets are not available. In this case, that is, the case in which there is no reliable way to decide the distance between a given point in state space and the attractor, Anderson (1997) argues that little is gained by application of explicit dynamical constraints. Smith (2000) goes so far as to say that at-

tempts at estimating PDFs of the states of numerical weather prediction models are futile, because model uncertainties make even useful approximations to an initial “true” PDF unattainable. Within this framework, in which the ocean or atmosphere is viewed as a highly complex but essentially deterministic system, there can be no useful meaning assigned to the notion of a true state or its PDF.

Both Smith (2000) and Anderson (1997) admit the possibility that in the case of some low-dimensional characterization of the attractor, the formalism of evolution of PDFs from constrained initial condition could be useful. No such characterization of the atmosphere is available at this time, but there may be some specific problems in ocean modeling; see, e.g., the discussion in the introduction on the bimodality of the Kuroshio, as well as Kumaran and Miller (1995), Masuda et al. (1999), and Sekine (1990), and references therein.

In the present work we have considered the perfect model setting as a first step toward the imperfect model case, in which the model contains errors that can be viewed as random variables, in the case of fixed but imperfectly known parameters, or stochastic processes as in the case of imperfectly known forcing fields. In this latter case, even if the underlying dynamics contain a low-dimensional attracting set, the model state will not always lie on that attracting set; one reason might be that some random forcing event might push it off. In such a case as the one examined here, white noise terms added to (14) could force the model state from the basin of attraction of one limit cycle to that of the other, and the stochastic system would have properties resembling the double well (see, e.g., MCB). The next step in the design of ensemble methods for fully stochastic models with imperfectly known parameters or initial conditions will involve ensembles of imperfect model runs.

Acknowledgments. The authors would like to thank Ron Errico, Son Hoang, Lenny Smith, and Zoltan Toth for helpful comments and discussion of earlier drafts of this manuscript. This work was supported by ONR Contract N00014-92-J-1595.

APPENDIX

Explicit Calculation of the Separatrix Surface

We assume that, if there are other attracting sets besides the two known stable limit cycles, their basins of attraction do not meet the cylinder defined by the two limit cycles and the lines that join them, as described in section 2. We cannot prove this rigorously, but our computational experience confirms this belief.

The reader should note that trajectories that do not lie on the separatrix must lie entirely in one basin or the other. In order to visualize the separatrix, as shown in Fig. 3, we began by imposing a regular grid on the

three-dimensional space in which the results are displayed. We then examined the 1000 trajectories with their basin assignments. Each time the projection of a trajectory into our three-dimensional subspace passed through a grid cube, we assigned the value 0 for dashed trajectories and 1 for solid trajectories. No number was assigned to cubes through which no trajectory passed. If multiple trajectories passed through a single cube, we took the average of the values so assigned. Figure 3 was the result of contouring the 0.5 surface in this three-dimensional space. If no number is assigned to a given grid cube, it is plotted as blank space. Figure 3 shows the results from drawing the 0.5 contour on our three-dimensional grid.

REFERENCES

- Anderson, J. L., 1996: Selection of initial conditions for ensemble forecasts in a simple perfect model forecast. *J. Atmos. Sci.*, **53**, 22–36.
- , 1997: The impact of dynamical constraints on the selection of initial conditions for ensemble predictions: Low-order model results. *Mon. Wea. Rev.*, **125**, 2969–2983.
- , 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.
- Ballabrera-Poy, J., P. Brasseur, and J. Verron, 2001: Dynamical evolution of the error statistics with the SEEK filter to assimilate altimetric data in eddy resolving ocean models. *Quart. J. Roy. Meteor. Soc.*, **127**, 233–253.
- Barkmeijer, J., M. van Guizen, and F. Bouttier, 1998: Singular vectors and estimates of the analysis-error covariance metric. *Quart. J. Roy. Meteor. Soc.*, **124**, 1695–1713.
- Bennett, A. F., 1992: *Inverse Methods in Physical Oceanography*. Cambridge University Press, 346 pp.
- Berner, J., and G. Branstator, 2000: Regime signatures in the phase-space tendencies and PDF of an AGCM. Preprints, *15th Conf. on Probability and Statistics in the Atmospheric Sciences*, Asheville, NC, Amer. Meteor. Soc., 92–95.
- Charney, J. G., and J. G. DeVore, 1979: Multiple flow equilibria in the atmosphere and blocking. *J. Atmos. Sci.*, **36**, 1205–1216.
- Ehrendorfer, M., 1994: The Liouville equation and its potential usefulness for the prediction of forecast skills. Part I: Theory. *Mon. Wea. Rev.*, **122**, 703–713.
- , and J. J. Tribbia, 1997: Optimal prediction of forecast error covariances through singular vectors. *J. Atmos. Sci.*, **54**, 286–313.
- , —, and R. M. Errico, 1996: Mesoscale predictability: An assessment through adjoint methods. *Proc. Seminar on Predictability*, Vol. I, Reading, United Kingdom, ECMWF, 157–183.
- Errico, R. M., and R. Langland, 1999: Notes on the appropriateness of “bred modes” for generating initial perturbations used in ensemble predictions. *Tellus*, **51A**, 431–441.
- 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.
- , and P. J. van Leeuwen, 1996: Assimilation of Geosat altimeter data for the Agulhas current using the ensemble Kalman filter with a quasigeostrophic model. *Mon. Wea. Rev.*, **124**, 85–96.
- Fisher, M., 1996: The specification of background error variances in the ECMWF variational analysis system. *Proc. Seminar on Data Assimilation*, Reading, United Kingdom, ECMWF, 645–652.
- Ghil, M., and S. Childress, 1987: *Topics in Geophysical Fluid Dynamics: Atmospheric Dynamics, Dynamo Theory and Climate Dynamics*. Springer-Verlag, 485 pp.
- Gravel, S., and J. Derome, 1993: A study of multiple equilibria in a

- β -plane and a hemispheric model of a barotropic atmosphere. *Tellus*, **45A**, 81–98.
- Guckenheimer, J., and P. Holmes, 1983: *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, 453 pp.
- Hamill, T. M., C. Snyder, and R. Morss, 2000: A comparison of probabilistic forecasts from bred, singular-vector, and perturbed observation ensembles. *Mon. Wea. Rev.*, **128**, 1835–1851.
- Hansen, A. R., and A. Sutera, 1986: On the probability density distribution of planetary-scale atmospheric wave amplitude. *J. Atmos. Sci.*, **43**, 3250–3265.
- Houtekamer, P. L., and J. Derome, 1995: Methods for ensemble prediction. *Mon. Wea. Rev.*, **123**, 2181–2196.
- Jazwinski, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic Press, 376 pp.
- Jiang, S., F.-F. Jin, and M. Ghil, 1995: Multiple equilibria, periodic, and aperiodic solutions in a wind-driven, double-gyre, shallow-water model. *J. Phys. Oceanogr.*, **25**, 764–786.
- Jin, F.-F., and M. Ghil, 1990: Intraseasonal oscillations in the extratropics: Hopf bifurcations and topographic instabilities. *J. Atmos. Sci.*, **47**, 3007–3022.
- Kalnay, E., and Z. Toth, 1996: The breeding method. *Proc. Seminar on Predictability*, Vol. I, Reading, United Kingdom, ECMWF, 69–82.
- Kumaran, S., and R. N. Miller, 1995: A comparison of parallelization techniques for a finite element quasigeostrophic model of regional ocean circulation. *Int. J. Supercomput. Appl.*, **4**, 256–279.
- Legras, B., and M. Ghil, 1985: Persistent anomalies, blocking and variations in atmospheric predictability. *J. Atmos. Sci.*, **42**, 433–471.
- , and R. Vautard, 1996: A guide to Liapunov vectors. *Proc. Seminar on Predictability*, Vol. I, Reading, United Kingdom, ECMWF, 143–156.
- Lermusiaux, P. F. J., and A. R. Robinson, 1999a: Data assimilation via error subspace statistical estimation. Part I: Theory and schemes. *Mon. Wea. Rev.*, **127**, 1385–1407.
- , and —, 1999b: Data assimilation via error subspace statistical estimation. Part II: Middle Atlantic Bight shelfbreak front simulations and ESSE validation. *Mon. Wea. Rev.*, **127**, 1408–1432.
- Lichtenberg, A. J., and M. A. Lieberman, 1983: *Regular and Stochastic Motion*. Springer-Verlag, 499 pp.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.
- , 1965: A study of the predictability of a 28-variable atmospheric model. *Tellus*, **17**, 321–333.
- Masuda, S., K. Akitomo, and T. Awaji, 1999: Effects of stratification and bottom topography on the Kuroshio path variation south of Japan. Part I: Dependence of the path selection on velocity. *J. Phys. Oceanogr.*, **29**, 2419–2431.
- Miller, R. N., E. F. Carter, and S. T. Blue, 1999: Data assimilation into nonlinear stochastic models. *Tellus*, **51A**, 167–194.
- Molteni, F., 1996: On the dynamics of planetary flow regimes. Part II: Results from a hierarchy of orographically forced models. *J. Atmos. Sci.*, **53**, 1972–1992.
- , R. Buizza, T. N. Palmer, and T. Petroliaigis, 1996: The ECMWF ensemble prediction system: Methodology and validation. *Quart. J. Roy. Meteor. Soc.*, **122**, 73–119.
- Moore, A. M., 1999: The dynamics of error growth and predictability in a model of the Gulf Stream. Part II: Ensemble prediction. *J. Phys. Oceanogr.*, **29**, 762–778.
- Palmer, T. N., 1999: A nonlinear dynamical perspective on climate prediction. *J. Climate*, **12**, 575–591.
- Pedlosky, J., 1981: Resonant topographic waves in barotropic and baroclinic flows. *J. Atmos. Sci.*, **38**, 2626–2641.
- Pham, D. T., J. Verron, and M. C. Roubaud, 1998: A singular evolutive extended Kalman filter for data assimilation in oceanography. *J. Mar. Syst.*, **16**, 323–340.
- Reynolds, C. A., and R. M. Errico, 1999: On the convergence of singular vectors toward Lyapunov vectors. *Mon. Wea. Rev.*, **127**, 2309–2323.
- Rozovski, B. L., 1990: *Stochastic Evolution Systems: Linear Theory and Applications to Nonlinear Filtering*. Kluwer Academic, 315 pp.
- Samelson, R., 2001: Periodic orbits and disturbance growth for baroclinic waves. *J. Atmos. Sci.*, **58**, 436–450.
- Schmeits, M. J., and H. A. Dijkstra, 2000: Physics of the 9-month variability in the Gulf Stream region: Combining data and dynamical systems analyses. *J. Phys. Oceanogr.*, **30**, 1967–1987.
- Sekine, Y., 1990: A numerical experiment on the path dynamics of the Kuroshio with reference to the formation of the large meander path south of Japan. *Deep-Sea Res.*, **37**, 359–380.
- Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 175 pp.
- Smith, L. A., 2000: Disentangling uncertainty and error: On the predictability of nonlinear systems. *Nonlinear Dynamics and Statistics*, A. Mees, Ed., Birkhauser, 31–64.
- Speich, S., H. Dijkstra, and M. Ghil, 1995: Successive bifurcations in a shallow-water model, applied to the wind-driven ocean circulation. *Nonlinear Proc. Geophys.*, **2**, 241–268.
- Szunyogh, I., E. Kalnay, and Z. Toth, 1997: A comparison of Lyapunov and optimal vectors in a low-resolution GCM. *Tellus*, **49A**, 200–227.
- Toth, Z., and E. Kalnay, 1993: Ensemble forecasting at NMC: The generation of perturbations. *Bull. Amer. Meteor. Soc.*, **74**, 2317–2330.
- , and —, 1996: Ensemble forecasting at NCEP. *Proc. Seminar on Predictability*, Vol. II, Reading, United Kingdom, 39–61.
- , and —, 1997: Ensemble forecasting at NCEP and the breeding method. *Mon. Wea. Rev.*, **125**, 3297–3319.
- , I. Szunyogh, and E. Kalnay, 1996: Singular, Lyapunov and bred vectors in ensemble forecasting. Preprints, *11th Conf. on Numerical Weather Prediction*, Norfolk, VA, Amer. Meteor. Soc., 53–55.
- , —, E. Kalnay, and G. Iyengar, 1999: Reply to: “Notes on the appropriateness of ‘bred modes’ for generating initial perturbations used in ensemble predictions.” *Tellus*, **51A**, 442–449.
- Trevisan, A., and F. Pancotti, 1998: Periodic orbits, Lyapunov vectors, and singular vectors in the Lorenz system. *J. Atmos. Sci.*, **55**, 390–398.