

Hailiang Du and Leonard A. Smith
Rising above chaotic likelihoods

Article (Published version)
(Refereed)

Original citation:

Du, Hailiang and Smith, Leonard A. (2017) *Rising above chaotic likelihoods*. [SIAM/ASA Journal on Uncertainty Quantification](#), 5 (1). pp. 246-258. ISSN 2166-2525
DOI: [10.1137/140988784](#)

© 2017, [Society for Industrial and Applied Mathematics](#)

This version available at: <http://eprints.lse.ac.uk/id/eprint/87083>
Available in LSE Research Online: March 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Rising Above Chaotic Likelihoods*

Hailiang Du[†] and Leonard A. Smith[‡]

Abstract. Berliner [*J. Amer. Statist. Assoc.*, 86 (1991), pp. 983–952] identified a number of difficulties in using the likelihood function within the Bayesian paradigm which arise both for state estimation and for parameter estimation of chaotic systems. Even when the equations of the system are given, he demonstrated “chaotic likelihood functions” both of initial conditions and of parameter values in the logistic map. Chaotic likelihood functions, while ultimately smooth, have such complicated small scale structure as to cast doubt on the possibility of identifying high likelihood states in practice. In this paper, the challenge of chaotic likelihoods is overcome by embedding the observations in a higher dimensional sequence space; this allows good state estimation with finite computational power. An importance sampling approach is introduced, where pseudo-orbit data assimilation is employed in the sequence space, first to identify relevant pseudo-orbits and then relevant trajectories. Estimates are identified with likelihoods orders of magnitude higher than those previously identified in the examples given by Berliner. The pseudo-orbit data assimilation importance sampler exploits the information both from the model dynamics and from the observations. While sampling from the relevant prior (here, the natural measure) will, of course, eventually yield an accountable sample, given the realistic computational resource this traditional approach would provide no high likelihood points at all. While one of the challenges Berliner posed is overcome, his central conclusion is supported. Chaotic likelihood functions for parameter estimation still pose a challenge; this fact helps clarify why physical scientists maintain a strong distinction between the initial condition uncertainty and parameter uncertainty.

Key words. likelihood, state estimation, chaos, parameter estimation, nonlinear dynamics

AMS subject classifications. 37M10, 62F10, 37D45, 70K99

DOI. 10.1137/140988784

1. Introduction. Nonlinear chaotic systems pose several challenges both for state estimation and for parameter estimation. Chaos as a phenomenon implies sensitive dependence on the initial condition: initially nearby states will eventually diverge in the future. The bifurcations of various chaotic systems [31] reveal how the behavior of the system differs as a parameter value changes. One might think that Bayesian analysis should be able to obtain good estimation both of initial conditions and of parameter values without much trouble.

*Received by the editors October 7, 2014; accepted for publication (in revised form) November 22, 2016; published electronically March 7, 2017.

<http://www.siam.org/journals/juq/5/98878.html>

Funding: This research was supported by the LSE's Grantham Research Institute on Climate Change and the Environment and the ESRC Centre for Climate Change Economics and Policy, funded by the Economic and Social Research Council and Munich Re. Additional support for the work of the first author was also provided by the National Science Foundation Award 0951576 “DMUU: Center for Robust Decision Making on Climate and Energy Policy (RDCEP)”.

[†]Center for Robust Decision Making on Climate and Energy Policy, University of Chicago, Chicago, IL 60637 (hdu@uchicago.edu).

[‡]Centre for the Analysis of Time Series, London School of Economics, London WC2A 2AE, UK (lenny@maths.ox.ac.uk).

Berliner [3] examined the log-likelihood function of estimates of initial conditions and parameter values for the logistic map, noting that chaotic systems can lead to “chaotic likelihood functions” and suggesting that Bayesian analysis would require prohibitively intensive computing. The failure of variational approaches, when applied to long window observations of chaotic systems [8, 13, 22], supports his point. Sensitivity to the initial conditions, on the other hand, suggests that information in the observations (even over a relatively short window) can lead to good estimates of the initial condition [14]. An importance sampling approach to Berliner’s challenge without “intensive computing” is deployed in this paper. Adopting pseudo-orbit data assimilation (PDA) [8, 13] recasts the problem into a higher dimensional sequence space, where truly high likelihood states are successfully located. The challenges of initial condition estimation and parameter estimation are dissimilar for chaotic systems. PDA does not easily generalize to parameter estimation, as it is unclear how to mathematically define a relevant subspace of the parameter space in which the high likelihood trajectories might exist. Thus challenges remain in identifying high likelihood parameter values given the initial condition; this asymmetry reflects differences between the initial conditions and parameter values. In terms of estimating initial conditions given the parameter values, however, Berliner’s challenge as originally stated is met and resolved.

2. Chaotic likelihood function of initial conditions. Berliner’s [3] formulation of the problem is adapted here. The logistic map is the system; in sections 2–4, the parameter $a = 4$ is known but the true initial state \tilde{x}_0 is not. In that case, the experiment is said to fall within the perfect model scenario.¹ In the first four sections of this paper, the mathematical system and the model are identical and the word model is not used. Once real systems are considered, either structural model error or parameter inaccuracy requires one to distinguish the system which generated the data from the model(s) employed to analyze it.

The evolution of system states² $x_i \in \mathbb{R}^m$ is then governed by the nonlinear dynamics $f : x_{i+1} = f(x_i)$, where for the logistic map

$$(1) \quad f(x_i) = ax_i(1 - x_i).$$

Assuming additive observational noise, δ_i , yields observations, $s_i = \tilde{x}_i + \delta_i$, where \tilde{x} is the true system state (Truth) and the observational noise, δ_i , is independent normally distributed, $\delta_i \sim N(0, \sigma^2)$. Under this assumption of normality, the log-likelihood function is

$$(2) \quad LLik(x_0) = - \sum_{i=1}^{n-1} (s_i - f^i(x_0))^2 / 2\sigma^2,$$

where $f^i(x)$ is the i th iteration of $f(x)$, s_i is the i th observation, and n is the duration of observations considered.

Figure 1 shows the chaotic likelihood structure of 1024 candidate values of x_0 drawn uniformly on the interval zero, one. Panel (a) plots the log-likelihood for each x_0 ; this can be

¹The perfect model scenario, when the system and the model are identical, eases use of a digital computer; one avoids the issue of “round-off error” by consistent use of the same computer and code. This implies, of course, that (1) no longer reflects the perfect model. For discussion, see [1, 28, 29] and references therein.

²For $m = 1$, the state x_i is a scalar.




Figure 1. Typical log-likelihood of 1024 states (uniformly distributed on $[0, 1]$) for the logistic map. The true initial condition $\tilde{x}_0 = \sqrt{2}/2$ (denoted by “x”), $\sigma = 0.1$, and $n = 32$. (a) Log-likelihood function, (b) relative log-likelihood to \tilde{x}_0 . States which have (relative) log-likelihood less than -400 are plotted on the -400 horizontal line. All logarithms are using the natural base. The maximum likelihood value found was $\sim e^{-109}$, which corresponds to log-likelihood of -132 relative to Truth.

contrasted with various panels in Berliner’s [3, Figure 3].³ Panel (b) shows the log-likelihood relative to that of the true trajectory of the system state.⁴ For the convenience of illustration, the same normalization used in panel (b) is applied in Figures 2–4 in this paper. From Figure 1, it is clear that no high likelihood states are identified. This is not a case of equifinality.⁵

Given the observational noise distribution, one can add random draws from the inverse of the observational noise distribution to the observation to obtain candidate estimates of initial condition. Figure 2(a) shows the relative log-likelihood of 1024 samples from inverse observational noise. No high likelihood states are identified in this way. (The maximum likelihood value found was $\sim e^{-90}$.) To illustrate the impact of making much more precise observations, consider a case where \tilde{x}_0 is known to be within a region of radius only $\sigma/10$. Figure 2(b) shows the relative log-likelihood of 1024 uniformly sampled states in the region around Truth with $\sigma/10$ radius. Yet again, no high likelihood state is identified. (The maximum likelihood value found increases to $\sim e^{-85}$.)

This difficulty here is not a shortcoming of the likelihood framework, as there are high likelihood states other than Truth. One may demonstrate that such high likelihood states exist

³Here log-likelihood functions based on a sequence of 32 observations are computed because the problem becomes more obvious when more observations are used. Berliner examined 15 and 10 observations. Shorter sequences of observations are examined in section 3.

⁴Note: given that only finite observations are considered the true state of the system is, with probability 1, not the maximum likelihood state.

⁵Equifinality occurs when many potential solutions to a task are each good, making it impossible to identify the true solution given the information in hand. In such cases the sampled likelihood functions are relatively flat. Equidismality arises when the sampled relative likelihood function is flat yet all solutions tested have vanishingly small likelihood given the information. Examining the relative likelihood obscures the difference; fortunately, the expected (distribution of) likelihood can be computed from the noise model alone without knowledge of the true initial condition.




Figure 2. Relative log-likelihood of 1024 states (a) sampled from inverse observational noise, (b) uniformly sampled from $[\tilde{x}_0 - \frac{\sigma}{10}, \tilde{x}_0 + \frac{\sigma}{10}]$. The true initial condition $\tilde{x}_0 = \sqrt{2}/2$ (denoted by “x”), $\sigma = 0.1$, and $n = 32$.

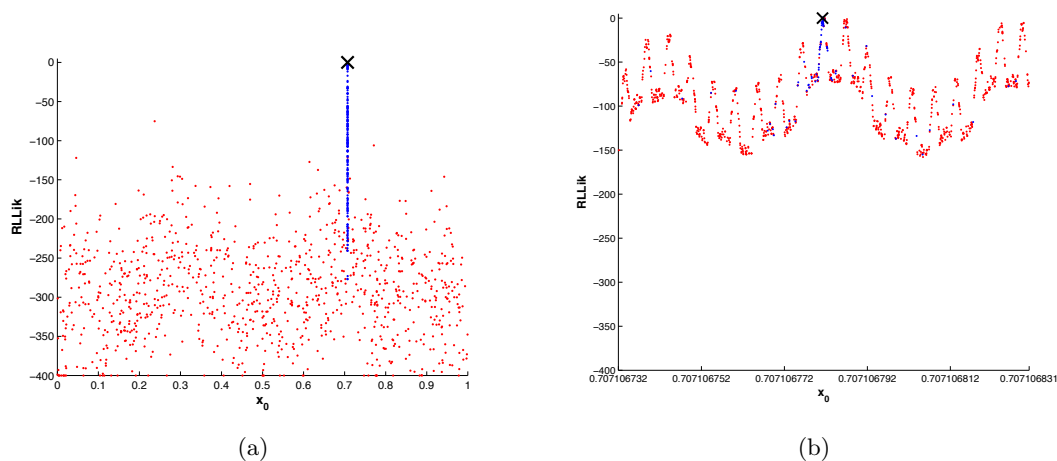


Figure 3. (a) Following Figure 1(b), but in this case including a set of 1024 states (blue), which are extremely close to \tilde{x}_0 , generated by spiral sampling around the \tilde{x}_0 (see footnote 6); (b) zoom in of (a).

by sampling the points on a logarithmic spiral approaching Truth (to machine precision).⁶ Figure 3 demonstrates that high likelihood states other than Truth do exist (i.e., some of the blue points). A smooth curve of the log-likelihood function is only observed within a radius of \tilde{x}_0 smaller than $\sim 10^{-7}$ (see Figure 3(b)).

Without knowing Truth, of course, this approach to identifying high likelihood points is inaccessible. The likelihood function is extremely jagged; as Berliner [3] stressed, finding even one high likelihood state by sampling the state space is prohibitively costly. That said, there is no sense in which “sensitivity to the initial conditions” can be taken to imply that the

⁶In these experiments 1024 points are generated by $\tilde{x}_0 + 2^{-(10 + \frac{60i}{1024})} \epsilon_i$, $i = 1, 2, \dots, 1024$ where ϵ_i is randomly drawn from $U(0, 1)$.

information in the initial condition is “forgotten” or “lost.” There is sufficient information in the observation segment to distinguish high likelihood initial states. Candidate states with nonvanishing log-likelihood relative to Truth can be found by extracting the information from the system dynamics using a relatively new approach to data assimilation: PDA.

3. Importance sampling via PDA.

3.1. Methodology. Uniform sampling in state space is not an efficient approach to locating high likelihood states. As the dimension of the system increases, the task becomes even more computationally impractical. The PDA importance sampler⁷ locates high likelihood states in the trajectory manifold by adopting the PDA approach [8, 13]. PDA takes advantage of the fact that solutions consistent with the known dynamics will lie along a lower dimensional manifold in the higher dimensional sequence space. A brief introduction of the PDA approach is given in the following paragraph (see [8, 13] for additional details).

Given a dynamical system of dimension m and a sequence of n observations,⁸ define a sequence space as the $m \times n$ dimensional space in which a single point can be thought of as a particular series of n states \mathbf{u}_i for $i = 0, \dots, n-1$, where \mathbf{u}_i is an m -dimensional vector. Most points in sequence space do not correspond to a trajectory of the system. Define a *pseudo-orbit*, $\mathbf{U} \equiv \{\mathbf{u}_0, \dots, \mathbf{u}_{n-2}, \mathbf{u}_{n-1}\}$, to be a point in the $m \times n$ dimensional sequence space for which $\mathbf{u}_{i+1} \neq f(\mathbf{u}_i)$ for one or more components of \mathbf{U} . Thus a pseudo-orbit corresponds to a sequence of states which is *not* a trajectory of the system. Each sequence of n observations $\mathbf{s}_i, i = 0, \dots, n-1$, defines a pseudo-orbit (a point in the sequence space). Call this an *observed pseudo-orbit*, $\mathbf{S} \equiv \{\mathbf{s}_0, \dots, \mathbf{s}_{n-2}, \mathbf{s}_{n-1}\}$, which with probability one will not be a trajectory.

Define the *mismatch* to be

$$(3) \quad e_i = |f(\mathbf{u}_i) - \mathbf{u}_{i+1}|.$$

By construction, trajectories have a mismatch of zero. The mismatch cost function is then given by

$$(4) \quad C(\mathbf{U}) = \sum_{i=0}^{n-1} e_i^2.$$

The most common form of PDA simply minimizes the mismatch cost function for \mathbf{U} in the $m \times n$ dimensional sequence space starting from the observed pseudo-orbit. If a gradient descent (GD) approach is adopted,⁹ then a minimum of the mismatch cost function can be obtained by solving the ordinary differential equation

$$(5) \quad \frac{d\mathbf{U}}{d\tau} = -\nabla C(\mathbf{U}),$$

⁷In high dimensional space the sampler targets the relevant lower dimensional trajectory manifold which is more efficient than sampling a hypersphere suggested by one observation. Even in the one dimensional logistic map this approach succeeds by using PDA to sample the trajectory manifold in the n -dimensional sequence space.

⁸For the logistic map, observations and system states share the same one dimensional space.

⁹Other methods for this minimization are available, GD is discussed here due to its simplicity and adequacy.

where τ denotes algorithmic time.¹⁰ In practice, the minimization is initialized with the observed pseudo-orbit, i.e., ${}^0\mathbf{U} = \mathbf{S}$, where the presuperscript zero on \mathbf{U} denotes the initial algorithmic time ($\tau = 0$) of the GD. Under GD, a point in sequence space moves towards the trajectory manifold under (5). The minimization algorithm requires differentiating the mismatch cost function (4), which gives

$$(6) \quad \frac{\partial C(\mathbf{U})}{\partial \mathbf{u}_i} = 2 \times \begin{cases} -(\mathbf{u}_{i+1} - f(\mathbf{u}_i))J(\mathbf{u}_i), & i = 0, \\ -(\mathbf{u}_i - f(\mathbf{u}_{i-1})) + (\mathbf{u}_{i+1} - f(\mathbf{u}_i))J(\mathbf{u}_i), & 0 < i < n-1, \\ -\mathbf{u}_i - f(\mathbf{u}_{i-1}), & i = n-1, \end{cases}$$

where $J(\mathbf{u}_i)$ is the Jacobian of f at \mathbf{u}_i . The ordinary differential equation (5) is solved using the Euler approximation in the examples below.

The mismatch cost function has no local minima other than points on the manifold, for which $C(\mathbf{U}) = 0$ (this defines the trajectory manifold¹¹) and of course every segment of trajectory lies on this manifold [13]. Denote the result of the GD minimization at algorithmic time $\tau = \alpha$ as ${}^\alpha\mathbf{U} \equiv {}^\alpha\mathbf{u}_0, \dots, {}^\alpha\mathbf{u}_{n-1}$. Here α indicates *algorithmic time* in GD (i.e., the number of iterations of the GD minimization). As $\alpha \rightarrow \infty$, the pseudo-orbit ${}^\alpha\mathbf{U}$ approaches a trajectory of the model asymptotically. In other words, the GD minimization takes the observed pseudo-orbit towards a model trajectory (${}^\infty\mathbf{U}$ is a point in sequence space on the trajectory manifold).

In practice, the GD algorithm is run for a finite time and a trajectory is not obtained. Nevertheless, the values of ${}^\alpha\mathbf{u}_0$ for large α define candidate trajectories using information from the observation window $0 \leq i \leq n-1$. For $i > 0$, the i -step preimages of the relevant component of ${}^\alpha\mathbf{u}_i$ provide candidates for the initial state. Complications arise from the fact that the logistic map is two-to-one; these complications have nothing to do with chaos per se.¹² In one-to-one chaotic maps the calculation of preimages is straightforward. The logistic map is a two-to-one map, and in most cases¹³ only one of the two preimages for each ${}^\alpha\mathbf{u}_i$ is relevant to \tilde{x}_{i-1} . In practice a threshold criterion to discard irrelevant preimages must be defined; a simple example would be to discard (with high probability) those preimages whose distance from the corresponding previous observation exceeds some threshold based on the properties of the observational noise (a 3σ criterion is adopted in the following section and shown to be adequate for the purpose of this paper).

3.2. Results. The green points in Figure 4 were located using an PDA importance sampler; note that some have relative log-likelihood near zero (maximum 1.3). As expected, the observations do not contain sufficient information to identify the state of the system at the time of the final observation with the same degree of precision as the initial state. This is

¹⁰The approach can be generalized to situations where the gradient is not known analytically [15]; improving the ability to work without gradient information would widen the application of the approach significantly.

¹¹Back substitution of the solution of $-\mathbf{u}_{n-1} - f(\mathbf{u}_{n-2}) = 0$ into (6) shows that the only critical points for $C(\mathbf{U})$ have $\mathbf{u}_i - f(\mathbf{u}_{i-1}) = 0$ for all i in $0 \leq i \leq n-1$. All points on the trajectory manifold have zero mismatch (are trajectories) and only points on the trajectory manifold have zero mismatch.

¹²They have nothing to do with chaos beyond the fact that one-to-one maps in one dimension cannot display chaotic dynamics, of course.

¹³It is not true in all cases, however. For discussion of the point, see [17].




Figure 4. (a) Following Figure 3(a), but in this case including the states located by PDA importance sampler which are plotted in green. (b) The forward images of those states plotted in (a) at time 31.

reflected in the fact that the green points are much less close to the true state at time 31 (Figure 4(b)) than the corresponding green points at time 0 (Figure 4(a)).

Two experiments were conducted to test the robustness of the PDA importance sampling approach. The first is based on 2048 different realizations of observations given the same initial condition, $\tilde{x}_0 = \sqrt{2}/2$, to examine consistency. The second considers 2048 different initial conditions to examine robustness over different states. Three different observation window lengths were used in each experiment. Tables 1 and 2 shows the results. States with greater likelihood than the true state are often identified.

Given uncertain observations, one can never identify Truth of a chaotic system unambiguously; this was noted by Berliner and MacEachern [2, 20], then Lalley [17, 18], and later explored by Judd and Smith [13]. Using the PDA approach, high likelihood states are indeed found: states with relative log-likelihood larger than minus one are found in every single experimental run. The fact that some PDA states have greater likelihood than Truth reflects the fact that Truth is not expected to be the most likely model state given the observations.

For each experimental run, the minimum distance between those states obtained by PDA importance sampler (whose relative log-likelihood is larger than minus one) and Truth is recorded. The minimum, maximum, and median statistical values of the minimum distance from Truth are reported in Tables 1 and 2. It is clear that the PDA importance sampler identifies states at higher quality (the minimum distance from Truth decreases) as the observation window length increases. This is expected inasmuch as more information from the dynamics is available given a longer window. Table 1 shows that the maximum value of the minimum distance among the 2048 different realizations is 1.49×10^{-10} for a window length of 32 and in Table 2 the maximum value of the minimum distance among different true initial conditions is 2.02×10^{-10} . The PDA importance sampler is both robust and efficient in this case.¹⁴

¹⁴Drawing samples uniformly from within a distance of 0.1 of Truth would require $\sim 10^8$ candidates in order to find a candidate within $\sim 2 \times 10^{-10}$ of Truth; even then that close candidate need not have high likelihood. The results of Tables 1 and 2 were obtained with only 1024 GD minimization iterations in each realization (each and every one of which identified high likelihood states close to Truth).

Table 1

Statistics of high likelihood states located by PDA importance sampler based on 2048 different realizations of observations (of $\tilde{x}_0 = \sqrt{2}/2$) for the logistic map: (i) statistics of the number of states (whose log-likelihood relative to the true state is larger than minus one); (ii) statistics of the number of states (whose relative log-likelihood is larger than zero); (iii) statistics of the minimum distance between the states (whose relative log-likelihood is larger than minus one) and Truth.

| Window length | # of RLLik > -1 | | | # of RLLik > 0 | | | Minimum distance to \tilde{x}_0 | | |
|---------------|-----------------|-----|--------|----------------|-----|--------|-----------------------------------|------------------------|------------------------|
| | Min | Max | Median | Min | Max | Median | Min | Max | Median |
| 32 | 6 | 15 | 8 | 0 | 14 | 6 | 2.00×10^{-15} | 1.49×10^{-10} | 9.98×10^{-12} |
| 16 | 2 | 11 | 8 | 0 | 11 | 6 | 1.57×10^{-10} | 7.63×10^{-6} | 5.13×10^{-7} |
| 8 | 2 | 7 | 7 | 0 | 7 | 5 | 8.58×10^{-8} | 4.55×10^{-2} | 2.56×10^{-4} |

Table 2

Statistics of high likelihood states located by PDA importance sampler based on 2048 different true initial states for the logistic map: (i) statistics of the number of states (whose log-likelihood relative to the true state is larger than minus one); (ii) statistics of the number of states (whose relative log-likelihood is larger than zero); (iii) statistics of the minimum distance between the states (whose relative log-likelihood is larger than minus one) and Truth.

| Window length | # of RLLik > -1 | | | # of RLLik > 0 | | | Minimum distance to \tilde{x}_0 | | |
|---------------|-----------------|-----|--------|----------------|-----|--------|-----------------------------------|------------------------|------------------------|
| | Min | Max | Median | Min | Max | Median | Min | Max | Median |
| 32 | 6 | 360 | 28 | 0 | 338 | 16 | 4.77×10^{-15} | 2.02×10^{-10} | 1.50×10^{-11} |
| 16 | 3 | 56 | 14 | 0 | 48 | 9 | 3.04×10^{-10} | 1.54×10^{-5} | 9.93×10^{-7} |
| 8 | 2 | 20 | 7 | 0 | 20 | 7 | 6.36×10^{-8} | 5.60×10^{-3} | 3.32×10^{-4} |

The experiments above solve by demonstration Berliner's "impossible" challenge: they demonstrate that truly high likelihood points can be located using dynamical information, easing Berliner's identification problem of the initial condition. Selecting an ensemble from this high likelihood set allows for informative forecasts which do not become useless until after those from the point forecasts illustrated by Berliner [3] become uninformative.

3.3. Equifinality and relative likelihood. Maximum likelihood estimation has been widely used [26] since its introduction by Fisher [10] in 1922. The "best" estimate is often chosen from a set of candidates, and only the relative likelihood of candidates within that sample is considered. Figure 5 shows the log-likelihood of 1024 states (the same set used in Figure 2(b)); the red dashed line is the median log-likelihood of those states. In this case, the problem is not one of which estimate to select (equifinality), but one of recognizing that each and every candidate has vanishing likelihood (equidismality). In practice, Truth is unknown, therefore, it cannot be used as a reference, as it is in Figure 2. Given the observations and the noise model, however, the expected log-likelihood of Truth can be determined¹⁵ and serve as a reference. Figure 5, the log-likelihood of 1024 states are plotted along with the expected log-likelihood of Truth (black dashed line). Figure 5 shows that it is not a case of equifinality but a case of equidismality. In cases where it is observed that all traditional candidate states

¹⁵The log-likelihood of Truth is $-\sum_{i=0}^{n-1} \delta_i^2 / 2\sigma^2$ (from (2)), where δ_i (observation noise) is *independently and identically distributed* $\sim N(0, \sigma^2)$ distributed. Let $Z = \sum_{i=0}^{n-1} \delta_i^2 / \sigma^2$, Z being a random variable following a chi-squared distribution with n degrees of freedom. Statistics of the log-likelihood of Truth can therefore be simply derived from Z .




Figure 5. Log-likelihoods of 1024 states uniformly sampled from $[\tilde{x}_0 - \frac{\sigma}{10}, \tilde{x}_0 + \frac{\sigma}{10}]$, the red dashed line is the median log-likelihood of those states, and the black dashed line is the expected log-likelihood of Truth, following Figure 2(b).

have vanishingly small log-likelihood relative to the expected log-likelihood of Truth, a PDA based importance sampling might prove valuable.

4. Implications for quantifying uncertainty in practice. Having solved one of Berliner’s “impossible” challenges for nonlinear systems, a reviewer asks for potential implications this solution might hold in practice. Two possibilities stand out. The first regards improved noise reduction on the initial condition of the forecast via a PDA importance sampler; this would indeed be expected to improve forecast skill. Berliner’s problem, however, focuses on identification of the point at the starting (oldest) observation, while weather forecasts are launched near the time of the end (most recent) observation. In operational weather forecasting, a relatively small set of initial conditions (~ 51) are launched as forecasts in a relatively high dimensional space ($\sim 10^7$) in a Monte Carlo fashion every forecast cycle (typically, every 6 hours); see [12, 19, 32]. As is long held in weather forecasting (see [11] for a discussion), uncertainty in the state at the start of the observation window in time collapses much more dramatically than that in the state at the end (present time); this is consistent with results shown in Figure 4. Figure 2 of [11] captures the challenge Berliner and MacEachern [2, 20], Hansen and Smith [11], and Lalley [18] foresaw. High likelihood states at the present time lie along the unstable manifold; such sets of states are more efficient and informative, in terms of exploring initial condition uncertainty, than traditional approaches like singular vectors [24] or bred vectors [33].

A second possibility lies in reanalysis [25]. A reanalysis uses the most modern weather model to revisit and improve the estimate of the state of the atmosphere of the distant past; here the opportunity for PDA importance sampler to make a contribution appears greater. A reanalysis takes a modern model and reconsiders past periods of time; reanalysis

is now common in atmospheric (weather) and oceanic contexts [16, 23, 34]. This task closely resembles that posed by Berliner. The results achieved in Tables 1 and 2 reflect the aims of reanalysis directly, albeit in a perfectly known one dimensional system. Interpretations in terms of scalar “noise reduction” are limited, as there can be information gained simply by restricting the state estimate to the model manifold.

Reanalysis has many applications, one of significant current interest is index-based insurance. In this case, uncertainty quantification of the index would be of significant interest. In short, the complications of evaluating insurance claims based on certifying damages in far flung regions of the world has led to the insurance policies triggered by a model-based meteorological index [9]. To the extent that the methods deployed in this paper better quantify (and reduce) the uncertainty in the wind speeds of a storm within a few days of its passing, the efficiency of index-based insurance could be improved.

Note from Tables 1 and 2 that the longer the window length the better the estimate (noise reduction) achieved. Noise reduction based upon PDA importance sampler might also help fill in missing historical observations, as well as identify model deficiency when identifying model trajectories intended to reflect past observations.

5. Distinguishing between initial conditions and parameters. The likelihood functions of initial conditions and that of parameter values have similar features [3]. There are, however, fundamental differences in the information available to address these two distinct estimation problems.

Given the structure of the model class, the model parameter value determines dynamical behaviors of the model (e.g., natural measures) which are not changed by the initial condition. Given the model and true parameter value, the search for unknown initial conditions is aided by the restriction the nature measure(s) places on candidate initial conditions in the state space (and, thereby, on trajectories in the sequence space).

It is unclear how to construct similar constraints on unknown parameter values in the parameter space given the true initial state (if they exist¹⁶). Uncertainty in the initial state differs from uncertainty in the parameter value. The information in a *measurement* of the initial condition uncertainty will decay with time and eventually become statistically indistinguishable from a random sample of the natural measure, while the information in the initial condition is preserved, arguably forever.

While assuming the parameter value is perfect might initially appear extreme, it appears less nonsensical given that one has already assumed that the model structure is perfect. Assuming the initial state is perfect indicates a noise free observation is possible. Let the model’s parameters be contained in the vector $\mathbf{a} \in \mathbb{R}^l$. A set of $l + 1$ sequential noise free observations $s_i, s_{i+1}, \dots, s_{i+l}$ would, in general, be sufficient to determine \mathbf{a} [21]. Thus if one noise free observation is obtainable, obtaining only a few more noise free observations would define the true parameter value precisely. Perhaps a more realistic way to put the problem is to estimate the parameter value(s) given realistic observations, without assuming the true initial condition is known. In that case, the goal is to locate high likelihood trajectories (Smith

¹⁶Neither is it clear how to construct the set of parameter values whose corresponding invariant measure contains the “true” initial state, or how to exploit this set; what is clear is how to exploit the existence of a trajectory manifold given a particular value of the parameter.

et al. [30] call these shadowing trajectories) for particular sets of parameter values. It is not yet clear how to solve such a problem. In fact, it is not clear if it is possible to constrain the solution in the parameter space in a manner analogous to the constraints in the space of the initial condition achieved by using a trajectory manifold in the state space.

Given a perfect model structure and knowing the true parameter value(s), the true initial state is a well-defined goal of the identification. Inasmuch as structural model errors imply no true parameter value exists [7, 14], it is unclear how one might define the true initial state and the goal of estimation must be rethought when the structure of the system differs from that of the model.

Despite the importance of model parameters, outside linear systems there is no general method of parameter estimation.¹⁷ Methods have been developed to obtain useful parameter values with some success: McSharry and Smith [21] and Smith [28] estimate model parameters by incorporating the global behavior of the model into the selection criteria; Creveling, Gill, and Abarbanel [5] have exploited synchronization for parameter estimation; Smith et al. [30] focused on the geometric properties of trajectories; Du and Smith [6] select parameter values based on the ignorance score of ensemble forecasts. Each of these methods, however, requires a large set of observations. Challenges remain when only a short sequence of observations is available. Berliner's second challenge still stands.

6. Conclusion. Berliner illustrated that even in the perfect model scenario traditional approaches are unable to provide good estimates of the initial condition for nonlinear chaotic systems. In large part, those failures are due to the inability of traditional approaches to skillfully meld the information in the dynamics of the nonlinear system itself with that in the observations. The importance sampling approach presented here uses PDA to combine information from both observations and dynamics more effectively, thereby locating high likelihood initial states; this achieves one aim Berliner (1991) argued to be impossible. As discussed in section 5, this achievement may prove useful in reanalysis studies.

Despite the similarity of state estimation and parameter estimation, there are fundamental differences between uncertainty in the initial state and uncertainty in parameter value. Significant obstacles remain in meeting Berliner's second challenge regarding parameter estimation while PDA overcomes his first challenge regarding the initial condition.

Acknowledgments. L.A.S. gratefully acknowledges the continuing support of Pembroke College, Oxford. We also acknowledge the insightful comments of Michael Stein and Mark Berliner in discussions of this work.

REFERENCES

- [1] C. BECK, *Effects of roundoff errors on chaotic dynamics*, in Signal Processing VI - Theories and Applications: Proceedings of EUSIPCO-92, J. Vandewalle, R. Boite, M. Moonen, and A. Oosterlinck, eds., Vol. 1, Elsevier, Amsterdam, 1992, pp. 183–186.
- [2] L. M. BERLINER AND S. N. MACEACHERN, *Examples of inconsistent Bayes procedures based on observations on dynamical systems*, Statist. Probab. Lett., 17 (1993), pp. 355–360.

¹⁷That is not to dismiss important advances in parameter estimation when attention is restricted to certain particular classes of model structure [4]; our claim here is that no general approach to parameter estimation is known which is applicable to nonlinear models in general.

- [3] M. BERLINER, *Likelihood and Bayesian prediction for chaotic systems*, J. Amer. Statist. Assoc., 86 (1991), pp. 938–952.
- [4] N. A. C. CRESSIE AND C. K. WIKLE, *Statistics for Spatio-Temporal Data*, Wiley Ser. Probab. Stat., Wiley, Hoboken, NJ, 2011.
- [5] D. CREVELING, P. GILL, AND H. ABARBANEL, *State and parameter estimation in nonlinear systems as an optimal tracking problem*, Phys. Lett. A, 372 (2008), pp. 2640–2644.
- [6] H. DU AND L. A. SMITH, *Parameter estimation through ignorance*, Phys. Rev. E (3), (2012), 016213.
- [7] H. DU AND L. A. SMITH, *Pseudo-orbit data assimilation. II: Assimilation with imperfect models*, J. Atmos. Sci., 71 (2014), pp. 483–495.
- [8] H. DU AND L. A. SMITH, *Pseudo-orbit data assimilation part I: The perfect model scenario*, J. Atmos. Sci., 71 (2014), pp. 469–482.
- [9] C. C. R. I. FACILITY, *Understanding CCRIF: A Collection of Questions and Answers*, CCRIF SPC, 2015.
- [10] R. FISHER, *On the mathematical foundations of theoretical statistics*, Philos. Trans. A., 222 (1922), pp. 309–368.
- [11] J. A. HANSEN AND L. A. SMITH, *Probabilistic noise reduction*, Tellus, 5 (2001), pp. 585–598.
- [12] B. HOSKINS, *The potential for skill across the range of the seamless weather-climate prediction problem: A stimulus for our science*, Q. J. Roy. Meteor. Soc., 139 (2013), pp. 573–584, <https://doi.org/10.1002/qj.1991>,
- [13] K. JUDD AND L. A. SMITH, *Indistinguishable states I: The perfect model scenario*, Phys. D, 151 (2001), pp. 125–141.
- [14] K. JUDD AND L. A. SMITH, *Indistinguishable states II: The imperfect model scenario*, Phys. D, 196 (2004), pp. 224–242.
- [15] K. JUDD, L. A. SMITH, AND A. WEISHEIMER, *Gradient free descent: Shadowing, and state estimation using limited derivative information*, Phys. D, 190 (2004), pp. 153–166.
- [16] E. KALNAY, M. KANAMITSU, R. KISTLER, W. COLLINS, D. DEAVEN, L. GANDIN, M. IREDELL, S. SAHA, G. WHITE, J. WOOLLEN, Y. ZHU, A. LEETMAA, R. REYNOLDS, M. CHELLIAH, W. EBISUZAKI, W. HIGGINS, J. JANOWIAK, K. C. MO, C. ROPELEWSKI, J. WANG, R. JENNE, AND D. JOSEPH, *The NCEP/NCAR 40-year reanalysis project*, Bull. Amer. Meteor. Soc., 77 (1996), pp. 437–471, [https://doi.org/10.1175/1520-0477\(1996\)077<0437:TNYRP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1996)077<0437:TNYRP>2.0.CO;2).
- [17] S. LALLEY, *Beneath the noise, chaos*, Ann. Statist., 27 (1999), pp. 461–479.
- [18] S. LALLEY, *Removing the noise from chaos plus noise*, in Nonlinear Dynamics and Statistics, A. Mees, ed., Birkhäuser, Boston, 2000.
- [19] M. LEUTBECHER AND T. PALMER, *Ensemble forecasting*, J. Comput. Phys., 227 (2008), pp. 3515–3539.
- [20] S. MACEACHERN AND L. BERLINER, *Asymptotic inference for dynamical systems observed with error*, J. Statist. Plann. Inference, 46 (1995), pp. 277–292.
- [21] P. MCSHARRY AND L. A. SMITH, *Better nonlinear models from noisy data: Attractors with maximum likelihood*, Phys. Rev. Lett, 83 (1999), pp. 4285–4288.
- [22] R. MILLER, M. GHIL, AND F. GAUTHIEZ, *Advanced data assimilation in strongly nonlinear dynamical systems*, J. Atmos. Sci., 51 (1994), pp. 1037–1056.
- [23] P. OKE, G. BRASSINGTON, D. GRIFFIN, AND A. SCHILLER, *The bluelink ocean data assimilation system (BODAS)*, Ocean Modell., 21 (2008), pp. 46–70.
- [24] T. N. PALMER, R. BUIZZA, F. MOLTENI, Y.-Q. CHEN, AND S. CORTI, *Singular vectors and the predictability of weather and climate*, R. Soc. Lond. Philos. Trans. Sect. A: Math. Phys. Eng. Sci., 348 (1994), pp. 459–475, <https://doi.org/10.1098/rsta.1994.0105>.
- [25] W. PARKER, *Computer simulation, measurement, and data assimilation*, British J. Philos. Sci., 68 (2017), pp. 273–304.
- [26] Y. PAWITAN, *In all likelihood: Statistical modelling and inference using likelihood*, Clarendon, Oxford, 2001.
- [27] V. PISARENKO AND D. SORNETTE, *Statistical methods of parameter estimation for deterministically chaotic time series*, Phys. Rev. E (3), 69 (2004), 036122.
- [28] L. A. SMITH, *Lacunarity and Chaos in Nature*, Ph.D. thesis, Columbia University, New York, 1987.
- [29] L. A. SMITH, *What might we learn from climate forecasts?*, Proc. Natl. Acad. Sci. USA, 4 (2002), pp. 2487–2492.

- [30] L. A. SMITH, M. CUELLAR, H. DU, AND K. JUDD, *Exploiting dynamical coherence: A geometric approach to parameter estimation in nonlinear models*, Phys. Lett. A, 374 (2010), pp. 2618–2623.
- [31] J. SPROTT, *Chaos and Time-Series Analysis*, Oxford University Press, Oxford, 2003.
- [32] Z. TOTH AND E. KALNAY, *Ensemble forecasting at NMC: The generation of perturbations*, Bull. Amer. Meteor. Soc., 74 (1993), pp. 2317–2330.
- [33] Z. TOTH AND E. KALNAY, *Ensemble Forecasting at NCEP and the Breeding Method*, Monthly Weather Rev., 125 (1997), pp. 3297–3319, [https://doi.org/10.1175/1520-0493\(1997\)125%3C3297:EFANAT%3E2.0.CO;2](https://doi.org/10.1175/1520-0493(1997)125%3C3297:EFANAT%3E2.0.CO;2),
- [34] S. M. UPPALA, P. W. KÄLLBERG, A. J. SIMMONS, U. ANDRAE, V. D. C. BECHTOLD, M. FIORINO, J. K. GIBSON, J. HASELER, A. HERNANDEZ, G. A. KELLY, X. LI, K. ONOGI, S. SAARINEN, N. SOKKA, R. P. ALLAN, E. ANDERSSON, K. ARPE, M. A. BALMASEDA, A. C. M. BELJAARS, L. V. D. BERG, J. BIDLOT, N. BORMANN, S. CAIRES, F. CHEVALLIER, A. DETHOF, M. DRAGOSAVAC, M. FISHER, M. FUENTES, S. HAGEMANN, E. HÓLM, B. J. HOSKINS, L. ISAKSEN, P. A. E. M. JANSSEN, R. JENNE, A. P. MCNALLY, J.-F. MAHFOUF, J.-J. MORCRETE, N. A. RAYNER, R. W. SAUNDERS, P. SIMON, A. STERL, K. E. TRENBERTH, A. UNTCH, D. VASILJEVIC, P. VITERBO, AND J. WOOLLEN, *The era-40 re-analysis*, Q. J. Roy. Meteor. Soc., 131 (2005), pp. 2961–3012, <https://doi.org/10.1256/qj.04.176>.