

Frequentist Perspective on Robust Parameter Estimation Using the Ensemble Kalman Filter



Sebastian Reich

Abstract Standard maximum likelihood or Bayesian approaches to parameter estimation for stochastic differential equations are not robust to perturbations in the continuous-in-time data. In this paper, we give a rather elementary explanation of this observation in the context of continuous-time parameter estimation using an ensemble Kalman filter. We employ the frequentist perspective to shed new light on two robust estimation techniques; namely subsampling the data and rough path corrections. We illustrate our findings through a simple numerical experiment.

Keywords Parameter estimation · Stochastic differential equations · Ensemble Kalman filter · Frequentist approach · Rough path theory

1 Introduction

In this note, we consider the well-studied problem of parameter estimation for stochastic differential equations (SDEs) from continuous-time observations $X_t^\dagger, t \in [0, T]$ [25]. It is well-known that the corresponding maximum likelihood estimator does not depend continuously on the observations $X_t^\dagger, t \in [0, T]$, which can result in a systematic estimation bias [27, 14]. In other words, the maximum likelihood estimator is not robust with respect to perturbations in the observations. Here, we revisit this problem from the perspective of online (time-continuous) parameter estimation [6, 11] using the popular ensemble Kalman filter (EnKF) and its continuous-time ensemble Kalman-Bucy filter (EnKBF) formulations [15, 10, 26]. As for the corresponding maximum likelihood approaches, the EnKBF does not depend continuously on the incoming observations $X_t^\dagger, t \geq 0$, with respect to the uniform norm topology on the space of continuous functions. This fact has been first investigated in [9] using rough path theory [16]. In particular, as already

S. Reich (✉)

Institute of Mathematics, University of Potsdam, Potsdam, Germany

e-mail: sebastian.reich@uni-postdam.de

© The Author(s) 2023

B. Chapron et al. (eds.), *Stochastic Transport in Upper Ocean Dynamics*,
Mathematics of Planet Earth 10, https://doi.org/10.1007/978-3-031-18988-3_15

237

demonstrated for the related maximum likelihood estimator in [14], rough path theory allows one to specify an appropriately generalised topology which leads to a continuous dependence of the EnKBF estimators on the observations. Here we expand the analysis of [9] to a frequentist analysis of the EnKBF in the spirit of [29], where the primary focus is on the expected behaviour of the EnKBF estimators over all admissible observation paths. One recovers that the discontinuous dependence of the EnKBF estimators on the driving observations results in a systematic bias from a frequentist perspective. This is also a well known fact for SDEs driven by multiplicative noise [23].

The proposed frequentist perspective naturally enables the study of known bias correction methods, such as subsampling the data [27], as well as novel de-biasing approaches in the context of the EnKBF.

In order to facilitate a rather elementary mathematical analysis, we consider only the very much simplified problem of parameter estimation for linear SDEs. This restriction allows us to avoid certain technicalities from rough path theory and enables a rather straightforward application of the numerical rough path approach put forward in [13]. As a result we are able to demonstrate that the popular approach of subsampling the data [2, 27, 5] can be well justified from a frequentist perspective. The frequentist perspective also suggests a rather natural approach to the estimation of the required correction term in the case an EnKBF is implemented without subsampling.

We end this introductory paragraph with a reference to [1], which includes a broad survey on alternative estimation techniques. We also point to [9] for an in-depth discussion of rough path theory in connection to filtering and parameter estimation.

The remainder of this paper is structured as follows. The problem setting and the EnKBF are introduced in the subsequent Sect. 2. The frequentist perspective and its implications on the specific implementations of an EnKBF in the context of low and high frequency data assimilation are laid out in Sect. 3. The importance of these considerations becomes transparent when applying the EnKBF to perturbed data in Sect. 4. Here again, we restrict attention to a rather simple model setting taken from [17] and also used in [9]. As a result we build a clear connection between subsampling and the necessity for a correction term in the case high frequency data is assimilated directly. A brief numerical demonstration is provided in Sect. 5, which is followed by a concluding remark in Sect. 6.

2 Ensemble Kalman Parameter Estimation

We consider the SDE parameter estimation problem

$$dX_t = f(X_t, \theta)dt + \gamma^{1/2}dW_t \quad (1)$$

subject to observations $X_t^\dagger, t \in [0, T]$, which arise from the reference system

$$dX_t^\dagger = f^\dagger(X_t^\dagger)dt + \gamma^{1/2}dW_t^\dagger, \tag{2}$$

where the unknown drift function $f^\dagger(x)$ typically satisfies $f^\dagger(x) = f(x, \theta^\dagger)$ and θ^\dagger denotes the true parameter value. Here we assume for simplicity that the unknown parameter is scalar-valued and that the state variable is d -dimensional with $d \geq 1$. Furthermore, W_t and W_t^\dagger denote independent standard d -dimensional Brownian motions and $\gamma > 0$ is the (known) diffusion constant.

Following the Bayesian paradigm, we treat the unknown parameter as a random variable Θ . Furthermore, we apply a sequential approach and update Θ with the incoming data X_t^\dagger as a function of time. Hence we introduce the random variable Θ_t which obeys the Bayesian posterior distribution given all observations X_τ^\dagger , $\tau \in [0, t]$, up to time $t > 0$. Furthermore, instead of exactly solving the time-continuous Bayesian inference problem as specified by the associated Kushner–Stratonovitch equation [6, 26], we define the time evolution of Θ_t by an application of the (deterministic) ensemble Kalman–Bucy filter (EnKBF) mean-field equations [10, 26], which take the form

$$d\Theta_t = \gamma^{-1}\pi_t \left[(\theta - \pi_t[\theta]) \otimes f(X_t^\dagger, \theta) \right] dI_t, \tag{3a}$$

$$dI_t = dX_t^\dagger - \frac{1}{2} \left(f(X_t^\dagger, \Theta_t) + \pi_t[f(X_t^\dagger, \theta)] \right) dt, \tag{3b}$$

where π_t denotes the probability density function (PDF) of Θ_t and $\pi_t[g]$ the associated expectation value of a function $g(\theta)$. The column vector I_t , defined by (3b), is called the innovation, while the row vector

$$K_t(\pi_t) = \gamma^{-1}\pi_t \left[(\theta - \pi_t[\theta]) \otimes f(X_t^\dagger, \theta) \right], \tag{4}$$

premultiplying the innovation in (3a) is called the gain. Here the notation $a \otimes b = ab^T$, where a, b can be any two column vectors, has been used. The initial condition $\Theta_0 \sim \pi_0$ is provided by the prior PDF of the unknown parameter.

A Monte-Carlo implementation of the mean-field equations (3) leads to the interacting particle system

$$d\Theta_t^{(i)} = \gamma^{-1}\pi_t^M \left[(\theta - \pi_t^M[\theta]) \otimes f(X_t^\dagger, \theta) \right] dI_t^{(i)}, \tag{5a}$$

$$dI_t^{(i)} = dX_t^\dagger - \frac{1}{2} \left(f(X_t^\dagger, \Theta_t^{(i)}) + \pi_t^M[f(X_t^\dagger, \theta)] \right) dt, \tag{5b}$$

$i = 1, \dots, M$, where expectations are now taken with respect to the empirical measure. That is,

$$\pi_t^M[g] = \frac{1}{M} \sum_{i=1}^M g(\Theta_t^{(i)}) \tag{6}$$

for given function $g(\theta)$, and all Monte-Carlo samples are driven by the same (fixed) observations X_t^\dagger . The initial samples $\Theta_0^{(i)}$, $i = 1, \dots, M$, are drawn identically and independently from the prior distribution π_0 .

We note in passing that there is also a stochastic variant of the innovation process [26] defined by

$$dI_t = dX_t^\dagger - f(X_t^\dagger, \Theta_t)dt - \gamma^{1/2}dW_t, \quad (7)$$

which leads to the Monte-Carlo approximation

$$dI_t^{(i)} = dX_t^\dagger - f(X_t^\dagger, \Theta_t^{(i)})dt - \gamma^{1/2}dW_t^{(i)} \quad (8)$$

of the innovation in (5).

Remark 1 There is an intriguing connection to the stochastic gradient descent approach to the estimation of θ^\dagger , as proposed in [30], which is written as

$$d\theta_t = \frac{\alpha_t}{\gamma} \nabla_\theta f(X_t^\dagger, \theta_t) d\tilde{I}_t, \quad (9a)$$

$$d\tilde{I}_t = dX_t^\dagger - f(X_t^\dagger, \theta_t)dt \quad (9b)$$

in our notation, where $\alpha_t > 0$ denotes the learning rate. We note that (9) shares with (3) the gain times innovation structure. However, while (3) approximates the Bayesian inference problem, formulation (9) treats the parameter estimation problem from an optimisation perspective. Both formulations share, however, the discontinuous dependence on the observation path X_t^\dagger , and the proposed frequentist analysis of the EnKBF (3) also applies in simplified form to (9). We also point out that (3) is affine invariant [18] and does not require the computation of partial derivatives.

We now state a numerical implementation with step-size $\Delta t > 0$ and denote the resulting numerical approximations at $t_n = n\Delta t$ by $\Theta_n \sim \pi_n$, $n \geq 1$. While a standard Euler–Maruyama approximation could be applied, the following stable discrete-time mean-field formulation of the EnKBF

$$\Theta_{n+1} = \Theta_n + K_n \left\{ (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) - \frac{1}{2} \left(f(X_{t_n}^\dagger, \Theta_n) + \pi_n[f(X_{t_n}^\dagger, \theta)] \right) \Delta t \right\} \quad (10)$$

is inspired by [3] with Kalman gain

$$K_n = \pi_n \left[(\theta - \pi_n[\theta]) \otimes f(X_{t_n}^\dagger, \theta) \right] \times \quad (11a)$$

$$\left(\gamma + \Delta t \pi_n \left[\left(f(X_{t_n}^\dagger, \theta) - \pi_n[f(X_{t_n}^\dagger, \theta)] \right) \otimes f(X_{t_n}^\dagger, \theta) \right] \right)^{-1}. \quad (11b)$$

It is straightforward to combine this time discretisation with the Monte-Carlo approximation (5) in order to obtain a complete numerical implementation of the EnKBF.

Remark 2 The rough path analysis of the EnKBF presented in [9] is based on a Stratonovich reformulation of (3) and its appropriate time discretisation. Here we follow the Itô/Euler–Maruyama formulation of the data-driven term in (3),

$$\int_0^T g(X_t^\dagger, t) dX_t^\dagger = \lim_{\Delta t \rightarrow 0} \sum_{i=1}^L g(X_{t_n}^\dagger, t_n)(X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) \tag{12}$$

for any continuous function $g(x, t)$ and $\Delta t = T/L$, as it corresponds to standard implementation of the EnKBF and is easier to analyse in the context of this paper.

The EnKBF provides only an approximate solution to the Bayesian inference problem for general nonlinear $f(x, \theta)$. However, it becomes exact in the mean-field limit for affine drift functions $f(x, \theta) = \theta Ax + Bx + c$.

Example 1 Consider the stochastic partial differential equation

$$\partial_t u = -U \partial_y u + \rho \partial_y^2 u + \dot{\mathcal{W}} \tag{13}$$

over a periodic spatial domain $y \in [0, L)$, where $\mathcal{W}(t, y)$ denotes space-time white noise, $U \in \mathbb{R}$, and $\rho > 0$ are given parameters. A standard finite-difference discretisation in space with d grid points and mesh-size Δy leads to a linear system of SDEs of the form

$$d\mathbf{u}_t = -(UD + \rho DD^T)\mathbf{u}_t dt + \Delta y^{-1/2} dW_t, \tag{14}$$

where $\mathbf{u}_t \in \mathbb{R}^d$ denotes the vector of grid approximations at time t , $D \in \mathbb{R}^{d \times d}$ a finite difference approximation of the spatial derivative ∂_y , and W_t the standard d -dimensional Brownian motion. We can now set $X_t = \mathbf{u}_t$, $\gamma = \Delta y^{-1}$ and identify either $\theta = U$ or $\theta = \rho$ as the unknown parameter in order to obtain an SDE of the form (1).

In this note, we further simplify our given inference problem to the case

$$f(x, \theta) = \theta Ax, \tag{15}$$

where $A \in \mathbb{R}^{d \times d}$ is a normal matrix with eigenvalues in the left half plane. That is $\sigma(A) \subset \mathbb{C}_-$. The reference parameter value is set to $\theta^\dagger = 1$. Hence the SDE (2) possesses a Gaussian invariant measure with mean zero and covariance matrix

$$C = -\gamma(A + A^T)^{-1}. \tag{16}$$

We assume from now on that the observations X_t^\dagger are realisations of (2) with initial condition $X_0^\dagger \sim \mathcal{N}(0, C)$.

Under these assumptions, the EnKBF (3) simplifies drastically, and we obtain

$$d\Theta_t = \frac{\sigma_t}{\gamma} (AX_t^\dagger)^\top dI_t, \quad (17a)$$

$$dI_t = dX_t^\dagger - \frac{1}{2} (\Theta_t + \pi_t[\theta]) AX_t^\dagger dt, \quad (17b)$$

with variance

$$\sigma_t = \pi_t \left[(\theta - \pi_t[\theta])^2 \right]. \quad (18)$$

Remark 3 For completeness, we state the corresponding formulation for the stochastic gradient descent approach (9):

$$d\theta_t = \frac{\alpha_t}{\gamma} (AX_t^\dagger)^\top d\tilde{I}_t, \quad (19a)$$

$$d\tilde{I}_t = dX_t^\dagger - \theta_t AX_t^\dagger dt. \quad (19b)$$

We find that the learning rate α_t takes the role of the variance σ_t in (17). However, we emphasise again that the same pathwise stochastic integrals arise from both formulations, and therefore, the same robustness issue of the resulting estimators $\theta_t, t > 0$, arises.

Similarly, the discrete-time mean-field EnKBF (10) reduces to

$$\Theta_{n+1} = \Theta_n + K_n \left\{ (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) - \frac{1}{2} (\Theta_n + \pi_n[\theta]) AX_{t_n}^\dagger \Delta t \right\} \quad (20)$$

with Kalman gain

$$K_n = \sigma_n (AX_{t_n}^\dagger)^\top \left(\gamma + \Delta t \sigma_n (AX_{t_n}^\dagger)^\top AX_{t_n}^\dagger \right)^{-1}. \quad (21)$$

Furthermore, since $X_t^\dagger \sim \mathcal{N}(0, C)$,

$$(AX_t^\dagger)^\top AX_t^\dagger = (A^\top A) : (X_t^\dagger \otimes X_t^\dagger) \approx (A^\top A) : C \quad (22)$$

for $d \gg 1$, and we may simplify the Kalman gain to

$$K_n = \sigma_n (AX_{t_n}^\dagger)^\top \left(\gamma + \Delta t \sigma_n (A^\top A) : C \right)^{-1}. \quad (23)$$

Here we have used the notation $A : B = \text{tr}(A^T B)$ to denote the Frobenius inner product of two matrices $A, B \in \mathbb{R}^{d \times d}$. The approximation (22) becomes exact in the limit $d \rightarrow \infty$, which we will frequently assume in the following section. Please note that

$$K_n = \frac{\sigma_n}{\gamma} (AX_{t_n}^\dagger)^T + \mathcal{O}(\Delta t) \tag{24}$$

under the stated assumptions.

Remark 4 The Stratonovitch reformulation of (17) replaces (17a) by

$$d\Theta_t = \frac{\sigma_t}{\gamma} \left\{ (AX_t^\dagger)^T \circ dI_t - \frac{\gamma}{2} \text{tr}(A) dt \right\}. \tag{25}$$

The innovation I_t remains as before. See Appendix B of [9] for more details. An appropriate time discretisation of the innovation-driven term replaces the Kalman gain (21) by

$$K_{n+1/2} = \sigma_n (AX_{t_{n+1/2}}^\dagger)^T \left(\gamma + \Delta t \sigma_n (AX_{t_{n+1/2}}^\dagger)^T AX_{t_{n+1/2}}^\dagger \right)^{-1}, \tag{26}$$

where

$$X_{t_{n+1/2}}^\dagger = \frac{1}{2} (X_{t_n}^\dagger + X_{t_{n+1}}^\dagger). \tag{27}$$

Please note that a midpoint discretisation of the data-driven term in (25) results in

$$(AX_{t_{n+1/2}}^\dagger)^T (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) = (AX_{t_n}^\dagger)^T (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) + \tag{28a}$$

$$\frac{1}{2} A^T : (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) \otimes (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) \tag{28b}$$

and that

$$\frac{1}{2} A^T : (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) \otimes (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) \approx \frac{\Delta t \gamma}{2} \text{tr}(A), \tag{29}$$

which justifies the additional drift term in (25). A precise meaning of the approximation in (29) will be given in Remark 5 below.

Alternatively, if one wishes to explicitly utilise the availability of continuous-time data X_t^\dagger , one could apply the following variant of (20):

$$\Theta_{n+1} = \Theta_n + \frac{\sigma_n}{\gamma} \int_{t_n}^{t_{n+1}} (AX_t^\dagger)^T dX_t^\dagger - \frac{1}{2} K_n A X_{t_n}^\dagger (\Theta_n + \pi_n[\theta]) \Delta t, \tag{30}$$

and following the Itô/Euler–Maruyama approximation (12), discretise the integral with a small inner step-size $\Delta\tau = \Delta t/L$, $L \gg 1$; that is,

$$\int_{t_n}^{t_{n+1}} (AX_t^\dagger)^\top dX_t^\dagger \approx \sum_{l=0}^{L-1} (AX_{\tau_l}^\dagger)^\top (X_{\tau_{l+1}}^\dagger - X_{\tau_l}^\dagger) \quad (31)$$

with $\tau_l = t_n + l\Delta\tau$. We note that

$$\sum_{l=0}^{L-1} (AX_{\tau_l}^\dagger)^\top (X_{\tau_{l+1}}^\dagger - X_{\tau_l}^\dagger) = (AX_{t_n}^\dagger)^\top (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) + \quad (32a)$$

$$A^\top : \left(\sum_{l=0}^{L-1} (X_{\tau_l}^\dagger - X_{t_n}^\dagger) \otimes (X_{\tau_{l+1}}^\dagger - X_{\tau_l}^\dagger) \right), \quad (32b)$$

which is at the heart of rough path analysis [13] and which we utilise in the following section.

3 Frequentist Analysis

It is well-known that the second-order contribution in (32) leads to a discontinuous dependence of the integral on the observed X_t^\dagger in the uniform norm topology on the space of continuous functions. Rough path theory fixes this problem by defining appropriately extended topologies and has been extended to the EnKBF in [9]. In this section, we complement the path-wise analysis from [9] by an analysis of the impact of second-order contribution on the EnKBF (17) from a frequentist perspective, which analyses the behaviour of EnKBF over all possible observations X_t^\dagger subject to (2). In other words, one switches from a strong solution concept to a weak one. While we assume that the observations satisfy (2), throughout this section, we will analyse the impact of a perturbed observation process on the EnKBF in Sect. 4.

We first derive evolution equations for the conditional mean and variance under the assumption that Θ_0 is Gaussian distributed with given prior mean m_{prior} and variance σ_{prior} . It follows directly from (17) that the conditional mean $\mu_t = \pi_t[\theta]$, that is the mean of Θ_t , satisfies the SDE

$$d\mu_t = \frac{\sigma_t}{\gamma} \left((AX_t^\dagger)^\top dX_t^\dagger - \mu_t (A^\top A) : (X_t^\dagger \otimes X_t^\dagger) dt \right), \quad (33)$$

which simplifies to

$$d\mu_t = \frac{\sigma_t}{\gamma} \left((AX_t^\dagger)^\top dX_t^\dagger - \mu_t (A^\top A) : C dt \right), \quad (34)$$

under the approximation (22). The initial condition is $\mu_0 = m_{\text{prior}}$. The evolution equation for the conditional variance, that is the variance of Θ_t , is given by

$$\frac{d}{dt}\sigma_t = -\frac{\sigma_t^2}{\gamma} (A^\top A) : (X_t^\dagger \otimes X_t^\dagger) \quad (35)$$

with initial condition $\sigma_0 = \sigma_{\text{prior}}$ and which again reduces to

$$\frac{d}{dt}\sigma_t = -\frac{\sigma_t^2}{\gamma} (A^\top A) : C \quad (36)$$

under the approximation (22).

We now perform a frequentist analysis of the estimator μ_t defined by (34) and (36), that is, we perform a weak analysis of the SDE (34) in terms of the first two moments of μ_t [29]. In the first step, we take the expectation of (34) over all realisations X_t^\dagger of the SDE (2), which we denote by

$$m_t := \mathbb{E}^\dagger[\mu_t]. \quad (37)$$

The associated evolution equation is given by

$$\frac{d}{dt}m_t = \frac{\sigma_t}{\gamma} (A^\top A) : \mathbb{E}^\dagger \left[X_t^\dagger \otimes X_t^\dagger \right] - \frac{\sigma_t}{\gamma} (A^\top A) : C m_t, \quad (38)$$

which reduces to

$$\frac{d}{dt}m_t = \frac{\sigma_t}{\gamma} (A^\top A) : C (1 - m_t) = \sigma_t (A^\top A) : (A + A^\top)^{-1} (1 - m_t). \quad (39)$$

In the second step, we also look at the frequentist variance

$$p_t := \mathbb{E}^\dagger[(\mu_t - m_t)^2]. \quad (40)$$

Using

$$d(\mu_t - m_t) = \frac{\sigma_t}{\gamma} \left\{ (A^\top A) : (X_t^\dagger \otimes X_t^\dagger - C) dt + \gamma^{1/2} (AX_t^\dagger)^\top dW_t^\dagger \right\} - \quad (41a)$$

$$\frac{\sigma_t}{\gamma} (A^\top A) : C (\mu_t - m_t) dt, \quad (41b)$$

we obtain

$$\frac{d}{dt} p_t = -\frac{\sigma_t}{\gamma} (A^T A) : C (2p_t - \sigma_t) + \quad (42a)$$

$$\frac{2\sigma_t}{\gamma} (A^T A) : \mathbb{E}^\dagger \left[(X_t^\dagger \otimes X_t^\dagger - C) (\mu_t - m_t) \right], \quad (42b)$$

which we simplify to

$$\frac{d}{dt} p_t = \frac{\sigma_t}{\gamma} (A^T A) : C (\sigma_t - 2p_t) = \sigma_t (A^T A) : (A + A^T)^{-1} (\sigma_t - 2p_t) \quad (43)$$

under the approximation (22). The initial conditions are $m_0 = m_{\text{prior}}$ and $p_0 = 0$, respectively. We note that the differential equations (36) and (43) are explicitly solvable. For example, it holds that

$$\sigma_t = \frac{\sigma_0}{1 + (A^T A) : (A^T + A)^{-1} \sigma_0 t} \quad (44)$$

and one finds that $\sigma_t \sim 1/((A^T A) : (A^T + A)^{-1} t)$ for $t \gg 1$. It can also be shown that $p_t \leq \sigma_t$ for all $t \geq 0$. Furthermore, this analysis suggests that the learning rate in the stochastic gradient descent formulation (19) should be chosen as

$$\alpha_t = \min \left\{ \bar{\alpha}, \frac{1}{(A^T A) : (A^T + A)^{-1} t} \right\}, \quad (45)$$

where $\bar{\alpha} > 0$ denotes an initial learning rate; for example $\bar{\alpha} = \sigma_0$.

We finally conduct a formal analysis of the ensemble Kalman filter time-stepping (20) and demonstrate that the method is first-order accurate with regard to the implied frequentist mean m_t . We recall (24) and conclude from (20) that the implied update on the variance σ_n satisfies

$$\sigma_{n+1} = \sigma_n - \frac{\sigma_n^2}{\gamma} (A^T A) : C \Delta t + \mathcal{O}(\Delta t^2), \quad (46)$$

which provides a first-order approximation to (36).

We next analyse the evolution equation (34) for the conditional mean μ_t and its numerical approximation

$$\mu_{n+1} = \mu_n + K_n \left\{ (X_{t_{n+1}}^\dagger - X_{t_n}^\dagger) - \mu_n A X_{t_n}^\dagger \Delta t \right\} \quad (47)$$

arising from (20). Here we follow [13] in order to analyse the impact of the data X_t^\dagger on the estimator. An in-depth theoretical treatment can be found in [9].

Comparing (47) to (34) and utilising (24), we find that the key quantity of interest is

$$J_{t_n, t_{n+1}}^\dagger := \int_{t_n}^{t_{n+1}} (AX_t^\dagger)^\top dX_t^\dagger, \tag{48}$$

which we can rewrite as

$$J_{t_n, t_{n+1}}^\dagger = A^\top : (X_{t_n}^\dagger \otimes X_{t_n, t_{n+1}}^\dagger) + A^\top : \mathbb{X}_{t_n, t_{n+1}}^\dagger. \tag{49}$$

Here, motivated by (32) and following standard rough path notation, we have used

$$X_{t_n, t_{n+1}}^\dagger := X_{t_{n+1}}^\dagger - X_{t_n}^\dagger \tag{50}$$

and the second-order iterated Itô integral

$$\mathbb{X}_{t_n, t_{n+1}}^\dagger := \int_{t_n}^{t_{n+1}} (X_t^\dagger - X_{t_n}^\dagger) \otimes dX_t^\dagger. \tag{51}$$

The difference between the integral (48) and its corresponding approximation in (47) is provided by $A^\top : \mathbb{X}_{t_n, t_{n+1}}^\dagger$ plus higher-order terms arising from (24). The iterated integral $\mathbb{X}_{t_n, t_{n+1}}^\dagger$ becomes a random variable from the frequentist perspective. Taking note of (2), we find that the drift, $f(x) = Ax$, contributes with terms of order $\mathcal{O}(\Delta t^2)$ to $\mathbb{X}_{t_n, t_{n+1}}^\dagger$ and the expected value of $\mathbb{X}_{t_n, t_{n+1}}^\dagger$ therefore satisfies

$$\mathbb{E}^\dagger[\mathbb{X}_{t_n, t_{n+1}}^\dagger] = \mathcal{O}(\Delta t^2), \tag{52}$$

since $\mathbb{E}^\dagger[W_{t_n, \tau}^\dagger] = 0$ for $\tau > t_n$, and

$$\mathbb{E}^\dagger[\mathbb{W}_{t_n, t_{n+1}}^\dagger] = \frac{1}{2} \mathbb{E}^\dagger[W_{t_n, t_{n+1}}^\dagger \otimes W_{t_n, t_{n+1}}^\dagger - [W_{t_n}^\dagger, W_{t_n, t_{n+1}}^\dagger]] - \frac{\Delta t}{2} I = 0, \tag{53}$$

where we have introduced the commutator

$$[W_{t_n}^\dagger, W_{t_n, t_{n+1}}^\dagger] := W_{t_n}^\dagger \otimes W_{t_n, t_{n+1}}^\dagger - W_{t_n, t_{n+1}}^\dagger \otimes W_{t_n}^\dagger. \tag{54}$$

Hence we find that, while (47) is not a first-order (strong) approximation of the SDE (34), the approximation becomes first-order in m_t when averaged over realisations X_t^\dagger of the SDE (2). More precisely, one obtains

$$\mathbb{E}^\dagger[J_{t_n, t_{n+1}}^\dagger] = (A^\top A) : C \Delta t + \mathcal{O}(\Delta t^2). \tag{55}$$

We note that the modified scheme (30) leads to the same time evolution in the variance σ_n while the update in μ_n is changed to

$$\mu_{n+1} = \mu_n + \frac{\sigma_n}{\gamma} \int_{t_n}^{t_{n+1}} (AX_t^\dagger)^\top dX_t^\dagger - K_n AX_{t_n}^\dagger \mu_n \Delta t. \tag{56}$$

This modification results in a more accurate evolution in the conditional mean μ_n , but because of (52) it does not impact to leading order the evolution of the underlying frequentist mean, $m_n = \mathbb{E}^\dagger[\mu_n]$. We summarise our findings in the following proposition.

Proposition 1 *The discrete-time EnKBF implementations (20) and (30) both provide first-order approximations to the time evolution of the frequentist mean, m_t , and the frequentist variance, p_t . In other words, both methods converge weakly with order one.*

We also note that the frequentist uncertainty is essentially data-independent and depends only on the time window $[0, T]$ over which the data gets observed. Hence, for fixed observation interval $[0, T]$, it makes sense to choose the step-size Δt such that the discretisation error (bias) remains on the same order of magnitude as $p_T^{1/2} \approx \sigma_T^{1/2}$. Selecting a much smaller step-size would not significantly reduce the frequentist estimation error in the conditional estimator μ_T .

Remark 5 We can now give a precise reformulation of the approximation (29):

$$\frac{1}{2} \mathbb{E}^\dagger \left[A^\top : (X_{t_n, t_{n+1}}^\dagger \otimes X_{t_n, t_{n+1}}^\dagger) \right] = \frac{\Delta t \gamma}{2} \text{tr}(A) + \mathcal{O}(\Delta t^2), \tag{57}$$

which is at the heart of the Stratonovich formulation (25) of the EnKFB [9].

4 Multi-Scale Data

We now have all the material in place to study the dependency of the EnKBF estimator on a set of observations $X_t^{(\epsilon)}$, $\epsilon > 0$, which approach the theoretical X_t^\dagger with respect to the uniform norm topology on the space of continuous functions as $\epsilon \rightarrow 0$. Since the second-order contribution in (32), that is (51), does not depend continuously on such perturbations, we demonstrate in this section that a systematic bias arises in the EnKBF. Furthermore, we show how the bias can be eliminated either via subsampling the data, which effectively amounts to ignoring these second-order contributions, or via an appropriate correction term, which ensures a continuous dependence on observations $X_t^{(\epsilon)}$ with respect to the uniform norm topology. More specifically, we investigate the impact of a possible discrepancy between the SDE model (1), for which we aim to estimate the parameter θ , and the data generating SDE (2). We therefore replace (2) by the following two-scale SDE [17]:

$$dX_t^{(\epsilon)} = AX_t^{(\epsilon)} dt + \frac{\gamma^{1/2}}{\epsilon} MP_t^{(\epsilon)} dt, \quad (58a)$$

$$dP_t^{(\epsilon)} = -\frac{1}{\epsilon} MP_t^{(\epsilon)} dt + dW_t^\dagger, \quad (58b)$$

where

$$M = \begin{pmatrix} 1 & \beta \\ -\beta & 1 \end{pmatrix}, \quad (59)$$

$\beta = 2$ and $\epsilon = 0.01$. The dimension of state space is $d = 2$ throughout this section. While we restrict here to the simple two-scale model (58), similar scenarios can arise from deterministic fast-slow systems [24, 7].

The associated EnKBF mean-field equations in the parameter Θ_t , which we now denote by $\Theta_t^{(\epsilon)}$ in order to explicitly record its dependence on the scale parameter $\epsilon \ll 1$, become

$$d\Theta_t^{(\epsilon)} = \frac{\sigma_t^{(\epsilon)}}{\gamma} (AX_t^{(\epsilon)})^T dI_t^{(\epsilon)}, \quad (60a)$$

$$dI_t^{(\epsilon)} = dX_t^{(\epsilon)} - \frac{1}{2} \left(\Theta_t^{(\epsilon)} + \pi_t^{(\epsilon)}[\theta] \right) AX_t^{(\epsilon)} dt, \quad (60b)$$

with variance

$$\sigma_t^{(\epsilon)} = \pi_t^{(\epsilon)} \left[(\theta - \pi_t^{(\epsilon)}[\theta])^2 \right] \quad (61)$$

and $\Theta_t^{(\epsilon)} \sim \pi_t^{(\epsilon)}$. The discrete-time mean-field EnKBF (20) turns into

$$\Theta_{n+1}^{(\epsilon)} = \Theta_n^{(\epsilon)} + K_n^{(\epsilon)} \left\{ \left(X_{n+1}^{(\epsilon)} - X_n^{(\epsilon)} \right) - \frac{1}{2} \left(\Theta_n^{(\epsilon)} + \pi_n^{(\epsilon)}[\theta] \right) AX_n^{(\epsilon)} \Delta t \right\} \quad (62)$$

with Kalman gain

$$K_n^{(\epsilon)} = \sigma_n^{(\epsilon)} (AX_{t_n}^{(\epsilon)})^T \left(\gamma + \Delta t \sigma_n^{(\epsilon)} (AX_{t_n}^{(\epsilon)})^T AX_{t_n}^{(\epsilon)} \right)^{-1}. \quad (63)$$

We also consider the appropriately modified scheme (30):

$$\Theta_{n+1}^{(\epsilon)} = \Theta_n^{(\epsilon)} + \frac{\sigma_n^{(\epsilon)}}{\gamma} \int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)} - \frac{1}{2} K_n^{(\epsilon)} AX_{t_n}^{(\epsilon)} \left(\Theta_n^{(\epsilon)} + \pi_n^{(\epsilon)}[\theta] \right) \Delta t. \quad (64)$$

In order to understand the impact of the modified data generating process on the two mean-field EnKBF formulations (62) and (64), respectively, we follow [17] and investigate the difference between $X_t^{(\epsilon)}$ and X_t^\dagger :




Fig. 1 SDE driven by mathematical vs. physical Brownian motion ($\epsilon = 0.01$). The top panel displays both X_t^\dagger (blue) and $X_t^{(\epsilon)}$ (red) over the long time interval $t \in [0, 10]$, while the lower panel provides a zoomed in perspective over the interval $t \in [0, 1]$

$$d(X_t^{(\epsilon)} - X_t^\dagger) = A(X_t^{(\epsilon)} - X_t^\dagger)dt + \frac{\gamma^{1/2}}{\epsilon} M P_t^{(\epsilon)} dt - \gamma^{1/2} dW_t^\dagger \tag{65a}$$

$$= A(X_t^{(\epsilon)} - X_t^\dagger)dt - \gamma^{1/2} dP_t^{(\epsilon)}. \tag{65b}$$

When $P_t^{(\epsilon)}$ is stationary, it is Gaussian with mean zero and covariance

$$\mathbb{E}_{\text{stat}} [P_t^{(\epsilon)} \otimes P_t^{(\epsilon)}] = \epsilon (M + M^T)^{-1} = \frac{\epsilon}{2} I. \tag{66}$$

Hence $P_t^{(\epsilon)} \rightarrow 0$ as $\epsilon \rightarrow 0$ and also

$$X_t^{(\epsilon)} \rightarrow X_t^\dagger \tag{67}$$

in L^2 uniformly in t , provided $\sigma(A) \subset \mathbb{C}_-$ and $X_0^{(\epsilon)} = X_0^\dagger$. This is illustrated in Fig. 1.

In order to investigate the problem further, we study the integral

$$J_{t_n, t_{n+1}}^{(\epsilon)} := \int_{t_n}^{t_{n+1}} (A X_t^{(\epsilon)})^T dX_t^{(\epsilon)} \tag{68}$$

and its relation to (48). As for (48), we can rewrite (68) as

$$J_{t_n, t_{n+1}}^{(\epsilon)} = A^T : (X_{t_n}^{(\epsilon)} \otimes X_{t_{n+1}}^{(\epsilon)}) + A^T : \mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)}. \tag{69}$$

We now investigate the limit of the second-order iterated integral

$$\mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)} = \int_{t_n}^{t_{n+1}} X_{t_n, t}^{(\epsilon)} \otimes dX_t^{(\epsilon)} \tag{70a}$$

$$= \frac{1}{2} X_{t_n, t_{n+1}}^{(\epsilon)} \otimes X_{t_n, t_{n+1}}^{(\epsilon)} - \frac{1}{2} \int_{t_n}^{t_{n+1}} [X_{t_n, t}^{(\epsilon)}, dX_t^{(\epsilon)}] \quad (70b)$$

as $\epsilon \rightarrow 0$ [17]. Here $[\cdot, \cdot]$ denotes the commutator defined by (54).

Proposition 2 *The second-order iterated integral $\mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)}$ satisfies*

$$\lim_{\epsilon \rightarrow 0} \mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)} = \mathbb{X}_{t_n, t_{n+1}}^\dagger + \frac{\Delta t}{2} \gamma M \quad (71)$$

Proof The proof follows [17] and can be summarised as follows:

$$\mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)} = \int_{t_n}^{t_{n+1}} X_{t_n, t}^{(\epsilon)} \otimes dX_t^{(\epsilon)} \quad (72a)$$

$$\rightarrow \int_{t_n}^{t_{n+1}} X_{t_n, t}^\dagger \otimes dX_t^\dagger - \gamma^{1/2} \int_{t_n}^{t_{n+1}} X_{t_n, t}^{(\epsilon)} \otimes dP_t^{(\epsilon)} \quad (72b)$$

$$= \mathbb{X}_{t_n, t_{n+1}}^\dagger - \gamma^{1/2} X_{t_n, t_{n+1}}^{(\epsilon)} \otimes P_{t_{n+1}}^{(\epsilon)} + \gamma^{1/2} \int_{t_n}^{t_{n+1}} dX_t^{(\epsilon)} \otimes P_t^{(\epsilon)} \quad (72c)$$

$$\rightarrow \mathbb{X}_{t_n, t_{n+1}}^\dagger + \gamma^{1/2} \int_{t_n}^{t_{n+1}} \left\{ AX_t^{(\epsilon)} + \frac{\gamma^{1/2}}{\epsilon} M P_t^{(\epsilon)} \right\} \otimes P_t^{(\epsilon)} dt \quad (72d)$$

$$\rightarrow \mathbb{X}_{t_n, t_{n+1}}^\dagger + \frac{\Delta t}{\epsilon} \gamma M \mathbb{E}_{\text{stat}} \left[P_{t_n}^{(\epsilon)} \otimes P_{t_n}^{(\epsilon)} \right] \quad (72e)$$

$$= \mathbb{X}_{t_n, t_{n+1}}^\dagger + \frac{\Delta t}{2} \gamma M. \quad (72f)$$

As discussed in detail in [9] already, Proposition 2 implies that the scheme (64) does not, in general, converge to the scheme (64) as $\epsilon \rightarrow 0$ since

$$J_{t_n, t_{n+1}}^\dagger = \lim_{\epsilon \rightarrow 0} J_{t_n, t_{n+1}}^{(\epsilon)} - \frac{\Delta t}{2} \gamma A^T : M. \quad (73)$$

This observation suggests the following modification

$$\Theta_{n+1}^{(\epsilon)} = \Theta_n^{(\epsilon)} + \frac{\sigma_n^{(\epsilon)}}{\gamma} \int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)} - \frac{\Delta t}{2} \sigma_n^{(\epsilon)} A^T : M - \quad (74a)$$

$$\frac{1}{2} K_n^{(\epsilon)} AX_{t_n}^{(\epsilon)} \left(\Theta_n^{(\epsilon)} + \pi_n^{(\epsilon)} [\theta] \right) \Delta t \quad (74b)$$

to (64). Please note that it follows from (70) that

$$\int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)} = A^T : \left(X_{t_{n+1/2}}^{(\epsilon)} \otimes X_{t_n, t_{n+1}}^{(\epsilon)} - \frac{1}{2} \int_{t_n}^{t_{n+1}} [X_{t_n, t}^{(\epsilon)}, dX_t^{(\epsilon)}] \right). \tag{75}$$

Proposition 3 *The discrete-time EnKBF (62) converges to (20) for fixed Δt as $\epsilon \rightarrow 0$. Similarly, (74) converges to (30) under the same limit.*

Proof The first statement follows from $\sigma_n^{(\epsilon)} = \sigma_n$, the limiting behaviour (67), and

$$\lim_{\epsilon \rightarrow 0} K_n^{(\epsilon)} = K_n. \tag{76}$$

The second statement additionally requires (73) to be substituted into (74) when taking the limit $\epsilon \rightarrow 0$.

Remark 6 The analogous adaptation of (74) to the gradient descent formulation (19) with X_t^\dagger replaced by $X_t^{(\epsilon)}$ becomes

$$\theta_{n+1}^{(\epsilon)} = \theta_n^{(\epsilon)} + \frac{\alpha_{t_n}}{\gamma} \left(\int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)} - \frac{\gamma \Delta t}{2} A^T : M - \tag{77a}$$

$$\theta_n^{(\epsilon)} (AX_{t_n}^{(\epsilon)})^T AX_{t_n}^{(\epsilon)} \Delta t \right). \tag{77b}$$

Alternatively, subsampling the data can be applied which leads to the simpler formulation

$$\theta_{n+1}^{(\epsilon)} = \theta_n^{(\epsilon)} + \frac{\alpha_{t_n}}{\gamma} (AX_{t_n}^{(\epsilon)})^T \left((X_{t_{n+1}}^{(\epsilon)} - X_{t_n}^{(\epsilon)}) - \theta_n^{(\epsilon)} AX_{t_n}^{(\epsilon)} \Delta t \right). \tag{78}$$

Remark 7 A two-scale SDE, closely related to (58), has been investigated in [8] in terms of the time integrated autocorrelation function of $P_t^{(\epsilon)}$ and modified stochastic integrals. In our case, the modified quadrature rule, here denoted by \diamond , has to satisfy

$$\int_{t_n}^{t_{n+1}} (AX_t^\dagger)^T \diamond dX_t^\dagger = \lim_{\epsilon \rightarrow 0} \int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)}, \tag{79}$$

and it is therefore related to the standard Itô integral via

$$\int_{t_n}^{t_{n+1}} (AX_t^\dagger)^T \diamond dX_t^\dagger = \int_{t_n}^{t_{n+1}} (AX_t^\dagger)^T dX_t^\dagger + \frac{\Delta t \gamma}{2} A^T : M. \tag{80}$$

Hence M plays the role of the integrated autocorrelation function of $P_t^{(\epsilon)}$ in our approach. We note that the modified quadrature rule reduces to the standard Stratonovitch integral if either $\beta = 0$ in (59) or A is symmetric. While the results from [8] could, therefore, also be used as a starting point for discussing the induced estimation bias, practical implementations would still require knowledge of the

integrated autocorrelation function of $P_t^{(\epsilon)}$ or, equivalently, the estimation of M in addition to observing $X_t^{(\epsilon)}$. We address this aspect next.

The numerical implementation of (74) requires an estimator for the generally unknown M in (73). This task is challenging as we only have access to $X_t^{(\epsilon)}$ without any explicit knowledge of the underlying generating process (58). While the estimator proposed in [9] is based on the idea of subsampling the data, the frequentist perspective taken in this note suggests the alternative estimator M_{est} defined by

$$\frac{\Delta t \gamma}{2} M_{\text{est}} = \mathbb{E}^\dagger[\mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)}], \tag{81}$$

which follows from (72f) and (52). That is, $\mathbb{E}^\dagger[\mathbb{X}_{t_n, t_{n+1}}^\dagger] = \mathcal{O}(\Delta t^2)$ for Δt sufficiently small. Note that second-order iterated integral $X_{t_n, t_{n+1}}^{(\epsilon)}$ satisfies (70) and is therefore easy to compute. In practice, the frequentist expectation value can be replaced by an approximation along a given single observation path $X_t^{(\epsilon)}$, $t \in [0, T]$, under the assumption of ergodicity.

An appropriate choice of the outer or sub-sampling step-size Δt [27] constitutes an important aspect for the practical implementation of the EnKBF formulation (62) for finite values of $\epsilon > 0$ [26]. Consistency of the second-order iterated integrals [13] implies

$$\mathbb{X}_{t_n, t_{n+2}}^{(\epsilon)} = \mathbb{X}_{t_n, t_{n+1}}^{(\epsilon)} + \mathbb{X}_{t_{n+1}, t_{n+2}}^{(\epsilon)} + X_{t_n, t_{n+1}}^{(\epsilon)} \otimes X_{t_{n+1}, t_{n+2}}^{(\epsilon)}. \tag{82}$$

A sensible choice of Δt is dictated by

$$\mathbb{E}^\dagger \left[X_{t_n, t_{n+1}}^{(\epsilon)} \otimes X_{t_{n+1}, t_{n+2}}^{(\epsilon)} \right] = \mathcal{O}(\Delta t^2), \tag{83}$$

that is, the sub-sampled data $X_{t_n}^{(\epsilon)}$ behaves to leading order like solution increments from the reference model (2) at scale Δt independent of the specific value of ϵ . Note that, on the other hand,

$$\mathbb{E}^\dagger \left[X_{\tau_l, \tau_{l+1}}^{(\epsilon)} \otimes X_{\tau_{l+1}, \tau_{l+2}}^{(\epsilon)} \right] = \mathcal{O}(\epsilon^{-1} \Delta \tau^2) \tag{84}$$

for an inner step-size $\Delta \tau \sim \epsilon$. In other words, a suitable step-size $\Delta t > 0$ can be defined by making

$$h(\Delta t) := \Delta t^{-2} \left\| \mathbb{E}^\dagger \left[X_{t_n, t_{n+1}}^{(\epsilon)} \otimes X_{t_{n+1}, t_{n+2}}^{(\epsilon)} \right] \right\| \tag{85}$$

as small as possible while still guaranteeing an accurate numerical approximation in (62).

Remark 8 The choice of the outer time step Δt is less critical for the EnKBF formulation (74) since it does not rely on sub-sampling the data and is robust with regard to perturbations in the data provided the appropriate M is explicitly available or has been estimated from the available data using (81). Furthermore, if A is symmetric, then it follows from (75) and the skew-symmetry of the commutator $[\cdot, \cdot]$ that

$$\int_{t_n}^{t_{n+1}} (AX_t^{(\epsilon)})^T dX_t^{(\epsilon)} = A : \left(X_{t_{n+1/2}}^{(\epsilon)} \otimes X_{t_n, t_{n+1}}^{(\epsilon)} \right), \quad (86)$$

which can be used in (74). The same simplification arises when M is symmetric. This insight is at the heart of the geometric rough path approach followed in [9] and which starts from the Stratonovich formulation (25) of the EnKBF. See also [28] on the convergence of Wong–Zakai approximations for stochastic differential equations. In all other cases, a more refined numerical approximation of the data-driven integral in (74) is necessary; such as, for example, (31). For that reason, we rely on the Itô/Euler–Maruyama interpretation of (68) in this note instead, that is the approximation (12).

5 Numerical Example

We consider the linear SDE (2) with $\gamma = 1$ and

$$A = \frac{-1}{2} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}. \quad (87)$$

We find that $C = I$ and $A^T A = 1/2I$. Hence $(A^T A) : C = 1$, and the posterior variance simply satisfies $\sigma_t = \sigma_0/(1 + \sigma_0 t)$ according to (44). We set $m_{\text{prior}} = 0$ and $\sigma_{\text{prior}} = 4$ for the Gaussian prior distribution of θ_0 , and the observation interval is $[0, T]$ with $T = 6$. We find that $\sigma_T = 0.16$. Solving (39) for given σ_t with initial condition $m_0 = 0$ yields

$$m_t = 1 - \frac{\sigma_t}{\sigma_0} \quad (88)$$

and $m_T = 0.96$. The corresponding curves are displayed in red in Fig. 2.

We implement the EnKBF schemes (20) and (30) with $t_n = n \Delta t$. The inner time-step is $\Delta \tau = 10^{-4}$ while $\Delta t = 0.06$, that is, $L = 600$. We repeat the experiment $N = 10^4$ times and compare the outcome with the predicted mean value of $m_T = 0.96$ and the posterior variance of $\sigma_T = 0.16$ in Fig. 2. The differences in the computed time evolutions of m_t and p_t are rather minor and support the idea that it is not necessary to assimilate continuous-time data beyond Δt . We

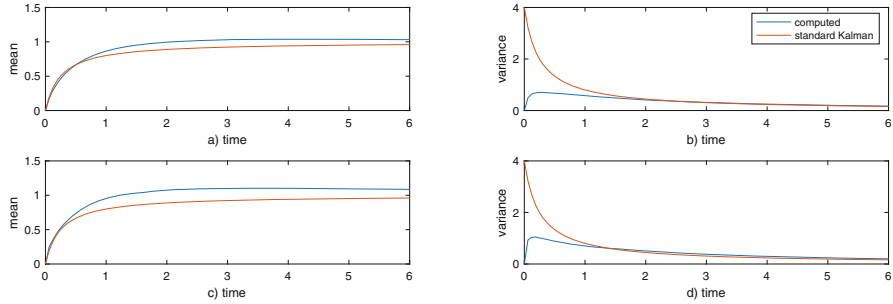


Fig. 2 (a–b) Frequentist mean, m_t and variance, p_t , from EnKBF implementation (20) with step-size $\Delta t = 0.06$; (c–d) Same results from EnKBF implementation (30) with inner time-step $\Delta \tau = \Delta t/600$. We also display the curves arising for σ_t and m_t from the standard Kalman theory using the approximation (22). Note that the posterior variance, σ_t , should provide an upper bound on the frequentist uncertainty p_t

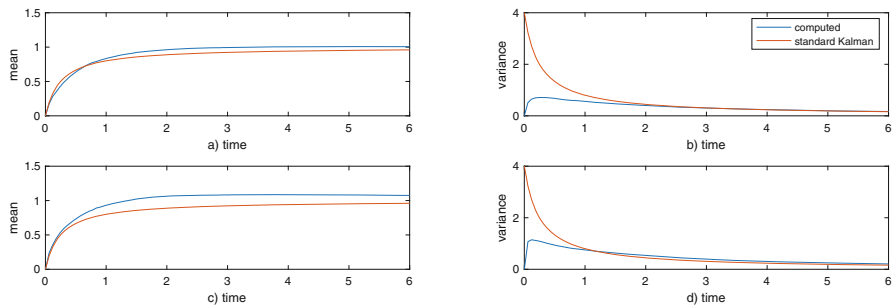


Fig. 3 Same experimental setting as in Fig. 2 but with the data now generated from the multi-scale SDE (58). Again, subsampling the data in intervals of $\Delta t = 0.06$ and high-frequency assimilation with step-size $\Delta \tau = 10^{-4}$ lead to very similar results in terms of their frequentist means and variances

also find that the simple prediction (88), based on standard Kalman filter theory, is not very accurate for this low-dimensional problem ($d = 2$). The corresponding approximation for σ_t provides, however, a good upper bound for p_t .

We now replace the data generating SDE model (2) by the multi-scale formulation (58) with $\epsilon = 0.01$ and $\beta = 2$. This parameter choice agrees with the one used in [9]. We again find that assimilating the data at the slow time-scale $\Delta t = 0.06$ leads to very similar results obtained from an assimilation at the fast time-scale $\Delta \tau = 10^{-4}$ with the EnKBF formulation (74), provided the correction term resulting from the second-order iterated integral (73) is included (See Fig. 3). We also verified numerically that $\Delta t = 0.06$ constitutes a nearly optimal step-size in the sense of making (85) sufficiently small while maintaining numerical accuracy. For example, reducing the outer step-size to $\Delta t = 0.02$ leads to $h(0.02) - h(0.06) \approx 10$ in (85).

6 Conclusions

In this follow-up note to [9], we have investigated the impact of subsampling and/or high-frequency data assimilation on the corresponding conditional mean estimators, μ_t , both for data generated from the standard SDE model and a modified multi-scale SDE. A frequentist analysis supports the basic finding that both approaches lead to comparable results provided that the systematic biases due to different second-order iterated integrals are properly accounted for. While the EnKBF is relatively easy to analyse and a full rough path approach can be avoided, extending these results to the nonlinear feedback particle filter [26, 9] will prove more challenging. Extensions to systems without a strong scale separation [4, 31] and applications to geophysical fluid dynamics [22, 12] are also of interest. In this context, the approximation quality of the proposed estimator (81) and the choice of the step-size Δt following (85) (and potentially $\Delta\tau$) will be of particular interest. Finally, while we have investigated the univariate parameter estimation problem, a semi-parametric parametrisation of the drift term f in (1), such as random feature maps [21], lead to high-dimensional parameter estimation problems and their statistics [19, 20]. This provides another fertile direction for future research.

Acknowledgments SR has been partially funded by Deutsche Forschungsgemeinschaft (DFG)—Project-ID 318763901—SFB1294 and Project-ID 235221301—SFB1114. He would also like to thank Nikolas Nüsken for many fruitful discussions on the subject of this paper.

References

1. A. Abdulle, G. Garegnani, G. A. Pavliotis, A. M. Stuart, and A. Zanoni. Drift estimation of multiscale diffusions based on filtered data. *Foundations of Computational Mathematics*, published online 2021/10/13: in press, 2021. <https://doi.org/10.1007/s10208-021-09541-9>.
2. Y. Ait-Sahalia, P. A. Mykland, and L. Zhang. How often to sample a continuous-time process in the presence of market microstructure noise. *The Review of Financial Studies*, 18: 351–416, 2005.
3. J. Amezcuca, E. Kalnay, K. Ide, and S. Reich. Ensemble transform Kalman-Bucy filters. *Q.J.R. Meteor. Soc.*, 140: 995–1004, 2014.
4. L. Arnold. Hasselmann’s program revisited: The analysis of stochasticity in deterministic climate models. In *Stochastic Climate Models*, pages 141–158. Birkhäuser Basel, 2001. <https://doi.org/10.1007/978-3-0348-8287-3>.
5. R. Azencott, A. Beri, A. Jain, and I. Timofeyev. Sub-sampling and parametric estimation for multiscale dynamics. *Communications in Mathematical Sciences*, 11: 939–970, 2013.
6. A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*, volume 60 of *Stoch. Model. Appl. Probab.* Springer, New York, 2009. <https://doi.org/10.1007/978-0-387-76896-0>.
7. P. Bálint and I. Melbourne. Statistical properties for flows with unbounded roof function, including the Lorenz attractor. *Journal of Statistical Physics*, 172: 1101–1126, 2018. <https://doi.org/10.1007/s10955-018-2093-y>.
8. S. Bo and A. Celani. White-noise limit of nonwhite nonequilibrium processes. *Physical Review E*, 88: 062150, 2013. <https://doi.org/10.1103/PhysRevE.88.062150>.

9. M. Coghi, T. Nilssen, N. Nüsken, and S. Reich. Rough McKean–Vlasov dynamics for robust ensemble Kalman filtering, 2021. arXiv:2107.06621.
10. C. Cotter and S. Reich. Ensemble filter techniques for intermittent data assimilation. *Radon Ser. Comput. Appl. Math.*, 13: 91–134, 2013. https://doi.org/10.1515/9783110282269_91.
11. D. Crisan, J. Diehl, P. K. Friz, H. Oberhauser, et al. Robust filtering: correlated noise and multidimensional observation. *The Annals of Applied Probability*, 23: 2139–2160, 2013.
12. J. Culina, S. Kravtsov, and A. H. Monahan. Stochastic parameterization schemes for use in realistic climate models. *Journal of the Atmospheric Sciences*, 68: 284 – 299, 2011. <https://doi.org/10.1175/2010JAS3509.1>.
13. A. M. Davie. Differential equations driven by rough paths: An approach via discrete approximation. *Applied Mathematics Research eXpress*, 2008, 2008. <https://doi.org/10.1093/amrx/abm009>. abm009.
14. J. Diehl, P. Friz, and H. Mai. Pathwise stability of likelihood estimators for diffusion via rough paths. *The Annals of Applied Probability*, 26: 2169–2192, 2016. <https://doi.org/10.1214/15-AAP1143>.
15. G. Evensen. *Data assimilation*. Springer-Verlag, Berlin, second edition, 2009. ISBN 978-3-642-03710-8. <https://doi.org/10.1007/978-3-642-03711-5>.
16. P. Friz and M. Hairer. *A course on rough paths*. Springer-Verlag, 2020.
17. P. Friz, P. Gassiat, and T. Lyons. Physical Brownian motion in a magnetic field as a rough path. *Transactions of the American Mathematical Society*, 367: 7939–7955, 2015.
18. A. Garbuno-Inigo, N. Nüsken, and S. Reich. Affine invariant interacting Langevin dynamics for Bayesian inference. *SIAM J. Appl. Dyn. Syst.*, 19: 1633–1658, 2020. <https://doi.org/10.1137/19M1304891>.
19. S. Ghosal and A. van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2017. <https://doi.org/10.1017/9781139029834>.
20. E. Giné and R. Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press, Cambridge, 2016. <https://doi.org/10.1017/CBO9781107337862>.
21. G. A. Gottwald and S. Reich. Supervised learning from noisy observations: Combining machine-learning techniques with data assimilation. *Physica D: Nonlinear Phenomena*, 423: 132911, 2021. ISSN 0167-2789. <https://doi.org/10.1016/j.physd.2021.132911>.
22. K. Hasselmann. Stochastic climate models Part I. Theory. *Tellus*, 28: 473–485, 1976. <https://doi.org/10.1111/j.2153-3490.1976.tb00696.x>.
23. N. Ikeda and S. Watanabe. *Stochastic differential equations and diffusion processes*. North Holland Publishing Company, Amsterdam-New York, 2nd edition, 1989.
24. D. Kelly and I. Melbourne. Deterministic homogenization for fast-slow systems with chaotic noise. *Journal of Functional Analysis*, 272: 4063–4102, 2017. <https://doi.org/10.1016/j.jfa.2017.01.015>.
25. Y. A. Kutoyants. *Statistical inference for ergodic diffusion processes*. Springer Science & Business Media, 2013.
26. N. Nüsken, S. Reich, and P. J. Rozdeba. State and parameter estimation from observed signal increments. *Entropy*, 21 (5): 505, 2019. <https://doi.org/10.3390/e21050505>.
27. A. Papavasiliou, G. Pavliotis, and A. Stuart. Maximum likelihood estimation for multiscale diffusions. *Stochastic Processes and their Applications*, 19: 3173–3210, 2009.
28. S. Pathiraja. L^2 convergence of smooth approximations of stochastic differential equations with unbounded coefficients, 2020. arXiv:2011.13009.
29. S. Reich and P. Rozdeba. Posterior contraction rates for non-parametric state and drift estimation. *Foundation of Data Science*, 2: 333–349, 2020. <https://doi.org/10.3934/fods.2020016>.
30. J. Sirignano and K. Spiliopoulos. Stochastic gradient descent in continuous time. *SIAM J. Financial Math.*, 8: 933–961, 2017. <https://doi.org/10.1137/17M1126825>.
31. J. Wouters and G. A. Gottwald. Stochastic model reduction for slow-fast systems with moderate time scale separation. *Multiscale Modeling & Simulation*, 17: 1172–1188, 2019.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

