



On the Mathematical Theory of Ensemble (Linear-Gaussian) Kalman-Bucy Filtering

Adrian N Bishop, Pierre del Moral

► To cite this version:

Adrian N Bishop, Pierre del Moral. On the Mathematical Theory of Ensemble (Linear-Gaussian) Kalman-Bucy Filtering. 2023. hal-03944397v2

HAL Id: hal-03944397

<https://inria.hal.science/hal-03944397v2>

Preprint submitted on 18 Jan 2023 (v2), last revised 7 Dec 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Mathematical Theory of Ensemble (Linear-Gaussian) Kalman-Bucy Filtering

Adrian N. Bishop^{*1} and Pierre Del Moral²

¹CSIRO; and University of Technology Sydney (UTS), Australia

²INRIA, Bordeaux Research Center, France

Abstract

The purpose of this review is to present a comprehensive overview of the theory of ensemble Kalman-Bucy filtering for continuous-time, linear-Gaussian signal and observation models. We present a system of equations that describe the flow of individual particles and the flow of the sample covariance and the sample mean in continuous-time ensemble filtering. We consider these equations and their characteristics in a number of popular ensemble Kalman filtering variants. Given these equations, we study their asymptotic convergence to the optimal Bayesian filter. We also study in detail some non-asymptotic time-uniform fluctuation, stability, and contraction results on the sample covariance and sample mean (or sample error track). We focus on testable signal/observation model conditions, and we accommodate fully unstable (latent) signal models. We discuss the relevance and importance of these results in characterising the filter's behaviour, e.g. its signal tracking performance, and we contrast these results with those in classical studies of stability in Kalman-Bucy filtering. We also provide a novel (and negative) result proving that the bootstrap particle filter cannot track even the most basic unstable latent signal, in contrast with the ensemble Kalman filter (and the optimal filter). We provide intuition for how the main results extend to nonlinear signal models and comment on their consequence on some typical filter behaviours seen in practice, e.g. catastrophic divergence.

Contents

1	Introduction	2
1.1	Purpose	3
1.2	Overview of the Main Topics and Literature	4
1.3	Aims and Contributions	7
1.4	Notation	8
2	Kalman-Bucy Filtering	9
3	Kalman-Bucy Diffusion Processes	12
4	Ensemble Kalman-Bucy Filtering	13
4.1	Vanilla Ensemble Kalman-Bucy Filter	13
4.2	'Deterministic' Ensemble Kalman-Bucy Filter	14
4.3	Transport-Inspired Ensemble Transport Filter	14
4.4	Nonlinear Ensemble Filtering in Practice	15
5	Theory in the Linear-Gaussian Setting	15
5.1	Fluctuation and Contraction Results for the Riccati Diffusion	18
5.2	Contraction Properties of Exponential Semigroups	23
5.3	Fluctuation and Stability of the Ensemble Kalman-Bucy Filter	26
6	Strong Results in One-Dimensional Illustrative Examples	29
7	Regularisations and Comparisons	35
7.1	Covariance Regularisation: Inflation	35
7.2	Discretisation Matters	37
7.3	Particle Filter Comparisons	38
8	Some Topics for Discussion	42
8.1	Comments on the Results Presented	42
8.2	Bridging the Gap to Nonlinear Ensemble Filtering	44
8.3	Other Related Literature	46
8.4	Some Open Problems	47

1 Introduction

Consider a time-invariant, continuous-time, signal and observation model of the form,

$$\begin{aligned} d\mathcal{X}_t &= a(\mathcal{X}_t) dt + R^{1/2} d\mathcal{V}_t \\ d\mathcal{Y}_t &= h(\mathcal{X}_t) dt + R_1^{1/2} d\mathcal{W}_t \end{aligned} \quad (1.1)$$

where $\mathcal{X}_t$ is the underlying signal (latent) process, $\mathcal{Y}_t$ is the observation signal, $a(\cdot)$ and $h(\cdot)$ are the signal and sensor model functions, and $\mathcal{V}_t$ and $\mathcal{W}_t$ are continuous-time Brownian motion (noise) signals. The filtering problem [4, 9] is concerned with estimating some statistic(s) of the signal $\mathcal{X}_t$ conditioned on the observations $\mathcal{Y}_s$, $0 \leq s \leq t$. For example, one may want to characterise fully the distribution of $\mathcal{X}_t$ given $\mathcal{Y}_t$, or one may seek some moments of this distribution. The conditional distribution of $\mathcal{X}_t$ given $\mathcal{Y}_s$, $0 \leq s \leq t$ is called the (optimal, Bayesian) filtering distribution. When the model functions $a(\cdot)$, $h(\cdot)$ are linear, the exact (optimal, Bayesian) solution to this problem is completely characterised by the first two moments of the filtering distribution and these moments are given by the celebrated Kalman-Bucy filter [78, 4, 16].

Apart from the most special of nonlinear models, there is in general no finite dimensional optimal filter [12, 9]. In practice, some filter approximations are needed. For example, one may consider a type of “extended” Kalman filter [4] based on linearisation of the nonlinear model and application of the classical Kalman-Bucy filter. This method works well in suitably regular, and sufficiently close to linear problems. This method does not handle well multiple modes in the true filtering distribution. So-called Gaussian-sum filters are another Kalman-filter-type/based approximation designed to handle in some sense multiple modes in the filtering distribution [4]. More recently there has been some focus on Monte Carlo integration methods for approximating the optimal Bayesian filter [47, 9]. Such methods, termed particle filters or sequential Monte Carlo filters/methods [64, 58, 57], have the advantage of not being subject to the assumption of linearity or Gaussianity in the model. These particle filters are consistent in the number of Monte Carlo samples, i.e. with infinite computational power these methods converge to the optimal nonlinear filter. However, typical particle filtering algorithms exhibit high computational costs with approximation errors that grow (with a fixed sample size) with the signal/observation dimensions [47, 119]. These methods are not scalable to the high-dimensional filtering or state estimation problems found in the geosciences and other areas [61, 79, 98, 122].

The ensemble Kalman-Bucy filter (generally abbreviated **EnKF**) [60, 61] is a type of Monte Carlo sample approximation of a class of linear (in the observations) filter in the spirit of the Kalman filter. The **EnKF** is a recursive algorithm for propagating and updating the sample mean and sample covariance of an approximated Bayesian filter [61]. The filter works via the evolution of a collection (i.e. an ensemble) of samples (i.e. ensemble members, or particles) that each satisfies a type of Kalman-Bucy update equation, linear in the observations. In classical Kalman-Bucy filtering [78, 4, 16], a gain function, that depends on the filter error covariance, is used to weight a predicted state estimate with the signal observations, see [61, 16]. In the **EnKF**, the error covariance in the gain function is replaced by a type of sample covariance. The result is a system of interacting particles in the spirit of a mean-field approximation of a certain McKean-Vlasov-type diffusion equation [108, 54]. We may refine this discussion by giving the relevant equations for a most basic form of **EnKF**. Let $(\mathcal{V}_t^i, \mathcal{W}_t^i, \mathcal{X}_0^i)$ with $1 \leq i \leq N+1$ be $(N+1)$ independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. The most basic ensemble Kalman filter, originally due to Evensen [34, 60, 61], is defined by,

$$d\mathcal{X}_t^i = a(\mathcal{X}_t^i) dt + R^{1/2} d\mathcal{V}_t^i + \hat{P}_t^h R_1^{-1} \left[d\mathcal{Y}_t - \left(h(\mathcal{X}_t^i) dt + R_1^{1/2} d\mathcal{W}_t^i \right) \right] \quad (1.2)$$

with $1 \leq i \leq N+1$ and the (particle) sample mean and the sample cross-covariance defining the

so-called Kalman gain matrix given by,

$$\hat{X}_t := \frac{1}{N+1} \sum_{i=1}^{N+1} \mathcal{X}_t^i \quad \text{and} \quad \hat{P}_t^h := \frac{1}{N} \sum_{i=1}^{N+1} \left[\mathcal{X}_t^i - \hat{X}_t \right] \left[h(\mathcal{X}_t^i) - \frac{1}{N+1} \sum_{i=1}^{N+1} h(\mathcal{X}_t^i) \right]' \quad (1.3)$$

and we may also write the standard sample covariance,

$$\hat{P}_t := \frac{1}{N} \sum_{i=1}^{N+1} \left[\mathcal{X}_t^i - \hat{X}_t \right] \left[\mathcal{X}_t^i - \hat{X}_t \right]' \quad (1.4)$$

In this work we study this most basic ensemble Kalman filter as described above, and also more sophisticated variants, including the method of Sakov and Oke [126], that exhibit less fluctuation due to sampling noise. Readers familiar with the Kalman filter will recognise immediately some structural similarities as discussed above. However, there is no evolution equation given above for the covariance as in the Kalman filter (e.g. no Riccati-type matrix flow equation). Instead, we replace the relevant covariance matrices with their sample-based counterparts.

Importantly, if the underlying model is linear and Gaussian, then the filtering distribution is Gaussian, and the **EnKF** propagates exactly the sample mean and covariance of the optimal Bayesian filter, and is provably consistent. If the model is nonlinear and/or non-Gaussian, then a standard implementation of the **EnKF** propagates a sample-based estimate of the filtering mean and covariance (but not the true posterior sample mean or covariance and with no results on consistency). In the context of estimation theory, we may contrast the notion of a state estimator (or observer) with the notion of a Bayesian filter. The goal of the former is to design an observer that tracks in some suitable (typically point-wise) sense the underlying signal and perhaps provides some usable measure of uncertainty on this estimate. The goal of the latter is to compute or approximate the true (Bayesian) filtering distribution (or some related statistics). In the nonlinear setting, even with infinite computational power, the **EnKF** methods do not converge to the optimal nonlinear filter; and indeed their limiting objects are not well understood in this setting. As discussed more technically later, ensemble Kalman filters are probably best viewed in practice as a type of (random) sample-based state estimator for nonlinear signal/observation models. However, in the special case of linear signal and observation models they are indeed provably consistent approximations of the optimal Bayesian filter.

In practice, the ensemble Kalman filtering methodology is applied in high-dimensional, non-linear state-space models, e.g. see [60, 61] and the application references listed later in this introduction. Empirically, this method has shown good tracking performance in these applications, see [61] and the application references listed later. This tracking behaviour of the **EnKF** when applied to practical models may be explainable by viewing the **EnKF** as a dynamic state estimator. The fluctuation, stability, and contraction properties of the **EnKF** studied in this article (albeit mainly for linear-Gaussian models) may be viewed in this context also, and provide some insight into the state estimate tracking behaviour seen in practice.

1.1 Purpose

The purpose of this review is to present a comprehensive overview of the theory of ensemble Kalman-Bucy filtering with an emphasis on rigorous results and behavioural characterisations for linear-Gaussian signal/observation models. We present a system of equations that describe the flow of individual particles, the flow of the sample covariance, and the flow of the sample mean in continuous-time ensemble filtering. We consider these equations and their characteristics in a number of popular **EnKF** varieties. Given these equations, *we study in detail some fluctuation, stability, and error contraction results for the various ensemble Kalman filtering equations.* We discuss the relevance and importance of these results in terms of characterising

the **EnKF** behaviour, and we contrast these results with those considered in classical studies of stability in Kalman-Bucy filtering.

We remark that classical studies of stability in (traditional, non-ensemble-type) Kalman-Bucy filtering are important because they rigorously establish the type of “tracking” properties desired in a filtering or estimation problem; and they establish intuitive, testable, model-based conditions (e.g. so-called model observability) for achieving these convergence properties. Classical results in Kalman-Bucy filtering also establish the (exponential) convergence of the error covariance to a fixed steady-state value computable from the model parameters. See the review [16, 18] for detailed results in the classical context and historical remarks. The results in this work seek to characterise in an analogous manner the practical performance and behaviour of ensemble Kalman filtering, and these results then provide guidance and intuition on the tracking, approximation error, and other properties of these practical methods. Notably, the stochastic fluctuation properties of ensemble Kalman methods also need to be established; and counterparts of this latter analysis do not arise at all in classical Kalman filtering analyses.

1.2 Overview of the Main Topics and Literature

In this subsection we touch on the main topics and related literature as it pertains to the **EnKF**. These topics include the fluctuation, stability, and contractive properties of the relevant **EnKF** stochastic equations. Later, toward the end of this article, we discuss some of these topics in the context of filtering and state estimation more broadly, and we touch on other related but somehow distinct results as they pertain to the **EnKF** more specifically.

The **EnKF** is a key numerical method for solving high-dimensional forecasting and data assimilation problems; see, e.g., [60, 61]. In particular, applications have been motivated by inference problems in ocean and atmosphere sciences [103, 105, 79, 113], weather forecasting [5, 6, 34, 71], environmental and ecological statistics [1, 76], as well as in oil reservoir simulations [62, 111, 130], and many others. This list is by no means exhaustive, nor the cited articles fully representative of the respective applications. We refer to (some of) the seminal methodology papers in [59, 71, 34, 7, 72, 66, 26, 5, 143, 6, 137, 126, 121, 145]. This long list is not exhaustive; see also the books [61, 79, 98, 122] for more background, and the detailed chronological list of references in Evensen’s text [61].

In continuous-time, we may broadly break down the class of **EnKF** methods into three distinct types; distinguished by the level of fluctuation added via sampling noise needed to ensure that the **EnKF** sample mean and covariance are consistent in the linear-Gaussian setting. The original form of the **EnKF** is the so-called ‘vanilla’ **EnKF** of Evensen [34, 61], see also [95]; and this method exhibits the most fluctuations due to sampling of both signal and observation noises. The next class is the so-called ‘deterministic’ **EnKF** of Sakov and Oke [126], see also [13, 121]; which exhibits (considerably) less fluctuation. In the continuous-time linear-Gaussian setting this class is representative of the so-called square-root **EnKF** methods [137, 96] (which differ somewhat in discrete-time, e.g. contrast [126] with [137], see also [96]). Finally, there has been recent interest in so-called transport-inspired **EnKF** methods [121, 133]; which apart from initialisation noise/randomisation are completely deterministic and whose analysis in the linear model setting follows closely that of the classical Kalman-Bucy filter, cf. [16]. These classes do not distinguish the totality of **EnKF** methodology (especially in nonlinear or non-Gaussian models); which may further consist of so-called covariance regularisation methods [7, 72, 66, 109, 60], etc. However, in the linear-Gaussian case, these three classes broadly capture the fundamentals.

As discussed later, the fully deterministic, transport-inspired **EnKF** method, see [121, 133], is a rather special case in the linear-Gaussian setting and is not studied in detail in this article where linear-Gaussian models are the focus. Nevertheless, we point to [43, 44] for certain mean-field consistency results, non-asymptotic fluctuation (e.g. finite sample size) results, and the long-time behaviour of this particular method in the case of a nonlinear signal model and linear observations. We also touch on this method briefly throughout; but when we refer to the general

EnKF we typically mean the so-called ‘vanilla’ [34, 95] or ‘deterministic’ [126, 13] methods (which will become clear as the article progresses).

Convergence to a mean-field limit, and large-sample asymptotics, of the discrete-time **EnKF** was studied in [100, 107, 91, 94], in the sense of taking the number of particles to infinity. The discrete-time square root form of the **EnKF** is accommodated in [91, 94], and nonlinear state-space models are accommodated in [100]. In the continuous-time, linear-Gaussian, setting, the convergence (in sample size) of the three broad classes of **EnKF** to the true Kalman-Bucy filter is more immediate; and follows from the sample mean and sample covariance evolution equations in [54, 19]. In this latter sense, we recover the fact that the **EnKF** is a consistent approximation of the optimal, Bayesian filter (i.e. the classical Kalman-Bucy filter) in the linear-Gaussian setting as discussed earlier. The mean-field limit of various **EnKF** methods in the continuous-time, nonlinear model setting is studied in [99, 43, 94].

We remark in the nonlinear model setting (discrete or continuous-time), see [100, 123, 99, 43, 94], the mean-field limiting equations (and distribution) are not easily related to the optimal filter. Moreover, in practice, one is typically interested in the non-asymptotic (in terms of ensemble size) fluctuation properties as well as the long time/stability behaviour of the particle-type filtering approximations.

The fluctuation analysis of the **EnKF** is studied in detail in the linear-Gaussian setting in [22, 21, 19]. In [22] a complete Taylor-type stochastic expansion of the sample covariance is given at any order with bounded remainder terms and estimates. Both non-asymptotic and asymptotic bias and variance estimates for the **EnKF** sample covariance and sample mean are given explicitly in [22]. These latter expansions directly imply an almost sure strong form of a central limit-type result on the sample covariance and sample mean at any time. The analysis in [22] is considered over the entire path space of the matrix-valued Riccati stochastic differential equation that describes the flow of the sample covariance. However, most of the non-asymptotic time-uniform results in [22] hold only when the underlying signal is stable. In [21, 19] we consider the case in which the underlying signal may be unstable, and we provide time-uniform, non-asymptotic moment estimates and time-uniform control over the fluctuation of the sample covariance and mean about their limiting Riccati and Kalman-Bucy filtering terms.

The emphasis of *time-uniformity* on the moment bounds and on the fluctuation bounds on the sample mean and sample covariance (about the true optimal Bayesian filtering mean and covariance) is important. If these bounds are allowed to grow in time, e.g. typically in this analysis one can easily obtain bounds that grow exponentially in time, then these bounds quickly become useless for any practical numerical application; e.g. an exponent > 200 may induce an exceedingly pessimistic bound greater than the estimated number of particles of matter in the visible universe. We remark also that our emphasis on *accommodating unstable (latent) signal models* is important because time-uniform fluctuation results in such cases (which are of real practical importance) are significantly more difficult to obtain under testable and realistic model assumptions (like the classical observability and controllability model assumptions in the control and filtering literature [4, 16]).

In [54], stability of the **EnKF** in continuous-time linear-Gaussian models is considered under the assumption that the underlying signal model is also stable. This latter assumption is in contrast with classical Kalman-Bucy filter stability results, which hold in the linear-Gaussian setting under the much weaker (and more natural) condition of signal detectability [140, 16, 18]. The classical Kalman-Bucy filter is stable as a result of the closed-loop stabilising properties of the so-called Kalman gain matrix, which is closely connected to the flow of the filter error covariance described by a Riccati differential equation. The **EnKF** analogue, in linear-Gaussian settings, is the sample covariance, and its random fluctuation properties (noted in the preceding paragraph) are the main source of difficulty in establishing the closed-loop filter stability in those models in which the underlying signal itself is unstable.

In [138], the authors analyse the long-time behaviour of the (discrete-time) **EnKF** in a class

of nonlinear systems, with finite ensemble size, using Foster-Lyapunov techniques. Applying the results of [138] to the basic linear-Gaussian filtering problem, the analysis and assumptions in [138] then also require stability of the underlying signal model. In a traditional sense, the conditions needed in [138] are hard to check, e.g. as compared to the classical observability or controllability-type model conditions in Kalman filtering analysis; but a range of examples are given in [138]. In [82], the long-time behaviour of the **EnKF** is analysed in both discrete and continuous time settings with similar conditions on the model as in [138]; and which again if linearised equates to a form of stability on the signal model.

We emphasise again that the type of analysis in [82, 138, 54] cannot handle unstable, or transient, signal models; i.e. signals with sample paths with at least one coordinate that may grow unbounded. In the context studied in [82, 138, 54] dealing with stable or bounded latent signal processes (e.g. the Lorenz-class of signal models [82, 138]), the important question on the filter stability or filtering error estimates relies on obtaining meaningful quantitative fluctuation constants decreasing with the number of ensemble members to achieve a desired performance. Of course, time uniformity of these bounds follows trivially in this setting from the boundedness properties of the latent signal process.

Covariance inflation is a mechanism used in practical methods to increase the positive-definiteness of the sample covariance matrix and essentially amplify its effect on the stabilisation properties of the Kalman gain matrix. In [82] time-uniform **EnKF** error boundedness results follow under a true signal stability condition and given a sufficiently large variance inflation regime. See also [139, 106] for related stability analysis in the presence of adaptive covariance inflation and projection techniques. In [19] in the continuous-time linear-Gaussian setting, the mechanism by which covariance inflation acts to stabilise the ensemble filter is exemplified, see also [24]. Covariance localisation is studied rigorously in [44] in the case of the fully deterministic, transport inspired ensemble filter [121, 133].

In the continuous-time, and linear-Gaussian setting, the first work to relax the assumption of underlying signal stability for the **EnKF** is in [22, 21, 19]. In those articles, latent signals with sample paths that may grow unbounded (to infinity exponentially fast) are accommodated. That work is based on both a fluctuation analysis of the sample covariance and the sample mean [22, 21, 19], followed by studies on the long-time behaviour, e.g. stability properties, of both the sample covariance and mean [21, 19]. Time-uniform fluctuation properties are given under a type of (strong) signal observability condition. In this setting time-uniformity of these results is non-trivial. This assumption is in keeping with classical Kalman-Bucy filtering and Riccati equation results; and does not require any form of underlying signal stability. As the authors of [138] note in their stability analysis, they use “few properties of the forecast [predicted] covariance matrix other than positivity”. As noted in [138], this lends generality to their results, but conversely places the burden back on the signal model assumptions (including those assumptions of true signal stability). Contrast this with the work in [22, 21, 19] where emphasis is placed on the fluctuation analysis of the sample covariance, with a primary aim of removing the stability assumptions needed on the underlying signal model. The time-uniform fluctuation and stochastic perturbation contributions in [22, 21, 19], were discussed earlier. Given this fluctuation analysis, the stability of the filter sample mean and sample covariance and their (time) asymptotic properties are studied in [22, 21, 19] without stability assumptions on the underlying signal model. These results rigorously establish the type of “tracking” properties desired by a filtering or estimation solution.

Although of lesser practical use in applications, strong results in the one-dimensional setting are also derived in [21] that converge, e.g. in the limit with the ensemble size, to those properties of the classical Kalman-Bucy filter. For example, we can recover the optimal exponential contraction and filter stability rates, etc. In the multidimensional setting, the decay rates to equilibrium are not sharp, and the stationary measures are not given in closed form.

1.3 Aims and Contributions

The main goal of this article is to: 1). present a novel formulation for ensemble filtering in linear-Gaussian, continuous-time, systems that lends itself naturally to analysis; 2). provide detailed fluctuation analysis of the ensemble Kalman-Bucy flow, the sample mean, and the stochastic Riccati equation describing the sample covariance; 3). study the stability of the resulting stochastic Riccati differential equation that describes the flow of the sample covariance; 4). study the stability of the continuous-time ensemble Kalman-Bucy update equation that is coupled to this stochastic Riccati equation, and which describes the flow of the sample mean (or the sample mean minus the true signal, i.e. the sample error signal). This article is primarily a review of the literature and results in these directions. The prime focal point of this review are the articles [54, 22, 21, 17, 19], which focus heavily on the linear-Gaussian model setting. In this review, *an emphasis is placed on deriving time-uniform fluctuation, stability, and contraction results under testable model conditions reminiscent of those classical observability and controllability-type model assumptions. Importantly, we do not generally assume the true underlying signal is stable in this review.*

Throughout this review we contrast and discuss the presented results with the broader literature on the rigorous mathematical behaviour of ensemble Kalman-type filtering. For example, we find easily that the sample covariance matrix in the broad class of **EnKF** methods considered is always under-biased when compared to the true covariance matrix. This may motivate, from a pure uncertainty quantification viewpoint, some form of covariance regularisation [7, 72, 66, 109, 60]. We provide detailed analysis illustrating the effect of inflation regularisation on stability (similarly to [82, 139, 106]). As another example, we provide strong intuition for so-called catastrophic filter divergence (studied previously in [67, 65, 83]) based on rigorous (heavy-tailed) fluctuation properties inherent to the relevant sample covariance matrices and their invariant distributions. We contrast the so-called ‘vanilla’ **EnKF** of [34, 61] with the ‘deterministic’ **EnKF** of Sakov and Oke [126] in terms of their fluctuation and sample noise characteristics and we show how this affects their respective sample behaviour and stability properties.

As with classical (non-ensemble) Kalman filtering, the importance of the results reviewed is in rigorously establishing the type of tracking and stability behaviour desired in filtering applications [4, 47, 9, 16]. For example, our results imply conditions under which the initial estimation errors are forgotten, and that the flow of the sample mean converges to the true Kalman filtering (conditional mean) state estimate (and thus the signal) in the average. In the case of the **EnKF**, there must be some emphasis placed on the stochastic behaviour of the ensemble (Monte Carlo) mean and covariance in order to establish filter stability. We also provide the analogue of the error covariance fixed point in classical Kalman filtering [4, 16]; whereby we state results that ensure the sample covariance matrix converges to an invariant, steady-state, distribution. We characterise the properties of this invariant distribution and relate this to the sample behaviour of the ‘vanilla’ **EnKF** [34, 61] and the ‘deterministic’ **EnKF** [126].

We focus on the linear, continuous-in-time, Gaussian setting in this review and note that in this case the sample mean and sample covariance are consistent approximations of the optimal Bayesian filtering mean and covariance. We emphasise that even in the linear-Gaussian case, the samples themselves are not in general independent. The analysis even in the linear setting is highly technical [54, 22, 21, 17, 19] and the results presented in this case are aimed as a step in the progression to more applied results and intuition in nonlinear model settings. There is some precedent for studying the relative properties, behaviour, or performance of ensemble Kalman filtering firstly with linear-Gaussian signal models [60]. For example, the seminal article [34] illustrated that a perturbation of the observations in the ensemble Kalman filter was necessary to recover a consistent covariance limit (to the true Kalman filter for linear-Gaussian systems); or to achieve the standard Monte Carlo error rate with a finite set of particles. The analysis (and even derivation) of ensemble square root filters for linear-Gaussian system models is standard

[127, 104], etc. Convergence of the ensemble Kalman filter in inverse problems is studied in [128] in the linear setting. We discuss connections and extensions of the results in this article to the nonlinear model setting toward the end.

We also briefly contrast the approximation capabilities of particle filtering (sequential Monte Carlo) methods [64, 58] with the EnKF. We give a revealing, and perhaps surprising, simple result illustrating the complete failure of the bootstrap particle filter [64] to track unstable linear-Gaussian latent signals. Compared to the EnKF, the fluctuation and stability of various particle filtering methods (e.g. see [49, 47, 39, 114, 142, 48, 56]) is a rather mature topic. Nevertheless, time-uniform particle filtering estimates rely on mixing-type, or certain contractive, conditions on the mutation transition which do not hold in general in the case of unstable linear-Gaussian models. We contrast this new (rather negative) particle filtering result with its (positive) counterpart for the EnKF.

Note that the analysis and proofs in [54, 22, 21, 17, 19], while motivated originally by ensemble Kalman-type filtering methods, are largely presented as independent technical results on certain general classes of matrix-valued Riccati diffusion equations and associated linear stochastic differential equations with random coefficients. In this review we emphasise the work in [54, 22, 21, 17, 19] via a series of results directly and solely stated in the context of ensemble Kalman-type filtering. Throughout we relate our results to the broader technical literature on ensemble Kalman filtering and we emphasise the practical significance of these results, e.g. via the tracking property of the filter, its stability, or via their error fluctuation or catastrophic divergence behaviour, among other topics. We also contrast the behaviour of the various classes of continuous-time EnKF methods.

1.4 Notation

We remark firstly that some care must be taken throughout to keep track of the font stylings; e.g. upright vs. calligraphic vs. script, etc. There is typically a relationship between like symbols appearing with different stylings.

Hatted terms $\hat{\cdot}$ should be viewed as being indexed to the ensemble size $N \geq 1$, i.e. $\hat{\cdot} := \cdot^N$. Time is indexed variously by $s, t, u, \tau \in [0, \infty[$. We write $c, c_n, c_\tau, c_{n,\tau}, c_{n,\tau}(Q), c_{n,\tau}(z, Q) \dots$ for some positive constants whose values may vary from result to result, and which only depend on the indexed/referenced parameters n, τ, z, Q , etc, as well as implicitly on the model parameters (A, H, R, R_1) introduced later. Importantly, these constants do not depend on the time horizon t , nor on the number of ensemble particles N .

Let $\mathbb{M}_d$ be the set of $(d \times d)$ real matrices with $d \geq 1$ and $\mathbb{M}_{d_1, d_2}$ the set of $(d_1 \times d_2)$ real matrices. Let $\mathbb{S}_d \subset \mathbb{M}_d$ be the subset of symmetric matrices, and $\mathbb{S}_d^0$, and $\mathbb{S}_d^+$ the subsets of positive semi-definite and definite matrices respectively. We write $A \geq B$ when $A - B \in \mathbb{S}_d^0$; and $A > B$ when $A - B \in \mathbb{S}_d^+$. We denote by 0 and I the null and identity matrices, for any $d \geq 1$. Given $R \in \partial\mathbb{S}_d^+ := \mathbb{S}_d^0 - \mathbb{S}_d^+$ we denote by $R^{1/2}$ a (non-unique) symmetric square root of R . When $R \in \mathbb{S}_d^+$ we choose the unique symmetric square root. We write A' the transpose of A , and $A_{\text{sym}} = (A + A')/2$ its symmetric part. We denote by $\text{Abse}(A) := \max \{\text{Re}(\lambda) : \lambda \in \text{Spec}(A)\}$ its spectral abscissa. We also denote by $\text{Tr}(A)$ the trace. When $A \in \mathbb{S}_d$ we let $\lambda_1(A) \geq \dots \geq \lambda_d(A)$ denote the ordered eigenvalues of A . We equip $\mathbb{M}_d$ with the spectral norm $\|A\| = \|A\|_2 = \sqrt{\lambda_1(AA')}$ or the Frobenius norm $\|A\| = \|A\|_{\text{Frob}} = \sqrt{\text{Tr}(AA')}$.

Let $\mu(A)$ denote a matrix logarithmic “norm” (which can be < 0), see [132]. The logarithmic norm is a tool to study the growth of solutions to ordinary differential equations and the error growth in approximation methods. For any square matrix $A \in \mathbb{M}_d$, the logarithmic norm is the smallest element in the set $\{h \in \mathbb{R} : \|\exp(At)\| \leq \exp(ht), t \geq 0\}$ where $\|\cdot\|$ is any matrix norm and the value $\mu(A)$ may be considered to be indexed to the matrix norm employed. For example, the (2-)logarithmic “norm”, or spectral log-norm, is given by $\mu(A) = \lambda_1(A_{\text{sym}})$. We have $\mu(\cdot) \geq \text{Abse}(\cdot)$ in general, but importantly we note that if $\text{Abse}(\cdot) < 0$, then there is a matrix norm $\|\cdot\|$ defining a logarithmic norm such that $\mu(\cdot) < 0$, see [132, Theorem 5].

2 Kalman-Bucy Filtering

Consider a time-invariant linear-Gaussian filtering model of the following form,

$$\begin{aligned} d\mathcal{X}_t &= A \mathcal{X}_t dt + R^{1/2} d\mathcal{V}_t \\ d\mathcal{Y}_t &= H \mathcal{X}_t dt + R_1^{1/2} d\mathcal{W}_t \end{aligned} \quad (2.1)$$

where $A \in \mathbb{M}_d$ and $H \in \mathbb{M}_{d_y, d}$ are the signal and sensor model matrices respectively, and $R \in \mathbb{S}_d^0$ and $R_1 \in \mathbb{S}_{d_y}^+$ are the respective signal and sensor noise covariance matrices. The noise inputs $\mathcal{V}_t$ and $\mathcal{W}_t$ are d and d_y -dimensional Brownian motions, and $\mathcal{X}_0$ is an d -dimensional Gaussian random variable (independent of $(\mathcal{V}_t, \mathcal{W}_t)$) with mean $\mathbb{E}(\mathcal{X}_0)$ and covariance $P_0 \in \mathbb{S}_d^0$.

We let $\mathcal{Y}_0 = 0$ and $\mathcal{Y}_t = \sigma(\mathcal{Y}_s, s \leq t)$ be the σ -algebra generated by the observations. The conditional distribution $\eta_t := \text{Law}(\mathcal{X}_t | \mathcal{Y}_t)$ of the signal states $\mathcal{X}_t$ given $\mathcal{Y}_t$ is Gaussian with a conditional mean and covariance given by

$$X_t := \mathbb{E}(\mathcal{X}_t | \mathcal{Y}_t) \quad \text{and} \quad P_t := \mathbb{E}([\mathcal{X}_t - X_t][\mathcal{X}_t - X_t]' | \mathcal{Y}_t).$$

The mean and the covariance obey the Kalman-Bucy and the Riccati equations

$$dX_t = A X_t dt + P_t H' R_1^{-1} (d\mathcal{Y}_t - H X_t dt) \quad (2.2)$$

$$\partial_t P_t = \text{Ricc}(P_t) \quad (2.3)$$

with the Riccati drift function from $\mathbb{S}_d^0$ into $\mathbb{S}_r$ defined for any $Q \in \mathbb{S}_d^0$ by

$$\text{Ricc}(Q) := AQ + QA' - QSQ + R \quad (2.4)$$

and with,

$$S := H' R_1^{-1} H \quad (2.5)$$

Importantly, the covariance of the conditional distribution $\text{Law}(\mathcal{X}_t | \mathcal{Y}_t)$ in this case does not depend on the observations $\mathcal{Y}_t$. The error $Z_t := (X_t - \mathcal{X}_t)$ satisfies

$$\begin{aligned} dZ_t &= (A - P_t S) Z_t dt + P_t H' R_1^{-1/2} d\mathcal{W}_t - R^{1/2} d\mathcal{V}_t \\ &\stackrel{\text{law}}{=} (A - P_t S) Z_t dt + (P_t S P_t + R)^{1/2} d\mathcal{B}_t \end{aligned} \quad (2.6)$$

where $\mathcal{B}_t$ is some independent d -dimensional Brownian motion. Here we make use of a martingale representation theorem, e.g. [80, Theorem 4.2], see also [55].

Let $\phi_t(Q) := P_t$ denote the flow of the matrix differential equation (2.3) with $P_0 = Q \in \mathbb{S}_d^0$. Let $\psi_t(z, Q) := Z_t$ denote the flow of the stochastic error (2.6) with $Z_0 = z = (x - \mathcal{X}_0) \in \mathbb{R}^d$ and $P_t = \phi_t(Q)$. Finally, we denote the flow of the Kalman-Bucy update (2.2) with $X_0 = x \in \mathbb{R}^d$ by $\chi_t(x, Q) := X_t$. This notation allows us to reference the flows $\psi_t(z, Q)$, $\phi_t(Q)$, $\chi_t(x, Q)$ with respect to their initialisation at $t = 0$ which is useful when we compare flows and study stability.

Throughout this section, we assume that $(A, R^{1/2})$ and (A, H) are controllable and observable pairs in the sense that

$$\left[R^{1/2}, AR^{1/2}, \dots, A^{r-1} R^{1/2} \right] \quad \text{and} \quad \begin{bmatrix} H \\ HA \\ \vdots \\ HA^{r-1} \end{bmatrix} \quad \text{have rank } d. \quad (2.7)$$

Note that if $R \in \mathbb{S}_d^+$ is positive definite, which is quite common in filtering problems, it follows that controllability holds trivially. We consider the observability and controllability Gramians

$(\mathcal{O}_t, \mathcal{C}_t(\mathcal{O}))$ and $(\mathcal{C}_t, \mathcal{O}_t(\mathcal{C}))$ associated with the triplet (A, R, S) and defined by

$$\mathcal{O}_t := \int_0^t e^{-A's} S e^{-As} ds \quad \text{and} \quad \mathcal{C}_t(\mathcal{O}) := \mathcal{O}_t^{-1} \left[\int_0^t e^{-(t-s)A'} \mathcal{O}_s R \mathcal{O}_s e^{-(t-s)A} ds \right] \mathcal{O}_t^{-1} \quad (2.8)$$

$$\mathcal{C}_t := \int_0^t e^{As} R e^{A's} ds \quad \text{and} \quad \mathcal{O}_t(\mathcal{C}) := \mathcal{C}_t^{-1} \left[\int_0^t e^{(t-s)A} \mathcal{C}_s S \mathcal{C}_s e^{(t-s)A} ds \right] \mathcal{C}_t^{-1}.$$

Given (2.7), for any finite $\tau > 0$, there exists some finite parameters $\varpi_{\pm}^{o,c}, \varpi_{\pm}^c(\mathcal{O}), \varpi_{\pm}^o(\mathcal{C}) > 0$ such that

$$\varpi_-^c \leq \|\mathcal{C}_\tau\| \leq \varpi_+^c \quad \text{and} \quad \varpi_-^o \leq \|\mathcal{O}_\tau\| \leq \varpi_+^o \quad (2.9)$$

$$\varpi_-^c(\mathcal{O}) \leq \|\mathcal{C}_\tau(\mathcal{O})\| \leq \varpi_+^c(\mathcal{O}) \quad \text{and} \quad \varpi_-^o(\mathcal{C}) \leq \|\mathcal{O}_\tau(\mathcal{C})\| \leq \varpi_+^o(\mathcal{C}). \quad (2.10)$$

The parameter τ is often called the interval of observability-controllability, see [30].

These rank conditions (2.7) ensure the existence and the uniqueness of a positive definite fixed-point matrix P_∞ solving the algebraic Riccati equation

$$\text{Ricc}(P_\infty) := AP_\infty + P_\infty A' - P_\infty S P_\infty + R = 0. \quad (2.11)$$

Indeed, if (2.7) holds, then $P_\infty \in \mathbb{S}_d^+$ and $\text{Absc}(A - P_\infty S) < 0$. We may relax the controllability assumption to just stabilisability, in which case $P_\infty \in \mathbb{S}_d^0$ and $\text{Absc}(A - P_\infty S) < 0$; see [88, 110, 92] and the convergence results in [90, 35]. Under just a detectability condition, it follows that $P_\infty \in \mathbb{S}_d^0$ and $\text{Absc}(A - P_\infty S) \leq 0$, i.e. $(A - P_\infty S)$, is only marginally stable, and convergence to this solution is given under mild additional conditions in [118, 36, 116]. In [140], given only detectability, the time-varying “closed loop” matrix $(A - \phi_t(Q)S)$ is shown to be stabilising, even when $(A - P_\infty S)$ is only marginally stable.

For any $s \leq t$ and $Q \in \mathbb{S}_d^0$ we define the state-transition matrix,

$$\mathcal{E}_{s,t}(Q) := \exp \left(\int_s^t (A - \phi_u(Q)S) du \right) \iff \partial_t \mathcal{E}_{s,t}(Q) = (A - \phi_u(Q)S) \mathcal{E}_{s,t}(Q). \quad (2.12)$$

When $s = 0$ we often write $\mathcal{E}_t(Q)$ instead of $\mathcal{E}_{0,t}(Q)$. The matrix $\mathcal{E}_t(Q)$ is the fundamental matrix. We have $\mathcal{E}_{s,t}(Q) = \mathcal{E}_t(Q) \mathcal{E}_s(Q)^{-1}$. The following convergence estimates follow from [16, 18]: For any $Q, Q_1, Q_2 \in \mathbb{S}_d^0$ and any $t \geq 0$ we have the local contraction inequalities

$$\|\mathcal{E}_t(Q)\| \leq c(1 + \|Q\|) \|\mathcal{E}_t(P_\infty)\| \quad \text{and} \quad \|\mathcal{E}_t(P_\infty)\| = \|e^{t(A-P_\infty S)}\| \leq c e^{-\alpha t} \quad (2.13)$$

for some finite $\alpha, c > 0$ and with P_∞ solving (2.11) and

$$\|\mathcal{E}_t(Q_2) - \mathcal{E}_t(Q_1)\| \leq c(Q_1, Q_2) e^{-2\alpha t} \|Q_2 - Q_1\| \quad (2.14)$$

for some finite constant $c(Q_1, Q_2) > 0$. In addition, there exists some parameter $\tau > 0$ such that for any $s \geq 0$ and any $t \geq \tau > 0$ we have the uniform estimates,

$$\|\mathcal{E}_{s,s+t}(Q)\| \leq c_\tau \|\mathcal{E}_t(P_\infty)\| \quad (2.15)$$

Note it is desirable to relate the decay of $\mathcal{E}_{s,s+t}(Q)$ to the decay at the fixed point $\|\mathcal{E}_t(P_\infty)\| = \|e^{t(A-P_\infty S)}\| \leq c e^{-\alpha t}$ (since as $t \rightarrow \infty$ it is clear that we cannot do better). See [18] for an explicit Floquet-type expression of $\mathcal{E}_t(Q)$ in terms of $\mathcal{E}_t(P_\infty)$.

The convergence and stability properties of the Kalman-Bucy filter and the associated Riccati equation are directly related to the contraction properties of the state-transition matrix $\mathcal{E}_{s,t}(Q)$. To get some intuition for this we note,

$$\psi_t(z, Q) = \mathcal{E}_{s,t}(Q) \psi_s(z, Q) + \int_s^t \mathcal{E}_{u,t}(Q) (\phi_u(Q) S \phi_u(Q) + R)^{1/2} d\mathcal{B}_u \quad (2.16)$$

and

$$\phi_t(Q) = \mathcal{E}_{s,t}(Q) \phi_s(Q) \mathcal{E}_{s,t}(Q)' + \int_0^t \mathcal{E}_{s,t}(Q) (\phi_s(Q) S \phi_s(Q) + R) \mathcal{E}_{s,t}(Q)' ds \quad (2.17)$$

for any $s \leq t$.

From [16], for any $t \geq \tau > 0$ and any $Q \in \mathbb{S}_d^0$ we have the uniform estimates

$$(\mathcal{O}_\tau(\mathcal{C}) + \mathcal{C}_\tau^{-1})^{-1} \leq \phi_t(Q) \leq \mathcal{O}_\tau^{-1} + \mathcal{C}_\tau(\mathcal{O}). \quad (2.18)$$

We also have

$$0 \leq \phi_t(Q) \leq P_\infty + e^{(A-P_\infty S)t}(Q - P_\infty)e^{(A-P_\infty S)'t} \quad (2.19)$$

The following stability result follows from [16, 18]: For any $Q_1, Q_2 \in \mathbb{S}_d^0$ and for any $t \geq 0$,

$$\|\phi_t(Q_1) - \phi_t(Q_2)\| \leq c(1 + \|Q_1\|^2 + \|Q_2\|^2) \|\mathcal{E}_t(P_\infty)\|^2 \|Q_2 - Q_1\| \quad (2.20)$$

and recall the exponential contraction estimate on $\|\mathcal{E}_t(P_\infty)\|$ in (2.13). Similarly, using (2.15), for any $s \geq 0$ and any $t \geq \tau > 0$, we have

$$\|\phi_{s,s+t}(Q_1) - \phi_{s,s+t}(Q_2)\| \leq c_\tau \|\mathcal{E}_t(P_\infty)\|^2 \|Q_2 - Q_1\| \quad (2.21)$$

Note that both (2.20) and (2.21) imply immediately that $\phi_t(Q) \rightarrow_{t \rightarrow \infty} P_\infty$ exponentially fast for any $Q \in \mathbb{S}_d^0$; e.g. by letting $Q_2 = P_\infty$.

Note that the uniform estimates with constants independent of the initial condition stated throughout, involve some arbitrarily small, positive time parameter τ , which can be directly related to the notion of a so-called observability/controllability interval introduced earlier; for further details on this topic we refer to [30, 16]. Contrast, for example, the stability results (2.20) and (2.21). The symbol τ is reserved for this arbitrary small time parameter throughout the article.

Results (e.g. bounds and convergence results) on the flow of the inverse of the solution of the Riccati equation are considered in [16] and are relevant for proving results on the flow of the Riccati equation itself; e.g. upper bounds on the flow of the inverse solution help to lower bound solutions of the Riccati flow. The flow of the inverse Riccati solution may also be of interest on its own as it relates to the flow of “information” (as the inverse of covariance).

Given the contraction properties on $\mathcal{E}_{s,t}(Q)$ it is often said the “deterministic part” of the filter error $\partial_t Z_t = (A - P_t S) Z_t$ is stable. From [16] we can be more explicit if desired, for example, for any $t \geq \tau$ we have the uniform estimate,

$$\sup_{Q \in \mathbb{S}_d^0} \|\mathbb{E}(\psi_t(z, Q) \mid \mathcal{X}_0)\| \leq c e^{-\alpha t} \|x - \mathcal{X}_0\| \quad (2.22)$$

for some rate $\alpha > 0$ and some finite constant $c > 0$. Moreover, the conditional probability of the following event

$$\|\psi_t(z, Q)\| \leq c(Q) \left(e^{-\alpha t} \|x - \mathcal{X}_0\| + \frac{e^2}{\sqrt{2}} \left[\frac{1}{2} + (\delta + \sqrt{\delta}) \right] \right) \quad (2.23)$$

given the state variable $\mathcal{X}_0$ is greater than $1 - e^{-\delta}$, for any $\delta \geq 0$. And, for any $t \geq 0$, $z_1, z_2 \in \mathbb{R}^d$, $Q_1, Q_2 \in \mathbb{S}_d^0$ and any $n \geq 1$ we have the almost sure local contraction estimate

$$\begin{aligned} & \mathbb{E}(\|\psi_t(z_1, Q_1) - \psi_t(z_2, Q_2)\|^n \mid \mathcal{X}_0)^{\frac{1}{n}} \\ & \leq c(Q_1, Q_2) e^{-\alpha t} \|z_1 - z_2\| + c_n(Q_1, Q_2) e^{-\alpha t} (1 + \|x_1 - \mathcal{X}_0\|) \|Q_1 - Q_2\| \end{aligned} \quad (2.24)$$

with some rate $\alpha > 0$ and the finite constants $c(Q_1, Q_2), c_n(Q_1, Q_2) > 0$.

3 Kalman-Bucy Diffusion Processes

For any probability measure η on $\mathbb{R}^d$ we let $\mathcal{P}_\eta$ denote the η -covariance

$$\eta \mapsto \mathcal{P}_\eta := \eta \left([\iota - \eta(\iota)][\iota - \eta(\iota)]' \right) \quad (3.1)$$

with the identity function $\iota(x) := x$ and the column vector $\eta(f) := \int f d\eta$ for some measurable function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

We now consider three different cases of a conditional nonlinear McKean-Vlasov-type diffusion process,

$$\begin{aligned} \text{(F1)} \quad d\mathcal{X}_t &= A \mathcal{X}_t dt + R^{1/2} d\mathcal{V}_t + \mathcal{P}_{\bar{\eta}_t} H' R_1^{-1} \left[d\mathcal{Y}_t - \left(H \mathcal{X}_t dt + R_1^{1/2} d\mathcal{W}_t \right) \right] \\ \text{(F2)} \quad d\mathcal{X}_t &= A \mathcal{X}_t dt + R^{1/2} d\mathcal{V}_t + \mathcal{P}_{\bar{\eta}_t} H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t + \bar{\eta}_t(\iota)}{2} \right) dt \right] \\ \text{(F3)} \quad d\mathcal{X}_t &= A \mathcal{X}_t dt + R \mathcal{P}_{\bar{\eta}_t}^{-1} (\mathcal{X}_t - \bar{\eta}_t(\iota)) dt + \mathcal{P}_{\bar{\eta}_t} H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t + \bar{\eta}_t(\iota)}{2} \right) dt \right] \end{aligned} \quad (3.2)$$

where

$$\bar{\eta}_t := \text{Law}(\mathcal{X}_t \mid \mathcal{Y}_t) \quad (3.3)$$

and thus the diffusions in (3.2) depend in some nonlinear fashion on the conditional law of the diffusion process itself. In all three cases $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$ are independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. These diffusions are time-varying Ornstein-Uhlenbeck processes [54] and consequently $\bar{\eta}_t$ is Gaussian; see also [16]. These Gaussian distributions have the same conditional mean $\bar{\eta}_t(\iota)$ and conditional covariance $\mathcal{P}_{\bar{\eta}_t}$.

Proposition 3.1 ([54, 16]). *We have*

$$\bar{\eta}_t := \text{Law}(\mathcal{X}_t \mid \mathcal{Y}_t) = \text{Law}(\mathcal{X}_t \mid \mathcal{Y}_t) =: \eta_t \quad (3.4)$$

and $X_t := \bar{\eta}_t(\iota) = \eta_t(\iota)$ and $P_t = \mathcal{P}_{\bar{\eta}_t} = \mathcal{P}_{\eta_t}$ where X_t and P_t correspond to the Kalman-Bucy filter update and Riccati equations in (2.2) and (2.3).

We may refer to this specific class (3.2) of McKean-Vlasov-type diffusion as a Kalman-Bucy diffusion process [16]. The case (F1) corresponds to the limiting object that is sampled in the continuous-time version of the ‘vanilla’ EnKF [61]; while (F2) is the continuous-time limiting object that is sampled in the ‘deterministic’ EnKF of [126], see also [121]; and (F3) is a fully deterministic transport-inspired equation [121, 133]. Note that in this case (F3) the existence of the inverse of $\mathcal{P}_{\bar{\eta}_t}$ is given by the positive-definiteness properties of the solution of the Riccati equation in (2.3). In the next section we detail the Monte-Carlo ensemble filters derived from these Kalman-Bucy diffusion processes.

Note we may define a generalised version of case (F3) by,

$$\begin{aligned} \text{(F3')} \quad d\mathcal{X}_t &= A \mathcal{X}_t dt + R \mathcal{P}_{\bar{\eta}_t}^{-1} (\mathcal{X}_t - \bar{\eta}_t(\iota)) dt \\ &\quad + \mathcal{P}_{\bar{\eta}_t} H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t + \bar{\eta}_t(\iota)}{2} \right) dt \right] + G_t \mathcal{P}_{\bar{\eta}_t}^{-1} (\mathcal{X}_t - \bar{\eta}_t(\iota)) dt \end{aligned} \quad (3.5)$$

for any skew symmetric matrix $G_t' = -G_t$ that may also depend $\bar{\eta}_t$. This added tuning parameter may be related to an optimality metric, when deriving this transport equation from an optimal transport beginning. We may also write similar generalised versions (F1') and (F2') by adding $G_t \mathcal{P}_{\bar{\eta}_t}^{-1} (\mathcal{X}_t - \bar{\eta}_t(\iota))$ to (F1) and (F2); though practically it likely makes little sense.

4 Ensemble Kalman-Bucy Filtering

Ensemble Kalman-Bucy filters (**EnKF**) coincide with the mean-field particle interpretation of the nonlinear diffusion processes defined in (3.2).

Let $(\mathcal{V}_t^i, \mathcal{W}_t^i, \mathcal{X}_0^i)$ with $1 \leq i \leq N+1$ be $(N+1)$ independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. Again, we consider three different cases of Kalman-Bucy-type interacting diffusion process,

$$\begin{aligned}
(\text{F1}) \quad d\mathcal{X}_t^i &= A \mathcal{X}_t^i dt + R^{1/2} d\mathcal{V}_t^i + \hat{P}_t H' R_1^{-1} \left[d\mathcal{Y}_t - \left(H \mathcal{X}_t^i dt + R_1^{1/2} d\mathcal{W}_t^i \right) \right] \\
(\text{F2}) \quad d\mathcal{X}_t^i &= A \mathcal{X}_t^i dt + R^{1/2} d\mathcal{V}_t^i + \hat{P}_t H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t^i + \hat{X}_t}{2} \right) dt \right] \\
(\text{F3}) \quad d\mathcal{X}_t^i &= A \mathcal{X}_t^i dt + R \hat{P}_t^{-1} \left(\mathcal{X}_t^i - \hat{X}_t \right) dt + \hat{P}_t H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t^i + \hat{X}_t}{2} \right) dt \right]
\end{aligned} \tag{4.1}$$

with $1 \leq i \leq N+1$ and the rescaled (particle) sample mean and covariance

$$\begin{aligned}
\hat{\eta}_t &:= \eta_t^N = \frac{1}{N+1} \sum_{i=1}^{N+1} \delta_{\mathcal{X}_t^i} \\
\implies \hat{X}_t &:= X_t^N = \frac{1}{N+1} \sum_{i=1}^{N+1} \mathcal{X}_t^i \quad \text{and} \quad \hat{P}_t := P_t^N = \frac{N+1}{N} \mathcal{P}_{\hat{\eta}_t}
\end{aligned} \tag{4.2}$$

In cases (F1) and (F2) we have $N \geq 1$ and in case (F3) we require $N \geq d$ for the almost sure invertibility of $\hat{P}_t$ (although in case (F3) one may substitute a pseudo-inverse of $\hat{P}_t$ without changing the mathematical analysis). A sampled version of case (F3') may also be derived in the same way.

The filters of (4.1) are mean-field approximations of those in (3.2). In (4.1) we see the utility of the Kalman-Bucy filter formulation in (3.2). In particular, in (4.1) we have eliminated the classical Riccati matrix differential equation completely, and replaced it with an ensemble of (interacting) particle flows and the computation of a sample covariance matrix from this ensemble. The sample mean and covariance of (4.2) can also be used for inference or decision making, etc.

4.1 Vanilla Ensemble Kalman-Bucy Filter

The vanilla **EnKF**, denoted by **VEnKF**, is associated with the first case (F1) of nonlinear process $\mathcal{X}_t$ in (3.2) and is defined by the Kalman-Bucy-type interacting diffusion process (F1) in (4.1). We then have the following key result.

Proposition 4.1 ([54]). *Let $N \geq 1$. The stochastic flow of the sample mean satisfies,*

$$d\hat{X}_t \stackrel{\text{law}}{=} \left(A - \hat{P}_t S \right) \hat{X}_t dt + \hat{P}_t H' R_1^{-1} d\mathcal{Y}_t + \frac{1}{\sqrt{N+1}} \left(R + \hat{P}_t S \hat{P}_t \right)^{1/2} d\mathcal{B}_t \tag{4.3}$$

where $\mathcal{B}_t$ is an independent d -dimensional Brownian motion.

The sample covariance evolves according to a so-called matrix-valued Riccati diffusion process of the form,

$$d\hat{P}_t \stackrel{\text{law}}{=} \text{Ric}(\hat{P}_t) dt + \frac{2}{\sqrt{N}} \left[\hat{P}_t^{1/2} d\mathcal{M}_t \left(R + \hat{P}_t S \hat{P}_t \right)^{1/2} \right]_{\text{sym}} \tag{4.4}$$

where $\mathcal{M}_t$ is a $(d \times d)$ -matrix with independent Brownian entries (also independent of $\mathcal{B}_t$).

We see that for the vanilla **EnKF**, the convergence of $\hat{X}_t \rightarrow X_t$ and $\hat{P}_t \rightarrow P_t$ as $N \rightarrow \infty$ follows immediately. This result follows via the martingale representation theorem, e.g. Theorem 4.2 in [80], see also [55].

4.2 ‘Deterministic’ Ensemble Kalman-Bucy Filter

The ‘deterministic’ **EnKF**, denoted **DEnKF**, is associated with the second case (F2) of nonlinear process $\mathcal{X}_t$ in (3.2), and is defined by the Kalman-Bucy-type interacting diffusion process (F2) in (4.1). This ‘deterministic’ epithet in the **DEnKF** follows because the update ‘part’ of the particle flow is deterministic and does not rely on the stochastic perturbations by $\mathcal{W}_t^i$ appearing in the **VEnKF**. This name and idea was taken from [126]; see also [13, 121] and [137, 96]. We have the following key result.

Proposition 4.2 ([21, 19]). *Let $N \geq 1$. The stochastic flow of the sample mean satisfies,*

$$d\hat{X}_t \stackrel{law}{=} \left(A - \hat{P}_t S \right) \hat{X}_t dt + \hat{P}_t H' R_1^{-1} d\mathcal{Y}_t + \frac{1}{\sqrt{N+1}} R^{1/2} d\mathcal{B}_t \quad (4.5)$$

where $\mathcal{B}_t$ is an independent d -dimensional Brownian motion.

The sample covariance evolves according to a so-called matrix-valued Riccati diffusion process of the form,

$$d\hat{P}_t \stackrel{law}{=} \text{Ricc}(\hat{P}_t) dt + \frac{2}{\sqrt{N}} \left[\hat{P}_t^{1/2} d\mathcal{M}_t R^{1/2} \right]_{\text{sym}} \quad (4.6)$$

where $\mathcal{M}_t$ is a $(d \times d)$ -matrix with independent Brownian entries (also independent of $\mathcal{B}_t$).

Again, for the **DEnKF**, the convergence of $\hat{X}_t \rightarrow X_t$ and $\hat{P}_t \rightarrow P_t$ as $N \rightarrow \infty$ follows immediately. Note the simplified diffusion weighting(s) in the case of the **DEnKF**, as compared to the **VEnKF**.

4.3 Transport-Inspired Ensemble Transport Filter

The fully deterministic ensemble transport filter **DEnTF** is associated with the third case (F3); defined by the Kalman-Bucy-type interacting diffusion process (F3) in (4.1). In this case, we have the special result.

Proposition 4.3 ([121, 133]). *Let $N \geq 1$. The flow of sample mean is given by,*

$$d\hat{X}_t = \left(A - \hat{P}_t S \right) \hat{X}_t dt + \hat{P}_t H' R_1^{-1} d\mathcal{Y}_t, \quad \hat{X}_0 := \frac{1}{N+1} \sum_{i=1}^{N+1} \mathcal{X}_0^i \quad (4.7)$$

The sample covariance evolves according to the deterministic Riccati equation,

$$d\hat{P}_t = \text{Ricc}(\hat{P}_t) dt, \quad \hat{P}_0 := \frac{N+1}{N} \mathcal{P}_{\hat{\eta}_0} \quad (4.8)$$

Note that the particle mean $\hat{X}_t$ and the particle covariance $\hat{P}_t$ associated with the particle interpretation (F3) discussed in (4.1) satisfy exactly the equations of the Kalman-Bucy filter with the associated deterministic Riccati equation.

The ‘randomness’ in this case only comes from the initial conditions. The stability analysis of this class of **DEnTF** model resumes to the one of the Kalman-Bucy filter and the associated Riccati equation. Thus, the results, e.g., in (2.20), (2.22), (2.23) and (2.24) hold immediately; see also [16] in the linear-Gaussian setting. In [43, 44] this filter is analysed in the case of a nonlinear signal, but fully observed (linear observation) model. The fluctuation analysis in this case can also be developed easily by combining certain stability results w.r.t. the initial state (see [16]) with conventional sample estimates based on independent copies of the initial states (see e.g. [23] for estimates associated with classical sample covariance estimates). Consequently, we do not consider this class of model going forward, but recommend [16, 43, 44].

When N is small compared to d , the inverse of the sample covariance defining the **DEnTF** is ill-posed and this is likely a limiting factor in the applicability of this method in high-dimensional applications with stochastic state evolutions. With non-Gaussian signal noise, one may also prefer the stochastic perturbation method in the **DEnKF**.

4.4 Nonlinear Ensemble Filtering in Practice

In practice, the ensemble Kalman filtering methodology is applied in high-dimensional, nonlinear state-space models, e.g. see [60, 61] and the application references listed in the introduction.

It is rather straightforward to extend the algorithmic particle methods in (4.1) to nonlinear systems as we now outline. Consider a time-invariant nonlinear diffusion model of the form,

$$\begin{aligned} d\mathcal{X}_t &= a(\mathcal{X}_t) dt + R^{1/2} d\mathcal{V}_t \\ d\mathcal{Y}_t &= h(\mathcal{X}_t) dt + R_1^{1/2} d\mathcal{W}_t \end{aligned} \quad (4.9)$$

where $a : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $h : \mathbb{R}^d \rightarrow \mathbb{R}^{d_y}$ are the nonlinear signal and sensor model functions of some sufficient regularity.

Let $(\mathcal{V}_t^i, \mathcal{W}_t^i, \mathcal{X}_0^i)$ with $1 \leq i \leq N+1$ be $(N+1)$ independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. We consider the three **EnKF** variants as before and define the flow of particles by,

$$\begin{aligned} \text{(NF1)} \quad d\mathcal{X}_t^i &= a(\mathcal{X}_t^i) dt + R^{1/2} d\mathcal{V}_t^i + \hat{P}_t^h R_1^{-1} \left[d\mathcal{Y}_t - \left(h(\mathcal{X}_t^i) dt + R_1^{1/2} d\mathcal{W}_t^i \right) \right] \\ \text{(NF2)} \quad d\mathcal{X}_t^i &= a(\mathcal{X}_t^i) dt + R^{1/2} d\mathcal{V}_t^i + \hat{P}_t^h R_1^{-1} \left[d\mathcal{Y}_t - \left(\frac{h(\mathcal{X}_t^i) + \hat{h}_t}{2} \right) dt \right] \\ \text{(NF3)} \quad d\mathcal{X}_t^i &= a(\mathcal{X}_t^i) dt + R \hat{P}_t^{-1} (\mathcal{X}_t^i - \hat{X}_t) dt + \hat{P}_t^h R_1^{-1} \left[d\mathcal{Y}_t - \left(\frac{h(\mathcal{X}_t^i) + \hat{h}_t}{2} \right) dt \right] \end{aligned} \quad (4.10)$$

with $1 \leq i \leq N+1$ and the (particle) sample mean $\hat{X}_t$ and sample covariance $\hat{P}_t$ defined as usual, e.g. see (4.2), and with the observation function sample mean and sample cross-covariance defined as,

$$\hat{h}_t := \frac{1}{N+1} \sum_{i=1}^{N+1} h(\mathcal{X}_t^i) \quad \text{and} \quad \hat{P}^h := \frac{1}{N} \sum_{i=1}^{N+1} [\mathcal{X}_t^i - \hat{X}_t] [h(\mathcal{X}_t^i) - \hat{h}_t]' \quad (4.11)$$

The mean-field limit of these interacting nonlinear conditional particle diffusion systems (4.10) is studied in [43, 94]. The (conditional) law of these mean field McKean-Vlasov diffusions may even be given in terms of a Kushner/Fokker-Planck-type partial differential equation, e.g. see [43, 94]. However, if the mean-field limit in this nonlinear setting is denoted by, say, $\mathcal{X}_t$, then it is certainly true that,

$$\text{Law}(\mathcal{X}_t | \mathcal{Y}_t) \neq \text{Law}(\mathcal{X}_t | \mathcal{Y}_t) =: \eta_t \quad (4.12)$$

in the nonlinear model setting. Said differently, even with infinite computational power, the **EnKF** methods as applied in this nonlinear model setting do not converge to the optimal nonlinear Bayes filter. As noted earlier, and again later, the **EnKF** in this nonlinear model setting is probably best viewed in practice as a type of (random) sample-based (point-valued) state estimator or a stochastic observer. In general it should not be seen as an approximation of the optimal Bayesian filter.

We discuss connections and extensions of our results to the nonlinear model setting, including different instances of the **EnKF** in these settings, in a later section (at the end of this article).

5 Theory in the Linear-Gaussian Setting

Going forward, we consider only the **VEnKF** (case (F1)) and **DEnKF** (case (F2)) since as noted the theory of the **DEnTF** in the linear-Gaussian setting reverts to that of the standard Kalman-Bucy filter as detailed in [16]. The parameter $\kappa \in \{0, 1\}$ will distinguish the two cases ($\kappa = 1$ in case (F1), and $\kappa = 0$ in case (F2)) throughout.

We may unify the analysis via the following representation,

$$d\hat{X}_t = (A - \hat{P}_t S) \hat{X}_t dt + \hat{P}_t H' R_1^{-1} d\mathcal{Y}_t + \frac{1}{\sqrt{N+1}} \Sigma_\kappa^{1/2}(\hat{P}_t) d\mathcal{B}_t \quad (5.1)$$

$$d\hat{P}_t = \text{Ric}(\hat{P}_t) dt + \frac{2}{\sqrt{N}} \left[\hat{P}_t^{1/2} d\mathcal{M}_t \Sigma_\kappa^{1/2}(\hat{P}_t) \right]_{\text{sym}} \quad (5.2)$$

with the mapping,

$$\Sigma_\kappa(Q) := R + \kappa Q S Q \quad \text{with} \quad \kappa = \begin{cases} 1 & \text{in case (F1)} \\ 0 & \text{in case (F2)} \end{cases} \quad (5.3)$$

Let $\hat{Z}_t := (\hat{X}_t - \mathcal{X}_t)$ and observe that

$$\begin{aligned} d\hat{Z}_t &= (A - \hat{P}_t S) \hat{Z}_t dt + \hat{P}_t H' R_1^{-1/2} d\mathcal{W}_t - R^{1/2} d\mathcal{V}_t + \frac{1}{\sqrt{N+1}} \Sigma_\kappa^{1/2}(\hat{P}_t) d\mathcal{B}_t \\ &\stackrel{\text{law}}{=} (A - \hat{P}_t S) \hat{Z}_t dt + \Omega_\kappa^{1/2}(\hat{P}_t) d\hat{\mathcal{B}}_t \end{aligned} \quad (5.4)$$

for some independent d -dimensional Wiener process $\hat{\mathcal{B}}_t$ and with,

$$\Omega_\kappa := \Sigma_1 + \frac{1}{N+1} \Sigma_\kappa \quad (5.5)$$

Note we often refer to the flows $\hat{Z}_t$ or Z_t as error flows.

We also underline that

$$\hat{Z}_t - Z_t = (\hat{X}_t - \mathcal{X}_t) - (X_t - \mathcal{X}_t) = \hat{X}_t - X_t \quad (5.6)$$

so that the difference between the noisy error flow $\hat{Z}_t$ and the classical Kalman-Bucy error flow Z_t is equal to the difference between the **EnKF** (sample mean) state estimate and the classical Kalman-Bucy state estimate.

Let $\hat{\phi}_t(Q) := \hat{P}_t$ denote the flow of the Riccati diffusion equation in (5.2) with $\hat{P}_0 = Q \in \mathbb{S}_d^0$. Let $\hat{\psi}_t(z, Q) := \hat{Z}_t$ denote the flow of the stochastic error (5.4) with $\hat{Z}_0 = z = (x - \mathcal{X}_0) \in \mathbb{R}^d$ and $\hat{P}_t = \hat{\phi}_t(Q)$. Finally, we denote the flow of the sample mean in (5.1) with $\hat{X}_0 = x \in \mathbb{R}^d$ by $\hat{\chi}_t(x, Q) := \hat{X}_t$.

We underline further that the difference between two error flows satisfies,

$$\hat{\psi}_t(z_1, Q_1) - \hat{\psi}_t(z_2, Q_2) = \hat{\chi}_t(x_1, Q_1) - \hat{\chi}_t(x_2, Q_2) \quad (5.7)$$

and is thus equal to the difference between the two corresponding sample means (with compatible starting points). Studying the difference between two error flows ($\hat{\psi}_t(z_1, Q_1) - \hat{\psi}_t(z_2, Q_2)$) subsumes the study of something like ($\hat{\chi}_t(x_1, Q_1) - \hat{\chi}_t(x_2, Q_2)$) which is the difference between the **EnKF** (sample mean) state estimate and the classical Kalman-Bucy state estimate (with different initial conditions).

For any $s \leq t$ and $Q \in \mathbb{S}_d^0$ we define the stochastic state-transition matrix,

$$\hat{\mathcal{E}}_{s,t}(Q) := \exp \left(\int_s^t (A - \hat{\phi}_u(Q) S) du \right) \iff \partial_t \hat{\mathcal{E}}_{s,t}(Q) = (A - \hat{\phi}_t(Q) S) \hat{\mathcal{E}}_{s,t}(Q) \quad (5.8)$$

As with the classical Kalman-Bucy filter, e.g. see (2.16) and (2.17), the convergence and stability properties of the ensemble Kalman-Bucy filter and the associated Riccati diffusion equation are directly related to the contraction properties of the stochastic state-transition matrix $\hat{\mathcal{E}}_{s,t}(Q)$. For example, the flow of the stochastic error equation (5.4) is given by,

$$\hat{\psi}_t(z, Q) = \hat{\mathcal{E}}_{s,t}(Q) \hat{\psi}_s(z, Q) + \int_s^t \hat{\mathcal{E}}_{u,t}(Q) \Omega_\kappa^{1/2}(\hat{\phi}_u(Q)) d\hat{\mathcal{B}}_u \quad (5.9)$$

and the stochastic flow of the matrix Riccati diffusion (5.2) is given implicitly by

$$\begin{aligned}\hat{\phi}_t(Q) &= \hat{\mathcal{E}}_{s,t}(Q) \hat{\phi}_s(Q) \hat{\mathcal{E}}_{s,t}(Q)' + \int_s^t \hat{\mathcal{E}}_{u,t}(Q) \Sigma_1(\hat{\phi}_u(Q)) \hat{\mathcal{E}}_{u,t}(Q)' du \\ &\quad + \frac{2}{\sqrt{N}} \int_s^t \hat{\mathcal{E}}_{u,t}(Q) \left[\hat{\phi}_u^{1/2}(Q) d\mathcal{M}_u \Sigma_\kappa^{1/2}(\hat{\phi}_u(Q)) \right]_{\text{sym}} \hat{\mathcal{E}}_{u,t}(Q)' \end{aligned} \quad (5.10)$$

for any $s \leq t$. We denote by $\hat{\Pi}_t$ the Markov semigroup of $\hat{\phi}_t(Q)$ defined for any bounded measurable function F on $\mathbb{S}_d$ and any $Q \in \mathbb{S}_d^0$ with the property that,

$$\hat{\Pi}_t(F)(Q) := \mathbb{E} \left[F(\hat{\phi}_t(Q)) \right] = \int \hat{\Pi}_t(Q, dP) F(P) \quad (5.11)$$

When Q is random with distribution $\Gamma(dQ)$ on $\mathbb{S}_d^+$, by Fubini's theorem we have,

$$(\Gamma \hat{\Pi}_t)(F) := \int \Gamma(dQ) \hat{\Pi}_t(F)(Q) := \mathbb{E} \left[F(\hat{\phi}_t(Q)) \right] = \int \left(\int \Gamma(dQ) \hat{\Pi}_t(Q, dP) \right) F(P) \quad (5.12)$$

This yields the formula

$$(\Gamma \hat{\Pi}_t)(dP) = \int \Gamma(dQ) \hat{\Pi}_t(Q, dP) \quad (5.13)$$

for the distribution of $\hat{\phi}_t(Q)$ on $\mathbb{S}_d^+$.

We then have the first result concerning the quadratic, matrix-valued, Riccati diffusion process (5.10).

Theorem 5.1. *For any $N \geq 1$ the Riccati diffusion (5.10) has a unique weak solution on $\mathbb{S}_d^0$. For $N \geq d + 1$ there exists a unique strong solution on $\mathbb{S}_d^+$. Moreover, $\hat{\Pi}_t(Q, dP)$ is a strongly Feller and irreducible semigroup with a unique invariant probability measure $\hat{\Gamma}_\infty$ on $\mathbb{S}_d^+$. This measure admits a positive density with respect to the natural Lebesgue measure on $\mathbb{S}_d$.*

Given the existence of a solution to the Riccati diffusion (5.2), it follows a solution for $\hat{X}_t$ in (5.1) or a solution $\hat{Z}_t$ in (5.4) exists and is unique. This result is proven in [19, Theorem 2.1].

Once the problem of existence and uniqueness is tackled, one major problem in this equation is the behavior at infinity: existence of a stationary measure and speed of convergence towards this stationary measure or even distance between two solutions starting at different points.

We will make wide use of the following two assumptions in the remainder of this article.

Assumption O. The matrix $S := H'R_1^{-1}H$ is strictly positive-definite, i.e. $S \in \mathbb{S}_d^+$. This is a strong form of observability, and it implies classical observability as defined in (2.7).

Assumption C. The pair $(A, R^{1/2})$ is controllable, as defined in (2.7).

Under both Assumptions O and C it follows that $P_\infty \in \mathbb{S}_d^+$ and $\text{Abse}(A - P_\infty S) < 0$, see the earlier discussion on this topic. We may relax the controllability Assumption C to just stabilisability. We discuss Assumption O more later as it (re-)appears throughout our presentation and is more restrictive than the classical detectability assumptions in classical Kalman filtering.

We emphasise the following:

Suppose Assumptions O and C hold. Then there exists some logarithmic norm, which we denote by $\bar{\mu}(\cdot) : \mathbb{M}_d \rightarrow \mathbb{R}$, with the property that $\bar{\mu}(A - P_\infty S) < 0$.

Proof of this statement follows from the fact that $\text{Absc}(A - P_\infty S) < 0$ under just detectability and stabilisability model conditions, and then an application of [132, Theorem 5]. The logarithmic norm $\bar{\mu}(\cdot)$ is not necessarily unique, but any particular chosen logarithmic norm $\bar{\mu}(\cdot)$ is indexed to the model parameters (A, H, R, R_1) . We use the notation $\bar{\mu}(\cdot)$ to distinguish the log-norms for which $\bar{\mu}(A - P_\infty S) < 0$ whenever $\text{Absc}(A - P_\infty S) < 0$ holds, or more specifically throughout this work whenever Assumptions O and C hold.

In prior work [54, 22, 21, 17, 19] and even the first draft of this article, we state certain results in terms of $\mu(A - P_\infty S)$, and under the assumption $\mu(A - P_\infty S) < 0$; for some, but we don't care which, logarithmic norm $\mu(\cdot)$. We knew of course that certain observability and controllability model conditions ensured $\text{Absc}(A - P_\infty S) < 0$. However, it was unclear that negativity of the spectral abscissa translated in general to $\mu(A - P_\infty S) < 0$ for any version of the logarithmic norm. We thus, in many results, began with the assumption $\mu(A - P_\infty S) < 0$ in prior work [54, 22, 21, 17, 19], and claimed somewhat informally that this amounts to asking for a strong form of observability and controllability (given its similarity to $\text{Absc}(A - P_\infty S) < 0$, but without actually giving testable model conditions). Owing to [132, Theorem 5], we can begin results simply with some form of observability and controllability assumption (typically we need the stronger observability Assumption O, for different reasons) and state results in terms of the special class of logarithmic norms $\bar{\mu}(A - P_\infty S) < 0$; which we know is negative because $\text{Absc}(A - P_\infty S) < 0$. This is a significant relaxation of the conditions precedent in many of the subsequent results and places these results back in the testable and relatable context of classical controllability and observability assumptions.

In Table 1 we denote the relevant flows and notation of interest going forward. This notation allows us to relate (for example) the flow of the approximation relative to the true object with respect to their initial conditions, e.g. fluctuation-type results: $\hat{\chi}_t(x, Q) - \chi_t(x, Q)$; or (for example) the flow of two approximated objects with respect to different initial positions, e.g. stability/contraction-type results: $\hat{\psi}_t(z_1, Q_1) - \hat{\psi}_t(z_2, Q_2)$.

True Flow	Approximated Flow	Description
χ_t	$\hat{\chi}_t$	Flow of the Kalman state estimate (the true conditional mean) (2.2); and the sample mean (5.1)
ϕ_t	$\hat{\phi}_t$	Flow of the Kalman state estimate error covariance, i.e. the Riccati differential equation (2.3); and the sample covariance (5.2)
ψ_t	$\hat{\psi}_t$	Flow of the Kalman state error (2.6); and the sample error (5.4)
$\mathcal{E}_{s,t}$	$\hat{\mathcal{E}}_{s,t}$	The state transition matrix (2.12); and the stochastic, approximated, state transition matrix (5.8)
N.A.	$\hat{\Pi}_t$	The Markov transition kernel for $\hat{\phi}_t$ on $\mathbb{S}_d^0$.

Table 1: Table of flow notations.

In Figure 1 we plot the flow of some of the subsequent sections and the main results. The presentation ordering is given mostly in terms of the dependencies and natural progression of the derivations. We discuss briefly the dependencies and reasoning as we progress.

5.1 Fluctuation and Contraction Results for the Riccati Diffusion

5.1.1 Fluctuation Properties of the Riccati Diffusion

In this section we consider the fluctuation of $\hat{\phi}_t(Q)$ about $\phi_t(Q)$ and of $\hat{\psi}_t(z, Q)$ about $\psi_t(z, Q)$.




Figure 1: Flowchart of the general result and topics in this article. Although detailed proof in most cases is not given, the arrows and top-down direction in the flowchart depict both the presentation of the results in this article, and the dependency in terms of the proof and derivation of the results

The fluctuation properties and moment boundedness properties of $\hat{\phi}_t(Q)$ and $\hat{\psi}_t(z, Q)$ depend naturally on the size on the fluctuation as determined by N .

Typically, we will write either of the the following expressions in stating our results,

$$"N \text{ is sufficiently large}" \quad \text{or} \quad "N \geq 1" \quad (5.14)$$

In case (F1) with $\kappa = 1$ there is often a minimum threshold on N needed to prove the results. In case (F1) this lower threshold on N may be large. In case (F2) with $\kappa = 0$, these same results typically hold; but moreover, we can often refine the relevant results and at the same time relax the conditions on N , often needing just $N \geq 1$. This is a significant analytical advantage of the DEnKF over the VEnKF. In some cases, this advantage is practically realised and provable (and not just a by-product of analysis methods). For example, we will show later that some moments of the VEnKF sample covariance in one-dimension provably do not exist in the steady-state without a sufficient number of particles; whereas in the DEnKF these moments always exist with $N \geq 1$.

In some cases, the results stated in this work are only known for the **DEnKF**. If we do not specify a particular case, or a value for $\kappa \in \{0, 1\}$, then the stated results may be assumed to hold for both the **VEnKF** and the **DEnKF**.

We start with the following under-bias estimate on the sample covariance which holds for both the **VEnKF** and the **DEnKF**.

Theorem 5.2. *For any $t \geq 0$, any $Q \in \mathbb{S}_d^0$, and any $N \geq 1$, we have the uniform under-bias estimate,*

$$\mathbb{E} \left[\widehat{\phi}_t(Q) \right] \leq \phi_t(Q) \leq c(1 + \|Q\|) I \quad (5.15)$$

for a finite constant $c > 0$ that doesn't depend on the time horizon.

We may refine this under-bias result as is done in [19]. For example, if we assume further that $S \in \mathbb{S}_d^+$, i.e. under Assumption O, then for any $t \geq 0$ we also have the refined bias estimates,

$$0 \leq \phi_t(Q) - \mathbb{E} \left[\widehat{\phi}_t(Q) \right] \leq c(Q) \frac{1}{N} I \quad (5.16)$$

when N is sufficiently large in case (F1), $\kappa = 1$; or for any $N \geq 1$ in case (F2), $\kappa = 0$. The proof of this refinement, and details on the constant $c(Q)$, is in [19, Theorem 2.3] and in [22].

We will see subsequently that Assumption O, i.e. the condition $S \in \mathbb{S}_d^+$, ensures that for any $n \geq 1$, the n -th moments of the trace of the sample covariance are uniformly bounded w.r.t. the time horizon (with a sufficient number of particles) even when *the matrix A is unstable*.

The next theorem concerns these time-uniform moment estimates on the stochastic Riccati flow in (4.4), i.e. on the flow of the sample covariance matrix.

Theorem 5.3. *Suppose Assumption O holds. For any $n \geq 1$, $t \geq 0$, any $Q \in \mathbb{S}_d^0$, and any N sufficiently large, we have the uniform estimate,*

$$\mathbb{E} \left[\|\widehat{\phi}_t(Q)\|^n \right]^{1/n} \leq c_n (1 + \|Q\|) \quad (5.17)$$

Furthermore, for any time horizon $t \geq \tau > 0$ we also have the uniform estimates

$$\mathbb{E} \left[\|\widehat{\phi}_t(Q)\|^n \right]^{1/n} \leq c_{n,\tau} \quad (5.18)$$

In addition, in case (F2), for any $N \geq 1$, any $n \geq 1$, $Q \in \mathbb{S}_d^0$, $t \geq 0$ and any $s \geq \tau > 0$ we have the refined estimates,

$$\mathbb{E} \left[\|\widehat{\phi}_t(Q)\|^n \right]^{1/n} \leq c(1 + \|Q\|) \left(1 + \sqrt{\frac{n}{N}}\right) \quad \text{and} \quad \mathbb{E} \left[\|\widehat{\phi}_s(Q)\|^n \right]^{1/n} \leq c_\tau \left(1 + \sqrt{\frac{n}{N}}\right) \quad (5.19)$$

The proof of this result is provided in [19, Theorem 2.2] where a precise description of the (finite) parameters $c_n, c_{n,\tau}, c, c_\tau > 0$ is also provided. The first estimate in (5.17) also holds without Assumption O, and even if $S = 0$, when $\text{Absc}(A) < 0$. The proof of this Theorem is based on a reduction of (4.4) to a scalar Riccati diffusion, a novel representation of its n -th powers, and a comparison of its moments to a judiciously designed deterministic scalar Riccati equation. We discuss this proof later, but this scalar reduction necessitates the condition $S \in \mathbb{S}_d^+$, i.e. Assumption O. The proof is conservative by nature (due to the scalar reduction and comparison).

Now we turn to quantifying the fluctuations of the matrix Riccati diffusions around their limiting (deterministic) values as found when N tends to ∞ . That is, we quantify the fluctuation of the **EnKF** sample covariance about the limiting covariance of the classical Kalman-Bucy filter.

Theorem 5.4. *Suppose Assumption O holds. For any $n \geq 1$, $t \geq 0$, any $Q \in \mathbb{S}_d^0$, and any N sufficiently large we have the uniform estimates,*

$$\mathbb{E} \left[\|\hat{\phi}_t(Q) - \phi_t(Q)\|^n \right]^{1/n} \leq c_n \frac{1}{\sqrt{N}} (1 + \|Q\|^7) \quad (5.20)$$

In case (F2), for any $N \geq 1$, any $n \geq 1$, $t \geq 0$, and any $Q \in \mathbb{S}_d^0$, we have

$$\mathbb{E} \left[\|\hat{\phi}_t(Q) - \phi_t(Q)\|^n \right]^{1/n} \leq c \frac{1}{\sqrt{N}} (1 + \|Q\|^5) \left(1 + \sqrt{\frac{n}{N}} \right)^5 \quad (5.21)$$

The estimates in Theorem 5.4 do not depend on $Q \in \mathbb{S}_d^0$ when $t \geq \tau$ for any $\tau > 0$ and with c_n, c replaced with $c_{n,\tau}, c_\tau$; e.g. similarly to (5.18) in Theorem 5.3.

The proof of the preceding Theorem is provided in [19, Theorem 2.3] and in [22]. The proof follows from a second-order expansion of the stochastic flow $\hat{\phi}_t$ about the deterministic flow ϕ_t and then an appropriate bounding of the first and second order stochastic terms. More generally, in [22] we consider a Taylor-type perturbation expansion of the form,

$$\hat{\phi}_t = \phi_t + \sum_{1 \leq k < n} \frac{N^{-k/2}}{k!} \varphi_t^{(k)} + \frac{1}{N^{n/2}} \hat{\varphi}_t^{(n)} \quad (5.22)$$

for any $n \geq 1$, and a stochastic flow $\varphi_t^{(k)}$ whose values don't depend on the ensemble size N , and a stochastic remainder term $\hat{\varphi}_t^{(n)}$. Odd order stochastic terms $\varphi_t^{(k)}$, with k odd, are zero mean (i.e. centred). This representation allows us in [22] to present sharp and non-asymptotic expansions of the matrix moments of the matrix Riccati diffusion with respect to N .

In [22] we provide uniform estimates of the stochastic flow $\varphi_t^{(k)}$ w.r.t. the time horizon *even when the matrix A is unstable*. These estimates are stronger than the conventional functional central limit theorems for stochastic processes. For example, these results imply the almost sure central limit theorem on the sample covariance,

$$\sqrt{N} \left[\hat{\phi}_t - \phi_t \right] \xrightarrow{N \rightarrow \infty} \varphi_t \quad (5.23)$$

Bias and variance estimates based on the expansion (5.22) are also given in [22]. See also in particular [22, Section 1.3] for detailed exposition of this functional central limit theorem and the bias and variance estimates. In the scalar case, we explore this expansion (5.22) up to second-order in detail in a later section to illustrate this form.

The under bias result (5.15) holds with any $N \geq 1$ in both the **VENKF** of case (F1), and in the **DENKF** of case (F2). This under-bias is a motivation for so-called sample covariance regularisation in practice; e.g. so-called sample covariance inflation or localisation methods [7, 72, 66, 109, 60]. Later we discuss the effects of inflation in particular.

As with the deterministic Riccati equation, we may bound the moments of the inverse of the stochastic Riccati flow $\hat{\phi}_t(Q)$ under stronger conditions on the number of particles N required; e.g. see [19]. It follows that with $Q \in \mathbb{S}_d^+$ and with additional conditions on N , that for $t \geq \tau > 0$ there exists a uniform positive definite lower bound on $\mathbb{E}[\hat{\phi}_t(Q)]$.

A number of basic corollaries follow the proofs in [22, 19], for instance, we have the monotone property,

$$\mathbb{S}_d^0 \ni Q_1 \leq Q_2 \implies \mathbb{E} \left[\hat{\phi}_t(Q_1) \right] \leq \mathbb{E} \left[\hat{\phi}_t(Q_2) \right] \quad (5.24)$$

and, for any $Q \in \mathbb{S}_d^0$, the fixed upper bound,

$$\mathbb{E} \left[\hat{\phi}_t(Q) \right] \leq P_\infty + \mathcal{E}_t(P_\infty) (Q - P_\infty) \mathcal{E}_t(P_\infty)' \quad (5.25)$$

These estimates hold for any $N \geq 1$ without any additional assumptions, as in Theorem 5.2.

Several spectral estimates can be deduced from the estimates (5.16), (5.20) and (5.21). For example, in case (F2), with $\kappa = 0$ and $N \geq 1$ then combining (5.21) with the n -version of the Hoffman-Wielandt inequality we have the uniform estimate,

$$\sup_{1 \leq i \leq r} \mathbb{E} \left[\left\| \lambda_i \left(\widehat{\phi}_t(Q) \right) - \lambda_i(\phi_t(Q)) \right\|^n \right]^{1/n} \leq c_n(Q) \frac{1}{\sqrt{N}} \quad (5.26)$$

Finally, it is worth noting briefly that all moment boundedness and fluctuation results stated in this section hold with any $N \geq 1$ and without further assumptions, if one replaces the constants $c, c_n, c_\tau, c_n(Q), \dots$ with functions that now depend on (and grow with) the time horizon $t \geq 0$. However, if these bounds depend exponentially on time (as is quite typical in analysis), an exponent of the form $(\alpha t) > 200$ induces an exceedingly pessimistic estimate larger than the estimated number of elementary particles of matter in the visible universe. In this sense, non-time-uniform bounds of this form are clearly impractical from a numerical user-case perspective.

5.1.2 Contraction and Long Time Properties of the Riccati Diffusion

With $Q \in \mathbb{S}_d^+$ we set $\Lambda(Q) := \|Q\|_2 + \|Q^{-1}\|_2$ and we consider the collection of Λ -norms on the set of probability measures Γ_1, Γ_2 on $\mathbb{S}_d^+$, indexed by $\hbar > 0$, and defined by,

$$\|\Gamma_1 - \Gamma_2\|_{\hbar, \Lambda} := \sup | \Gamma_1(F) - \Gamma_2(F) | \quad (5.27)$$

In the above display, the supremum is taken over all measurable function F on $\mathbb{S}_d$ such that

$$\|F\|_{\Lambda} := \sup_{Q \in \mathbb{S}_d^+} \frac{|F(Q)|}{1 + \hbar \Lambda(Q)} \leq 1 \quad (5.28)$$

It is known that the deterministic Riccati equation that describes the flow of the covariance matrix in classical Kalman-Bucy filtering tends to a fixed point P_∞ for any initial point $Q \in \mathbb{S}_d^0$ when the (time-invariant) model (2.1) is detectable and stabilisable; e.g. see (2.20) and [16]. The next result is the analogue of this idea in the EnKF setting and describes the stability of the flow of the sample covariance.

Theorem 5.5. *Assume the fluctuation parameter N is sufficiently large such that $\mathbb{E}[\|\widehat{\phi}_t(Q)\|]$ and $\mathbb{E}[\|\widehat{\phi}_t^{-1}(Q)\|]$ are uniformly bounded (e.g. as in Theorem 5.3 for bounds on $\mathbb{E}[\|\phi_t(Q)\|]$). Then, there exists some finite constants $c, \alpha, \hbar > 0$ such that for any $t \geq 0$ and probability measures Γ_1, Γ_2 on $\mathbb{S}_d^+$, we have the Λ -norm contraction inequality*

$$\|\Gamma_1 \widehat{\Pi}_t - \Gamma_2 \widehat{\Pi}_t\|_{\hbar, \Lambda} \leq c e^{-\alpha t} \|\Gamma_1 - \Gamma_2\|_{\hbar, \Lambda} \quad (5.29)$$

Of course, setting $\Gamma_2 = \widehat{\Gamma}_\infty$ where $\widehat{\Gamma}_\infty$ is the unique invariant probability measure described in Theorem 5.1 implies that for any initial probability measure $Q \sim \Gamma$ on $\mathbb{S}_d^+$ we have that $\widehat{\phi}_t(Q)$ tends to be distributed according to $\widehat{\Gamma}_\infty$. The proof of the above theorem is provided in [19, Theorem 2.4] and is based on matrix-valued Lyapunov and minorisation conditions (choosing the Lyapunov candidate, $\Lambda(\cdot)$).

For one-dimensional models, the article [21] provides explicit analytical expressions for the reversible measure of $\widehat{P}_t$ in terms of the model parameters. As expected, heavy tailed reversible measures arise when $\kappa = 1$, and weighted Gaussian distributions when $\kappa = 0$. The article [21] also provides sharp exponential decay rates to equilibrium, in the sense that the decay rates tend to those of the limiting deterministic Riccati equation when N tends to ∞ .

In a later section, we explore the one-dimensional case in more detail and explicitly examine the invariant measures in each model $\kappa \in \{0, 1\}$. The contrast between the steady-state invariant measures in each case $\kappa \in \{0, 1\}$ provides some insight into various phenomenon seen in practice we believe, e.g. so-called catastrophic divergence, and fluctuations of the sample covariance, etc. We also state the strong $\mathbb{L}_n$ -type contraction of $\widehat{\phi}_t(Q)$ in both cases (F1) and (F2).

5.2 Contraction Properties of Exponential Semigroups

Recall that the stability properties of the deterministic ($N = \infty$) semigroups $\mathcal{E}_{s,t}(Q)$ associated with the classical Kalman-Bucy filter are rather well understood; e.g. see (2.13), (2.15), and (2.14) and also [16, 18]. We emphasise that in the deterministic case, stability of the matrix-valued Riccati differential equation, e.g. as in (2.20), follows from the contraction properties of $\mathcal{E}_{s,t}(Q)$ in (2.13); see [16, 18] for the derivation. Some intuition for this follows from the implicit form for the solution in (2.17). Similarly, in classical Kalman-Bucy filter, the stability properties of the error flow (2.6) are related to the contraction properties of the state-transition matrix $\mathcal{E}_{s,t}(Q)$. Again, the intuition follows from the solution form in (2.16). The stability properties of the classical Kalman-Bucy error flow are given in, e.g., (2.22) and (2.24); see [16].

We come now to the contractive properties of $\widehat{\mathcal{E}}_{s,t}(Q)$ defined in (5.8). The stability of $\widehat{\mathcal{E}}_{s,t}(Q)$ will naturally play a role in the derivation of contraction results on, e.g., the sample error flow $\widehat{\psi}_t(z, Q)$, see (5.9). Indeed, we also require stability of $\widehat{\mathcal{E}}_{s,t}(Q)$ to derive fluctuation results on the sample error flow $\widehat{\psi}_t(z, Q)$. Note we did not need stability of the exponential semigroup to derive fluctuation results on the sample covariance $\widehat{\phi}_t(z, Q)$ earlier.

Firstly, we remark that if $S \in \mathbb{S}_d^+$, then up to a change of basis we can always assume that $S = I$. Then, for any $s, t \in [0, \infty[$ we immediately have the rather crude almost sure estimate

$$\mu(A) < 0 \quad \implies \quad \left\| \widehat{\mathcal{E}}_{s,s+t}(Q) \right\|_2 \leq e^{t\mu(A)} \xrightarrow{t \rightarrow \infty} 0 \quad (5.30)$$

for any logarithmic norm. Note again that if $\text{Absc}(A) < 0$, then $\mu(A) < 0$ for some log-norm. In any case, in general, asking for A to be stable is a very strong and restrictive condition. We typically seek contraction results on $\widehat{\mathcal{E}}_{s,t}(Q)$ that accomodate arbitrary $A \in \mathbb{M}_d$ matrices; in particular, we seek to accomodate unstable signal matrices A , i.e. matrices with (some) non-negative eigenvalues. To this end, fix $Q \in \mathbb{S}_d^0$ and consider the process $\widehat{\mathcal{A}}$ defined by

$$\widehat{\mathcal{A}} : t \in [0, \infty[\mapsto \widehat{\mathcal{A}}_t := A - \widehat{\phi}_t(Q)S \quad (5.31)$$

We write $\mathcal{A}$ for the analogous process driven by $\phi_t(Q)$, i.e. with $N = \infty$; which we know under just detectability conditions is a time-varying stabilising matrix process [140].

We seek to characterise, in a useful manner, the fluctuation of the stochastic process $\widehat{\mathcal{A}}$ about $\mathcal{A}$; with the hope that the contractive properties of $\widehat{\mathcal{E}}_{s,t}$ can then be in some sense related to the established contractive properties of $\mathcal{E}_{s,t}$.

For example, given Assumption O and $\kappa = 0$, combining (5.19) (5.21) and (2.15) with Krause's inequality [87] for any $nd \geq 1$ we have the uniform fluctuation estimate,

$$\mathbb{E} \left[\left\| \text{SpecDist} \left(\mathcal{A}_t, \widehat{\mathcal{A}}_t \right) \right\|^{nd} \right]^{1/(nd)} \leq c_n(Q) \frac{1}{\sqrt{N}} \quad (5.32)$$

where we define the optimal matching distance between the spectrum of matrices $A, B \in \mathbb{M}_d$ by

$$\text{SpecDist}(A, B) = \min_{\text{perm}(\cdot)} \max_{1 \leq i \leq d} |\lambda_i(A) - \lambda_{\text{perm}(i)}(B)| \quad (5.33)$$

where the minimum is taken over the set of $d!$ permutations of $\{1, \dots, d\}$. This spectral estimate is of interest on its own, but is not immediately usable for controlling the contraction properties of the exponential semigroups.

By Theorem 5.3 and Theorem 5.4, under Assumption O, the collection of processes $(\mathcal{A}, \widehat{\mathcal{A}})$ satisfy the following regularity properties:

- *Case $\kappa \in \{1, 0\}$:* For any $n \geq 1$, $t \geq 0$, $Q \in \mathbb{S}_d^0$, and any N sufficiently large we have the uniform estimates

$$\sqrt{N} \mathbb{E} \left[\left\| \mathcal{A}_t - \widehat{\mathcal{A}}_t \right\|^n \right]^{\frac{1}{n}} \leq c_n (1 + \|Q\|^7) \quad \text{and} \quad \mathbb{E} \left[\left\| \widehat{\mathcal{A}}_t \right\|^n \right]^{\frac{1}{n}} \leq c_n (1 + \|Q\|) \quad (5.34)$$

- *Case $\kappa = 0$:*

For any $n \geq 1$, $t \geq 0$, $Q \in \mathbb{S}_d^0$, and any $N \geq 1$ we have the uniform estimates

$$\begin{aligned} \sqrt{N} \mathbb{E} \left[\left\| \mathcal{A}_t - \hat{\mathcal{A}}_t \right\|^n \right]^{\frac{1}{n}} &\leq c(1 + \|Q\|^5) \left(1 + \frac{\sqrt{n}}{\sqrt{N}}\right)^5, \\ \text{and} \\ \sqrt{N} \mathbb{E} \left[\left\| \hat{\mathcal{A}}_t \right\|^n \right]^{\frac{1}{n}} &\leq c(1 + \|Q\|) (1 + \sqrt{n}) \end{aligned} \quad (5.35)$$

The stability properties of stochastic semigroups associated with a general collection of stochastic flows $(\mathcal{A}, \hat{\mathcal{A}})$ satisfying fluctuation and moment boundedness properties in a general form accommodating both (5.34) and (5.35) have been developed in our prior work [17]. Several local-type contraction estimates can now be derived.

Theorem 5.6. *Let $\kappa \in \{1, 0\}$ and suppose Assumptions O and C hold. Then for any increasing sequence $0 \leq s \leq t_k \uparrow_{k \rightarrow \infty} \infty$, and for any $Q \in \mathbb{S}_d^0$, the probability of the following event*

$$\limsup_{k \rightarrow \infty} \frac{1}{t_k} \log \|\hat{\mathcal{E}}_{s, t_k}(Q)\| \leq \frac{1}{2} \bar{\mu}(A - P_\infty S) < 0 \quad \text{is greater than } 1 - \nu \quad (5.36)$$

for any $\nu \in]0, 1[$, as soon as N is sufficiently large (as a function of $\nu \in]0, 1[$).

This log-Lyapunov estimate (5.36) immediately implies the semigroup $\hat{\mathcal{E}}_{s, t_k}(Q)$ is exponentially contracting with a high probability (in both cases $\kappa \in \{1, 0\}$); given a sufficient number of particles, and the observability and controllability Assumptions O and C.

A number of reformulations of this result that offer insight individually are worth stating:

- Let $\kappa \in \{1, 0\}$. For any $0 \leq s \leq t_{k_1} \uparrow_{k_1 \rightarrow \infty} \infty$, there exists a sequence $N := N_{k_2} \uparrow_{k_2 \rightarrow \infty} \infty$ such that we have the almost sure Lyapunov estimate

$$\limsup_{k_2 \rightarrow \infty} \limsup_{k_1 \rightarrow \infty} \frac{1}{t_{k_1}} \log \|\hat{\mathcal{E}}_{s, s+t_{k_1}}(Q)\| \leq \frac{1}{2} \bar{\mu}(A - P_\infty S) \quad (5.37)$$

- Let $\kappa \in \{1, 0\}$. Then, for any increasing sequence of times $0 \leq s \leq t_k \uparrow_{k \rightarrow \infty} \infty$, the probability of the following event,

$$\left\{ \begin{array}{l} \forall 0 < \nu_2 \leq 1 \quad \exists l \geq 1 \quad \text{such that} \quad \forall k \geq l \quad \text{it holds that} \\ \frac{1}{t_k} \log \|\hat{\mathcal{E}}_{s, t_k}(Q)\| \leq \frac{1}{2} (1 - \nu_2) \bar{\mu}(A - P_\infty S) \end{array} \right\} \quad (5.38)$$

is greater than $1 - \nu_1$, for any $\nu_1 \in]0, 1[$, as soon as N is sufficiently large (as a function of $n \geq 1$ and $\nu_1 \in]0, 1[$).

- Let $\kappa \in \{1, 0\}$. Consider any $s \geq 0$, any increasing sequence of time horizons $t_k \uparrow_{k \rightarrow \infty} \infty$, and any sequence $N := N_{k_2, n} \uparrow_{k_2 \rightarrow \infty} \infty$ such that $\sum_{k_2 \geq 1} 1/\sqrt{N_{k_2, n}} < \infty$ for some $n \geq 1$. Then, we have the almost sure Lyapunov estimate,

$$\left\{ \begin{array}{l} \forall 0 < \nu \leq 1 \quad \exists l_1, l_2 \geq 1 \quad \text{such that} \quad \forall k_1 \geq l_1, \forall k_2 \geq l_2 \quad \text{it holds that} \\ \frac{1}{t_{k_1}} \log \|\hat{\mathcal{E}}_{s, s+t_{k_1}}(Q)\| \leq \frac{1}{2} (1 - \nu) \bar{\mu}(A - P_\infty S) \end{array} \right\} \quad (5.39)$$

The first dot-point result captured by (5.37) is derived from (5.36) in Theorem 5.6 via the Borel-Cantelli lemma. The next two dot-point results provide some reformulation of the supremum limit estimates (5.36) and (5.37) in terms of random relaxation time horizons and random relaxation-type fluctuation parameters. The last reformulation in (5.39) underlines the fact that after some random time (i.e. determined by l_1), and given some randomly sufficiently large number of particles (determined by l_2) the semigroup $\widehat{\mathcal{E}}_{s,t}(Q)$ is exponentially contractive. We have no direct control over the parameters l_1 and l_2 in (5.39) which depend on the randomness in any realisation.

Stronger results hold if we restrict $\kappa = 0$, i.e. in case (F2). We have the following immediate corollary of our prior work in [17] and the earlier fluctuation analysis leading to (5.35):

Theorem 5.7. *Let $\kappa = 0$ and suppose Assumptions O and C hold. Then, for any $n \geq 1$, $s \geq 0$, $Q \in \mathbb{S}_d^0$, there is some time horizons $\mathbf{t}_n < \widehat{\mathbf{t}}_n \rightarrow_{\mathbf{N} \rightarrow \infty} \infty$ such that for any $\mathbf{t}_n \leq t \leq \widehat{\mathbf{t}}_n$ we have*

$$\frac{1}{t} \log \mathbb{E} \left[\|\widehat{\mathcal{E}}_{s,s+t}(Q)\|^n \right] \leq \frac{n}{4} \bar{\mu}(A - P_\infty S) < 0 \quad (5.40)$$

whenever $\mathbf{N}$ is sufficiently large such that $\widehat{\mathbf{t}}_n > \mathbf{t}_n$; see [17] for details on these time parameters.

Importantly, in this last result we have $\widehat{\mathbf{t}}_n \rightarrow_{\mathbf{N} \rightarrow \infty} \infty$ and thus we can control (via $\mathbf{N}$) the horizon on which the semigroup $\widehat{\mathcal{E}}_{s,t}(Q)$ is asymptotically $\mathbb{L}_n$ -stable for any $n \geq 1$ when $\kappa = 0$. In other words, the estimate (5.40) ensures that the stochastic semigroup $\widehat{\mathcal{E}}_{s,t}(Q)$ is stable on arbitrary long finite time horizons, as soon as $\kappa = 0$, and when the ensemble size is sufficiently large. We have the following fact immediate from Theorem 5.7:

- Assume $\kappa = 0$. For any $n \geq 1$, $s \geq 0$, we have

$$\limsup_{\mathbf{N} \rightarrow \infty} \frac{1}{\widehat{\mathbf{t}}_n} \log \mathbb{E} \left[\|\widehat{\mathcal{E}}_{s,s+\widehat{\mathbf{t}}_n}(Q)\|^n \right] \leq \frac{n}{4} \bar{\mu}(A - P_\infty S)$$

Combining Theorem 5.6 and Theorem 5.7 we may draw the basic (qualitative) conclusion that, after some initial time period, and given enough particles, the (noisy) exponential semigroups $\widehat{\mathcal{E}}_{s,t}(Q)$ are exponentially contractive (in some sense, e.g. almost-sure or $\mathbb{L}_n$ -type) at a rate related to a logarithmic norm $\bar{\mu}(A - P_\infty S)$.

We remind the reader again that weak detectability and stabilisability assumptions ensure $\text{Absc}(A - P_\infty S) < 0$ and consequently, via the earlier discussion and [132, Theorem 5], there exists some logarithmic norm such that $\bar{\mu}(A - P_\infty S) < 0$. Assumptions O and C imply weak detectability and stabilisability.

Finally, we also have the following new result which extends the exponential decay results for one-dimensional models presented in [21] to the determinant of the matrix-valued Riccati diffusions considered herein. This is a type of stochastic Liouville formula.

Theorem 5.8. *Suppose Assumptions O and C hold. Then, for any $n \geq 1$, $t \geq 0$, any $Q \in \mathbb{S}_d^+$, and $\mathbf{N}$ sufficiently large we have the exponential decay estimate*

$$\begin{aligned} \mathbb{E} \left[\det(\widehat{\mathcal{E}}_t(Q))^n \right]^{1/n} &= \mathbb{E} \left[\exp \left(n \int_0^t \text{Tr}(A - \widehat{\phi}_s(Q)S) ds \right) \right]^{1/n} \\ &\leq c_n(Q) \exp \left(-t \sqrt{\text{Tr}(\widehat{R}_n \widehat{S}_n)} \right) \end{aligned} \quad (5.41)$$

with

$$\widehat{R}_n := R \left(1 - \frac{1}{\mathbf{N}}(2n + d + 1) \right) > 0 \quad \text{and} \quad \widehat{S}_n := S \left(1 - \frac{1}{\mathbf{N}}(2n + d + 1)\kappa \right) > 0 \quad (5.42)$$

In addition, there exists some function $\widehat{\nu}_n$ with $\lim_{\mathbf{N} \rightarrow \infty} \widehat{\nu}_n = 0$ such that

$$\mathbb{E} \left[\det(\widehat{\mathcal{E}}_t(Q))^n \right]^{1/n} \leq c_n(Q) \exp \left(-t(1 - \widehat{\nu}_n) \sqrt{\text{Tr}(A)^2 + \text{Tr}(RS)} \right) \quad (5.43)$$

The proof of this theorem is in [19, Theorem 2.7]. In the one-dimensional case, $d = 1$, this result collapses to capture the strong exponential contraction results presented in [21]. Indeed in one dimension, Theorem 5.8 can be seen as a significant improvement over both Theorem 5.6 and Theorem 5.7 in both theoretical development and practical usability.

In the scalar case, strong stability results on the stochastic Riccati flow $\hat{\phi}_t$ analogous to the deterministic setting, e.g. (2.20), also follow from Theorem 5.8; see also [21] and the results and illustrative examples in a later section in this article.

5.3 Fluctuation and Stability of the Ensemble Kalman-Bucy Filter

In this section we consider the fluctuation of the sample mean $(\hat{\chi}_t(x, Q) := \hat{X}_t$ with $x \in \mathbb{R}^d$ and $\hat{P}_0 = Q \in \mathbb{S}_d^0$; or more typically the sample mean error $\hat{\psi}_t(z, Q) := \hat{Z}_t = (\hat{X}_t - \mathcal{X}_t)$ with $\hat{Z}_0 = (x - \mathcal{X}_0) = z \in \mathbb{R}^d$. We also consider the contraction properties of the error flow of $\hat{\psi}_t(z, Q)$. This flow may be related to the Ornstein-Uhlenbeck process (5.4) and whose solution can be written more generally as in (5.9).

The first result is a fluctuation result of the ensemble sample mean about the Kalman-Bucy filter estimate, i.e. the true conditional mean; and also a conditional bias, or fluctuation, result on the conditional expectation of the ensemble sample mean given the observation sequence, with respect to the true conditional mean given by the Kalman-Bucy filter.

The first result is given under the strong assumption that the latent signal is stable, i.e. $\text{Absc}(A) < 0$, and this result holds for both the **VEnKF** and the **DEnKF**.

Theorem 5.9. *Let $\kappa \in \{1, 0\}$ and suppose Assumption O holds and $\text{Absc}(A) < 0$. For any $n \geq 1$, any $x \in \mathbb{R}^d$, any $Q \in \mathbb{S}_d^0$, and for $N \geq 1$ sufficiently large, we have the fluctuation estimate,*

$$\mathbb{E} \left[\left\| \hat{\chi}_t(x, Q) - \chi_t(x, Q) \right\|^n \right]^{1/n} \leq c_n(x, Q) \frac{1}{\sqrt{N}} \quad (5.44)$$

We also have the conditional bias estimate,

$$\mathbb{E} \left[\left\| \mathbb{E} [\hat{\chi}_t(x, Q) | \mathcal{Y}_t] - \chi_t(x, Q) \right\|^n \right]^{1/n} \leq c_n(x, Q) \frac{1}{N} \quad (5.45)$$

Proof of the fluctuation estimate (5.44) is given in [54]. Proof of the conditional bias estimate (5.45) is given in [40, Theorem 2.4]. The latter result (5.45) is used in [40] to study the estimation of the log-normalization constant associated with a class of continuous-time filtering models.

The next theorem concerns time-uniform moment estimates on the sample mean error; and the fluctuation of the sample mean error around its limiting value (found when N tends to ∞). The next result relaxes the assumption that the latent signal be stable.

Theorem 5.10. *Consider only case (F2) and suppose Assumptions O and C hold. Then for any $n \geq 1$, $z \in \mathbb{R}^d$, $Q \in \mathbb{S}_d^0$, there exists a time $\hat{\mathbf{t}}_n \rightarrow_{N \rightarrow \infty} \infty$ such that for any $0 \leq t \leq \hat{\mathbf{t}}_n$ we have,*

$$\mathbb{E} \left[\left\| \hat{\psi}_t(z, Q) \right\|^n \right]^{1/n} \leq c_n(z, Q) \quad (5.46)$$

and

$$\mathbb{E} \left[\left\| \hat{\psi}_t(z, Q) - \psi_t(z, Q) \right\|^n \right]^{1/n} \leq c_n(z, Q) \frac{1}{\sqrt{N}} \quad (5.47)$$

See [17] for details on the time parameter $\hat{\mathbf{t}}_n \rightarrow_{N \rightarrow \infty} \infty$.

Note again the difference $(\hat{\psi}_t(z_1, Q_1) - \psi_t(z_2, Q_2))$ resumes to that of $(\hat{\chi}_t(x_1, Q_1) - \chi_t(x_2, Q_2))$. Thus, e.g., (5.47) is comparable to (5.44), under different antecedent conditions.

Unlike Theorem 5.3 and Theorem 5.4, the proof of both Theorem 5.9 and Theorem 5.10 requires contraction properties to be established a priori for the stochastic transition matrix

$\widehat{\mathcal{E}}_{s,t}(Q)$ defined in (5.8). Hence, in Theorem 5.9 we rely on $\text{Absc}(A) < 0$ which ensures the contractive property holds for $\widehat{\mathcal{E}}_{s,t}(Q)$, see (5.30). In Theorem 5.10, we rely on Theorem 5.7 which establishes the $\mathbb{L}_n$ -contractivity of $\widehat{\mathcal{E}}_{s,t}(Q)$ without asking for A to be stable, but only in the case of the **DEnKF** with $\kappa = 0$, at least in the multi-dimensional setting.

The proof of Theorem 5.10 is provided in [21] in the one-dimensional setting where a detailed description of the (finite) parameters $c_n(z, Q) > 0$ are provided. The multi-dimensional result follows using similar proof methods to those used in [21] in combination with the contraction properties of the transition matrix $\widehat{\mathcal{E}}_{s,t}(Q)$ established in Theorem 5.7. In the one-dimensional setting studied in [21], contraction of $\widehat{\mathcal{E}}_{s,t}(Q)$ is given under very general model conditions which also accommodate both the **VEnKF** and the **DEnKF**. Consequently, in one dimension Theorem 5.10 holds on an infinite time horizon for any $t \geq 0$ and with any $\kappa \in \{0, 1\}$.

One may consider a perturbation expansion of the sample mean flow as

$$\widehat{\psi}_t = \psi_t + \sum_{1 \leq k < n} \frac{N^{-k/2}}{k!} \vartheta_t^{(k)} + \frac{1}{N^{n/2}} \widehat{\vartheta}_t^{(n)} \quad (5.48)$$

for any $n \geq 1$, and some stochastic flow $\vartheta_t^{(k)}$ that does not depend on the ensemble size N , and some stochastic remainder term $\widehat{\vartheta}_t^{(n)}$. This implies the almost sure central limit theorem on the sample mean,

$$\sqrt{N} [\widehat{\psi}_t - \psi_t] \xrightarrow{N \rightarrow \infty} \vartheta_t \quad (5.49)$$

See in particular [22, Section 1.3] for detailed exposition of this functional central limit theorem.

Uniform propagation of chaos follows from the proceeding central limit theorems and the development in this subsection. In particular we have,

$$\text{Law}(\mathcal{X}_t^i) \xrightarrow{N \rightarrow \infty} \text{Law}(\mathcal{X}_t), \quad \forall t \geq 0, \quad \forall i \in \{1, \dots, N+1\} \quad (5.50)$$

in some suitable metric (e.g. Wasserstein).

Now we turn to the stability of the error flow $\widehat{\psi}_t(z, Q)$ and its contraction properties. The subsequent study on the stability of $\widehat{\psi}_t(z, Q)$ relies again on the contraction of $\widehat{\mathcal{E}}_{s,t}$ studied previously.

The following uniform error contraction estimate follows from (5.9) and Theorem 5.7,

$$\sup_{Q \in \mathbb{S}_d^0} \left\| \mathbb{E} [\widehat{\psi}_t(z, Q) | \mathcal{X}_0] \right\| \leq c e^{t \alpha \bar{\mu}(A - P_\infty S)} \|x - \mathcal{X}_0\| \quad (5.51)$$

and holds for the **DEnKF**, with $\kappa = 0$, for some $\alpha, c > 0$, and under conditions compatible with the conditions in Theorem 5.7. This contraction result is analogous to (2.22) for the classical Kalman-Bucy filter; but under stronger conditions dictated by the available results on the contraction properties of $\widehat{\mathcal{E}}_{s,t}$ stated in Theorem 5.7. In particular, our methods prove this contraction (5.51) only in the case of the **DEnKF**, with $\kappa = 0$, with N sufficiently large, and on time horizons compatible with those detailed in Theorem 5.7.

If $\text{Absc}(A) < 0$ and Assumption O holds, then (5.51) holds on any infinite time horizon for both the **VEnKF** and **DEnKF**; because in this case $\widehat{\mathcal{E}}_{s,t}(Q)$ is contractive from (5.30). This is analogous to the setting of Theorem 5.9, as compared to that of Theorem 5.10; in line with the earlier discussion on the conditions leading to stability of $\widehat{\mathcal{E}}_{s,t}(Q)$.

The next results on the stability of $\widehat{\psi}_t(z, Q)$ similarly follow immediately from those stability results in the preceding section, but are stated at the level of the process $\widehat{\psi}_t(z, Q)$ itself, rather than the stochastic exponential semigroup $\widehat{\mathcal{E}}_{s,t}$.

Theorem 5.11. *Let $\kappa \in \{1, 0\}$ and suppose Assumptions O and C hold. Then for any increasing sequence of times $t_k \uparrow_{k \rightarrow \infty} \infty$, any $z_1 \neq z_2$ and any $Q \in \mathbb{S}_d^0$, the probability of the following event*

$$\limsup_{k \rightarrow \infty} \frac{1}{t_k} \log \|\widehat{\psi}_{t_k}(z_1, Q) - \widehat{\psi}_{t_k}(z_2, Q)\| < \alpha \bar{\mu}(A - P_\infty S) \quad \text{is greater than } 1 - \nu \quad (5.52)$$

for any $\nu \in]0, 1[$ and some $\alpha > 0$, as soon as $\mathbf{N}$ is sufficiently large (as a function of ν).

Two reformulations of this result may shed insight individually and are worth highlighting:

- Let $\kappa \in \{1, 0\}$. For $0 \leq t_{k_1} \uparrow_{k_1 \rightarrow \infty} \infty$, there exists a sequence $\mathbf{N} := \mathbf{N}_{k_2} \uparrow_{k_2 \rightarrow \infty} \infty$ such that we have the almost sure Lyapunov estimate

$$\limsup_{k_2 \rightarrow \infty} \limsup_{k_1 \rightarrow \infty} \frac{1}{t_{k_1}} \log \|\hat{\psi}_{t_{k_1}}(z_1, Q) - \hat{\psi}_{t_{k_1}}(z_2, Q)\| < \alpha \bar{\mu}(A - P_\infty S) \quad (5.53)$$

- Let $\kappa \in \{1, 0\}$. Consider any increasing sequence of time horizons $t_k \uparrow_{k \rightarrow \infty} \infty$, and any sequence $\mathbf{N} := \mathbf{N}_{k_2, n} \uparrow_{k_2 \rightarrow \infty} \infty$ such that $\sum_{k_2 \geq 1} 1/\sqrt{\mathbf{N}_{k_2, n}} < \infty$ for some $n \geq 1$. Then, we have the almost sure Lyapunov estimate,

$$\left\{ \begin{array}{l} \forall 0 < \nu \leq 1 \quad \exists l_1, l_2 \geq 1 \quad \text{such that} \quad \forall k_1 \geq l_1, \forall k_2 \geq l_2 \quad \text{it holds that} \\ \frac{1}{t_{k_1}} \log \|\hat{\psi}_{t_{k_1}}(z_1, Q) - \hat{\psi}_{t_{k_1}}(z_2, Q)\| \leq \alpha(1 - \nu) \bar{\mu}(A - P_\infty S) \end{array} \right\} \quad (5.54)$$

Again we emphasise that the reformulation in (5.54) highlights that after some random time (i.e. determined by l_1), and given a random sufficiently large number of particles (determined by l_2) the difference of error flows (or sample means; see (5.7)) is exponentially stable.

We have stronger $\mathbb{L}_n$ -type stability results in those settings analogous to Theorems 5.9 and 5.10 and in line with the discussion after Theorem 5.10 on the conditions for of $\hat{\mathcal{E}}_{s,t}(Q)$.

Theorem 5.12. *Let $\kappa \in \{1, 0\}$ and suppose Assumption O holds and $\text{Absc}(A) < 0$. Then for any $n \geq 1$, any $z_1 \neq z_2$, and any $Q \in \mathbb{S}_d^0$ we have the stability estimate,*

$$\mathbb{E} \left[\|\hat{\psi}_t(z_1, Q) - \hat{\psi}_t(z_2, Q)\|^n \right]^{1/n} \leq c_n(z_1, z_2, Q) \|z_1 - z_2\| e^{t \bar{\mu}(A - P_\infty S)} \quad (5.55)$$

whenever $\mathbf{N}$ is sufficiently large.

In the case (F2), i.e. for the **DEnKF** only, when $\kappa = 0$, we can relax the strong assumption that the latent signal be stable.

Theorem 5.13. *Let $\kappa = 0$ and suppose Assumptions O and C hold. Then for any $n \geq 1$, any $z_1 \neq z_2$, and any $Q \in \mathbb{S}_d^0$, there exists some time horizons $\mathbf{t}_n < \hat{\mathbf{t}}_n \rightarrow_{\mathbf{N} \rightarrow \infty} \infty$ such that for any $\mathbf{t}_n \leq t \leq \hat{\mathbf{t}}_n$ we have the stability estimate,*

$$\mathbb{E} \left[\|\hat{\psi}_t(z_1, Q) - \hat{\psi}_t(z_2, Q)\|^n \right]^{1/n} \leq c_n(z_1, z_2, Q) \|z_1 - z_2\| e^{t \bar{\mu}(A - P_\infty S)} \quad (5.56)$$

whenever $\mathbf{N}$ is sufficiently large such that $\hat{\mathbf{t}}_n > \mathbf{t}_n$; see [17] for details on these time parameters.

We emphasise again that $\hat{\mathbf{t}}_n \rightarrow_{\mathbf{N} \rightarrow \infty} \infty$. With regards to qualitative reasoning, we may combine Theorem 5.11 and Theorem 5.13 and draw the basic (qualitative) conclusion that, after some initial time period, and given enough particles, the difference in (noisy) error flows ($\hat{\psi}_t(z_1, Q) - \hat{\psi}_t(z_2, Q)$), or the difference in sample means ($\hat{\chi}_t(x_1, Q) - \hat{\chi}_t(x_2, Q)$), is exponentially stable (in some sense) with a rate related to $\bar{\mu}(A - P_\infty S)$.

In the scalar case $d = 1$, stronger stability results on the error flow $\hat{\psi}_t(z, Q)$ follow from the contraction properties in Theorem 5.8 under weaker model and ensemble size assumptions. The strong $\mathbb{L}_n$ -type stability results in the scalar $d = 1$ case are quantitative and hold over infinite horizons for both the **VEnKF** and the **DEnKF**, i.e. with $\kappa \in \{0, 1\}$, with unstable latent signals, with differing initial variance states, and with exponential rates that collapse to the optimal deterministic rates (explicitly computable when $d = 1$) as $\mathbf{N} \rightarrow \infty$. See [21]; and the results, and illustrative examples in the next section.

6 Strong Results in One-Dimensional Illustrative Examples

Throughout this section we let $d = 1$ and $R \wedge S > 0$. The latter condition $R \wedge S > 0$ is both necessary and sufficient for observability and controllability to hold in one dimension; and besides, in some cases, conditions on $\mathbf{N} \geq 1$ no other conditions are needed in this section (and we emphasise that the latent signal may be unstable). The purpose of this section is to illustrate in more detailed quantitative terms some of the more abstract or qualitative results given in the general multi-variate setting. In some cases, the derivation of a multi-variate counterpart of a result in this section remains an open problem. In the scalar setting, the analysis of the **EnKF** is rather complete in the linear-Gaussian case.

When $P_0 \in [0, \infty[$, the deterministic Riccati equation defined on $[0, \infty[$, in the classical Kalman-Bucy filter, satisfies the quadratic differential equation (2.3) which may be written also as,

$$\partial_t P_t = \text{Ricc}(P_t) = -S(P_t - \varrho_+)(P_t - \varrho_-), \quad (6.1)$$

with the equilibrium states (ϱ_-, ϱ_+) defined by

$$S \varrho_- := A - \sqrt{A^2 + RS} < 0 < S \varrho_+ := A + \sqrt{A^2 + RS} \quad (6.2)$$

With $P_0 \in [0, \infty[$, we have $P_t \rightarrow_{t \rightarrow \infty} P_\infty = \varrho_+$. It follows that,

$$A - P_\infty S = -\sqrt{A^2 + RS} \quad (6.3)$$

and thus simplifying, e.g. (2.13), we have the equality,

$$\mathcal{E}_t(Q) = c_t(Q) \mathcal{E}_t(P_\infty) \leq c(Q) \mathcal{E}_t(P_\infty) \quad \text{and} \quad \mathcal{E}_t(P_\infty) = e^{-t\sqrt{A^2 + RS}} \quad (6.4)$$

where $-\sqrt{A^2 + RS}$ may be viewed explicitly as the optimal semigroup contraction rate in the scalar case. The explicit form of the constants $c_t(Q)$, $c(Q)$ is also available in the scalar case, see [21] and also the general Floquet-type multivariate result in [18].

The Riccati drift function $\text{Ricc}(\cdot)$ is also the derivative of the double-well potential function

$$F(Q) = -\frac{S}{3} Q (Q - \zeta_-) (Q - \zeta_+) \quad (6.5)$$

with the roots

$$\zeta_- := \frac{3A}{2S} - \left[\left(\frac{3A}{2S} \right)^2 + \frac{3R}{S} \right]^{1/2} < 0 < \zeta_+ := \frac{3A}{2S} + \left[\left(\frac{3A}{2S} \right)^2 + \frac{3R}{S} \right]^{1/2} \quad (6.6)$$

In this situation, the general Riccati diffusion (5.2) describing the flow of the sample covariance in both case (F1) and case (F2) resumes to the Langevin-Riccati drift-type diffusion process,

$$d\hat{P}_t = \partial F(\hat{P}_t) dt + \frac{1}{\sqrt{\mathbf{N}}} \hat{P}_t^{1/2} \Sigma_\kappa^{1/2}(\hat{P}_t) d\mathcal{M}_t \quad (6.7)$$

with the mapping Σ_κ defined in (5.3). Recall that case (F1) corresponds to the vanilla **EnKF**, denoted by **VEnKF**, and case (F2) corresponds to the ‘deterministic’ **EnKF**, denoted by **DEnKF**. Also observe that $\partial F > 0$ on the open interval $]0, \zeta_+[$ and $\partial F(0) = R > 0 = \sigma(0)$ so that the origin is repellent and instantaneously reflecting.

At any time $t \geq 0$ we may comment on the boundedness of certain moments of the sample variance and the fluctuation of the sample variance and sample mean about their limiting (classical Kalman-Bucy variance and mean) values.

For example, we have the under-bias result $\mathbb{E}[\hat{\phi}_t(Q)] \leq \phi_t$ with any $\mathbf{N} \geq 1$ in both the **VEnKF** of case (F1), and in the **DEnKF** of case (F2). This under-bias motivates so-called variance/covariance

regularisation methods in practice; e.g. so-called sample covariance inflation or localisation methods. Later we discuss the effects of inflation in particular. However, more generally, in the scalar case we have the result of Theorem 5.3 with explicit expressions on the ensemble size, i.e. we have for any $n \geq 1$, $t \geq 0$, $Q \in [0, \infty[$, and any $N \geq 1 \vee 2\kappa(n-1)$, the uniform estimates,

$$\mathbb{E} \left[\widehat{\phi}_t(Q)^n \right]^{1/n} \leq c_n(Q) \quad (6.8)$$

We also have bounds on the inverse Riccati flow (leading to lower bounds on the sample covariance) under stronger conditions on N ; see [21]. We remark here, and again later when we explicitly examine the invariant measure for $\widehat{\phi}_t$, that these conditions on N while explicit, may still be conservative (in the case of the **VENKF**). From Theorem 5.4 and the scalar exposition in [21] we have the uniform fluctuation estimate $\mathbb{E}[\widehat{\phi}_t(Q) - \phi_t(Q)^n]^{1/n} \leq c_n(Q)/\sqrt{N}$ with the explicit $N \geq 1 \vee 2\kappa(n-1)$. The constant $c_n(Q)$ is also studied in [21] with $d = 1$ in explicit detail.

Note that we may expand the stochastic flow of the sample variance as in (5.22). Exploring this idea further in the scalar case for illustrative purposes, we may write the first and second-order fluctuations as,

$$\widehat{\varphi}_t := \sqrt{N}[\widehat{\phi}_t - \phi_t], \quad (6.9)$$

$$\widehat{\varphi}_t^{(2)} := \sqrt{N}[\widehat{\varphi}_t - \varphi_t] \quad (6.10)$$

where in the second line we emphasise the superscript $\cdot^{(2)}$ is an order index (not a power) and where,

$$\varphi_t(Q) := \int_0^t (\partial \phi_{t-s})(\phi_s(Q)) \Sigma_\kappa^{1/2}(\phi_s(Q)) d\mathcal{M}_s \quad (6.11)$$

and the derivatives $\partial^k \phi_t$ of any order are explicitly given in [21]. In this case, $\partial \phi_t(Q) = \mathcal{E}_t^2(Q)$ (where the superscript here is now a power). We then have,

$$\widehat{\phi}_t = \phi_t + \frac{1}{\sqrt{N}} \varphi_t + \frac{1}{N} \widehat{\varphi}_t^{(2)} \quad (6.12)$$

The natural central limit theorem follows, i.e. $\sqrt{N}[\widehat{\phi}_t - \phi_t] \xrightarrow{N \rightarrow \infty} \varphi_t$. The (non-)asymptotic variance is estimated in [21, 22].

The expansion (6.12) allows ones to better understand the bias properties of the sample covariance $\widehat{\phi}_t$. Writing the third-order fluctuation as,

$$\widehat{\varphi}_t^{(3)} := \sqrt{N}[\widehat{\varphi}_t^{(2)} - \varphi_t^{(2)}/2] \quad (6.13)$$

and expanding and taking expectations,

$$\mathbb{E}[\widehat{\phi}_t] = \phi_t + \frac{1}{2N} \mathbb{E}[\varphi_t^{(2)}] + \frac{1}{N^{3/2}} \mathbb{E}[\widehat{\varphi}_t^{(3)}] \quad (6.14)$$

and limits we have the dominating (N -order-asymptotic) bias is given by

$$N \left(\mathbb{E}[\widehat{\phi}_t] - \phi_t \right) \xrightarrow{N \rightarrow \infty} \frac{1}{2} \int_0^t (\partial^2 \phi_{t-s})(\phi_s(Q)) \Sigma_\kappa(\phi_s(Q)) ds < 0 \quad (6.15)$$

which is always negative (agreeing with the under-bias result $\mathbb{E}[\widehat{\phi}_t(Q)] \leq \phi_t$). See [21, 22] for further exploration of these general expansions. A detailed study of these expansions may aid in the development and tuning of (adaptive) sample covariance regularisation methods.

Significantly generalising Theorem 5.10 in the scalar case [21], we have for any $t \geq 0$, and any $N > 2(4n+1)(1+4\kappa)$, the uniform bound $\mathbb{E}[\widehat{\psi}_t(z, Q)^n]^{1/n} \leq c_n(z, Q)$. We also have the

generalisation that for any $t \geq 0$, and any $N > 2(6n + 1)(1 + 4\kappa)$, the uniform fluctuation estimate,

$$\mathbb{E} \left[\widehat{\psi}_t(Q) - \psi_t(Q)^n \right]^{1/n} \leq \frac{c_n(z, Q)}{\sqrt{N}} \quad (6.16)$$

holds.

The expansion (5.48) of the sample mean (error) may be explored similarly to the above expansion of the sample covariance. The first order terms ϑ_t in (5.48) related to the central limit theorem are studied in [22, Section 1.3].

The infinitesimal generator of the diffusion (6.7) on $]0, \infty[$ is given in Sturm-Liouville form by the equation

$$L(f) = \frac{2}{N} \iota \Sigma_\kappa e^V \partial (e^{-V} \partial f) \quad \text{with} \quad V(\cdot) = -\frac{N}{2} \int_\delta^\cdot \partial F(x) \iota^{-1}(x) \Sigma_\kappa^{-1}(x) dx \quad (6.17)$$

for any $\delta > 0$ and where we recall the identity function $\iota(x) := x$. This implies that a reversible measure of the Riccati diffusion (5.2) in the scalar $d = 1$ case is given by the formula

$$\widehat{\Gamma}_\infty(dx) \propto 1_{]0, \infty[}(x) \frac{N}{4 \iota(x) \Sigma_\kappa(x)} \exp(-V(x)) dx \quad (6.18)$$

In case (F1) corresponding to the **VEnKF**, we have that L is reversible w.r.t. the probability measure $\widehat{\Gamma}_\infty$ on $]0, \infty[$ defined by,

$$\widehat{\Gamma}_\infty(dx) \propto 1_{]0, \infty[}(x) \exp \left(N \frac{A}{\sqrt{RS}} \tan^{-1} \left(x \sqrt{\frac{S}{R}} \right) \right) \left(\frac{x}{R + Sx^2} \right)^{\frac{N}{2}} \frac{1}{x(R + Sx^2)} dx \quad (6.19)$$

See also [21] for alternate derivations/forms of this heavy tailed invariant measure. The heavy tailed nature of the stationary measure implies that for the n -th moment to exist one requires $N > 0 \vee 2(n - 2)$. As expected this condition on N is generally weaker than that required for n -th moment boundedness at any time $t \geq 0$ in (6.8) in terms of the **VEnKF**. In Figure 2 we plot the line defined by $(2n - 4)/N$ for various N values. With any $N \geq 1$, we have existence of the first two moments.




Figure 2: Existence of moments for the **VEnKF**. Each line corresponds to some number N with N moving from 1 to 50 fanning left to right. The ‘x-axis’ corresponds to moment orders n and a moment n exists whenever the line $(2n - 4)/N$ is strictly less than one.

Higher-order moments even in one dimension are still troublesome (for the **VEnKF**, $\kappa = 1$). In fact, the diffusion $\widehat{P}_t$ for the sample variance in case (F1) does not have any exponential

moments in the stationary regime for any finite $N \geq 1$. That is, for any $t \geq 0$ and any finite $\alpha > 0$ we have

$$\text{Law}(Q) = \hat{\Gamma}_\infty \implies \mathbb{E} \left[\exp \left(\alpha \|\hat{\phi}_t(Q)\| \right) \right] = \infty \quad (6.20)$$

for any $N \geq 1$.

We also remark that the heavy tailed nature of this stationary distribution, in the case of the **VEnKF**, implies that numerical stability in practice may be worrisome. In the stationary regime, it is realistic to expect samples from the tails in this case, and these may be large enough and/or frequent enough to cause numerical divergence. This property may lead to so-called catastrophic divergence as studied in, e.g., [83]. In [67, 65, 83] mechanisms for catastrophic divergence are studied in complex nonlinear systems. Here we argue that even in linear systems, the heavy-tailed nature of the invariant measure of the sample covariance may lead to samples numerically large enough to cause numerical catastrophe in any practical computing system.

In case (F2) corresponding to the **DEnKF**, we have that L is reversible w.r.t. the probability measure $\hat{\Gamma}_\infty$ on $]0, \infty[$ defined by,

$$\hat{\Gamma}_\infty(dx) \propto 1_{]0, \infty[}(x) x^{\frac{N}{2}-1} \exp \left(-\frac{SN}{4R} \left(x - 2 \frac{A}{S} \right)^2 \right) dx \quad (6.21)$$

Note this measure has Gaussian tails, and we contrast this with the heavy tailed nature of (6.19). This is significant, since it implies that the sample variance (and mean) of this **DEnKF** will exhibit smaller fluctuations than the **VEnKF**, and that all moments (including exponential moments) exist in this case for any choice of $N \geq 1$. This latter result is consistent with Theorem 5.3 at any time $t \geq 0$ (and in the general multivariate setting). We can also expect better numerical stability (e.g. less outliers); including better time-discretisation properties [74] in case (F2). These better fluctuation properties are already apparent in the preceding results (e.g. see Theorem 5.3, 5.4 and 5.13) in the full multi-dimensional setting.

As an illustrative example, take $A = 20$ (i.e. the underlying signal model is highly unstable), $R = S = 1$ and $N = 6$. In Figure 3 we compare the invariant measure for the flow of the sample variance in each case.




Figure 3: The invariant measure of the sample variance of the ‘vanilla’ **EnKF** in case (F1), versus that of the ‘deterministic’ **EnKF** in case (F2).

We see in Figure 3 the heavy tails of the invariant measure (6.19) for the vanilla **EnKF** sample variance, and conversely the Gaussian-type tails in the case (6.21) of the ‘deterministic’ **EnKF**. Note also the positioning of the mode/mean in each case. In case (F1) of the **VEnKF**, n -th order moments exist only when $(2n - 4)/N$ is strictly less than one (in this case for $n < 5$); while all moments exist in case (F2) for the **DEnKF**.

The benefit and real interest in the scalar case is the ability to explicate the convergence rates, e.g. as in (6.4). We finally turn to the convergence/stability properties of the **EnKF** sample variance and sample mean. In the case of the sample variance, we know from Theorem 5.5 that convergence of $\hat{\phi}_t$ to its invariant measure $\hat{\Gamma}_\infty$ (e.g. as depicted in Figure 3 and described by (6.19) or (6.21)) holds if $N > 4 + (\kappa S)/(2R)$. Proof of this condition on N follows from Theorem 5.5, the original multivariate statement of the same result in [19, Theorem 2.4] and bounds on the mean of the sample variance flow and its inverse [21, 19]. In [21] we also consider contraction and stability properties of the distribution of the sample covariance with respect to a particular Wasserstein metric; as opposed to the Λ -norm contraction used in Theorem 5.5. An interesting result from [21] is that when $\kappa = 0$, and for stable signal models (i.e. $A < 0$), the Riccati diffusion (6.7) (describing the flow of the sample covariance) may converge faster to its invariant measure in (6.21), than the deterministic Riccati (6.1) does to its fixed point in (6.2).

In one-dimensional ($d = 1$) settings we may say more on the (stochastic) stability of the **EnKF** sample covariance ϕ_t and sample mean ψ_t based on the contraction properties of the stochastic transition matrix $\hat{\mathcal{E}}_t(Q)$ defined in (5.8). It follows from Theorem 5.8 that we have the exponential decay estimate with $N > 4 \vee (4n - 2)\kappa$ which comes from [19, Theorem 2.7],

$$\mathbb{E} \left[\hat{\mathcal{E}}_t(Q)^n \right]^{1/n} = \mathbb{E} \left[\exp \left(n \int_0^t (A - \hat{\phi}_s(Q)S) ds \right) \right]^{1/n} \leq c_n(Q) \exp \left(-t \sqrt{\hat{R}_n \hat{S}_n} \right) \quad (6.22)$$

where $\hat{R}_n$ and $\hat{S}_n$ follow from (5.42). In addition, there exists some function (of N) $\lim_{N \rightarrow \infty} \hat{\nu}_n = 0$ such that

$$\mathbb{E} \left[\hat{\mathcal{E}}_t(Q)^n \right]^{1/n} = c_n(Q) \exp \left(-t (1 - \hat{\nu}_n) \sqrt{A^2 + RS} \right) \quad (6.23)$$

which we may relate (or contrast) with the exact contraction rate of the exponential semigroup associated with the deterministic Riccati equation in (6.4) describing the true filter variance in the classical Kalman-Bucy filter. The rate parameter $\hat{\nu}_n$ is different between the **VEnKF** and **DEnKF**. Details on the parameter $\hat{\nu}_n$ are given in [21] but importantly for both $\kappa \in \{0, 1\}$ we recover naturally the convergence rate of the deterministic Riccati flow in (6.4).

The exponential decay of the exponential semigroup $\hat{\mathcal{E}}_t(Q)$ plays a central role in the stability of the pair of processes $(\hat{\phi}_t, \hat{\psi}_t)$. For large time horizons the Lyapunov exponent can be estimated by the formula

$$\frac{1}{t} \log \hat{\mathcal{E}}_t(Q) = \frac{1}{t} \int_0^t (A - \hat{\phi}_s(Q)S) ds \xrightarrow{t \rightarrow \infty} A - \hat{\Gamma}_\infty(\iota)S, \quad (6.24)$$

where $\hat{\Gamma}_\infty$ denotes the reversible measure (6.19) or (6.21). We also have the following estimates of the Lyapunov exponent (6.24) from [21], and that relate also to the under-bias $\mathbb{E}[\hat{\phi}_t(Q)] \leq \phi_t$. Let $\kappa = 0$ and let $\text{Law}(Q) = \hat{\Gamma}_\infty$ be the reversible probability measure defined in (6.21). Then, for any $t \geq 0$, we have

$$N > 4 \implies -\sqrt{A^2 + RS} \leq A - \mathbb{E}[\hat{\phi}_t(Q)]S \leq -\sqrt{A^2 + RS \left(1 - \frac{4}{N}\right)} < 0 \quad (6.25)$$

Similarly assuming $\kappa = 1$ with $\text{Law}(Q) = \hat{\Gamma}_\infty$ and $\hat{\Gamma}_\infty$ defined in (6.19) we have for any $t \geq 0$,

$$N > 4 \implies -\sqrt{A^2 + RS} \leq A - \mathbb{E}[\hat{\phi}_t(Q)]S \leq -\frac{\sqrt{A^2 + RS(1 - (4/N)^2)} - 4A/N}{1 + 4/N} < 0 \quad (6.26)$$

As noted, the left hand inequalities in the preceding two equations follows immediately from the under-bias result $\mathbb{E}[\hat{\phi}_t(Q)] \leq \phi_t$

From the contraction properties on $\mathbb{E}[\hat{\mathcal{E}}_t(Q)^n]$ we may deduce, in the scalar setting, strong stability results on the stochastic Riccati flow $\hat{\phi}_t$ analogous to the deterministic setting, e.g.

(2.20). Similarly, strong stability results on the error flow $\hat{\psi}_t$ follow from the contraction properties of $\mathbb{E}[\hat{\mathcal{E}}_t(Q)^n]$. Importantly, in the scalar $d = 1$ case of $\hat{\psi}_t(z, Q)$ we may relax the multivariate results like Theorem 5.11 and Theorem 5.13 which require more restrictive model (e.g. the strong observability/stability $\mu(A - P_\infty S) < 0$ condition) and ensemble (particle) size assumptions.

From [21, Theorem 5.10] we have that for any $N > 4 \vee 4\kappa(n - 1)$,

$$\mathbb{E} \left[\|\hat{\phi}_t(Q_1) - \hat{\phi}_t(Q_2)\|^n \right]^{1/n} \leq c_n \|Q_1 - Q_2\| \exp \left(-t(1 - \hat{\nu}_n) \sqrt{A^2 + RS} \right) \quad (6.27)$$

for some function (of N) $\lim_{N \rightarrow \infty} \hat{\nu}_n = 0$. Note we have found no analogue of this result in the multivariate setting.

From [21, Theorem 6.1] we have that for any $N > 2(4n + 1)(1 + 4\kappa)$,

$$\mathbb{E} \left[\|\hat{\psi}_t(z_1, Q_1) - \hat{\psi}_t(z_2, Q_2)\|^n \right]^{1/n} \leq c_n(z_1, z_2, Q_1, Q_2) (\|z_1 - z_2\| + \|Q_1 - Q_2\|) \exp \left(-t(1 - \hat{\nu}_n) \sqrt{A^2 + RS} \right) \quad (6.28)$$

for some function (of N) $\lim_{N \rightarrow \infty} \hat{\nu}_n = 0$. We may contrast this result with the more restrictive Theorem 5.13 in the multivariate setting. Note in the scalar setting we accommodate both the **VEnKF** and **DEnKF**, different initial variance states, and we recover, over fully infinite horizons, a continuous relationship with the optimal stability rates of (6.4).

The constants in (6.23), (6.27), and (6.28) are given explicitly in terms of the model parameters in [21]. We remark that across these three stability results, the details of $\hat{\nu}_n$ vary [21], but importantly we recover the optimal (classical Kalman-Bucy) rates $\lim_{N \rightarrow \infty} \hat{\nu}_n = 0$.

We consider an illustration of the fluctuation and stability properties of the sample variance in the different **EnKF** variants. Consider again the model leading to Figure 3, and let $\hat{P}_0 = 0$. The deterministic Riccati flow ($N = \infty$, in (6.1)) of the classical Kalman-Bucy filter and with the chosen model parameters ($A = 20$, $R = S = 1$) is given in Figure 4, along with 100 sample paths of the sample variances for both the **VEnKF** and the **DEnKF** (with $N = 6$).




Figure 4: Flow of the deterministic Riccati equation, and 100 sample paths of the **VEnKF** sample variance of case (F1), and 100 sample paths of the **DEnKF** sample variance of case (F2).

Note in Figure 4 the drastically reduced fluctuations in the ‘deterministic’ **EnKF** sample variance sample paths. At equilibrium, these fluctuations are related to the invariant measures of the two **EnKF** varieties in (6.19) and (6.21).

In Figure 5 we plot the flow of the first two central moments and the 3rd through the 9th standardised central moments for both the **VEnKF** and **DEnKF** sample variance distribution. Recall that $N = 6$ in this example, and we expect moments of the **VEnKF** sample variance in case (F1) to exist up to $n = 4$ with $n = 5$ the boundary case; while all moments exist for the **DEnKF** of case (F2).




Figure 5: Flow of the sample variance moments for the **VEnKF** and the **DEnKF**.

We note in Figure 5 that the sample variance moments for the **VEnKF** in case (F1) begin to destabilize around the 5th/6th moments as expected. Importantly, the mean of the sample variance for the **VEnKF** is very negatively biased in this case, while the mean of the **DEnKF** in case (F2) is quite accurate. Note also the very large variance in the sample variance for the **VEnKF**.

Lastly, observe that (6.7) has non-globally Lipschitz coefficients in case (F1). The drift is quadratic, while the diffusion has a polynomial growth of order $3/2$. It follows by [74] that the naive Euler-type time-discretization may blow up, regardless of the boundedness properties of the limiting (continuous-time) diffusion.

7 Regularisations and Comparisons

7.1 Covariance Regularisation: Inflation

Let $(\mathcal{V}_t^i, \mathcal{W}_t^i, \mathcal{X}_0^i)$ with $1 \leq i \leq N+1$ be $(N+1)$ independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. Now consider a modification of the individual particle update equations in the two cases of interest,

$$(F1) \quad d\mathcal{X}_t^{i,\varepsilon} = A \mathcal{X}_t^{i,\varepsilon} dt + R^{1/2} d\mathcal{V}_t^i + \left(\hat{P}_t^\varepsilon + \varepsilon T \right) H' R_1^{-1} \left[d\mathcal{Y}_t - \left(H \mathcal{X}_t^{i,\varepsilon} dt + R_1^{1/2} d\mathcal{W}_t^i \right) \right] \quad (7.1)$$

$$(F2) \quad d\mathcal{X}_t^{i,\varepsilon} = A \mathcal{X}_t^{i,\varepsilon} dt + R^{1/2} d\mathcal{V}_t^i + \left(\hat{P}_t^\varepsilon + \varepsilon T \right) H' R_1^{-1} \left[d\mathcal{Y}_t - H \left(\frac{\mathcal{X}_t^{i,\varepsilon} + \hat{X}_t^\varepsilon}{2} \right) dt \right]$$

where $\varepsilon \in [0, \infty[$, and $T \in \mathbb{S}_r^0$ is some given reference matrix. Here, $\hat{P}_t^\varepsilon = \hat{\phi}_t^\varepsilon(Q)$ denotes the sample covariance-type function of $\mathcal{X}_t^{i,\varepsilon}$ given by,

$$\begin{aligned} \hat{\eta}_t^\varepsilon &:= \eta_t^{N,\varepsilon} = \frac{1}{N+1} \sum_{i=1}^{N+1} \delta_{\mathcal{X}_t^{i,\varepsilon}} \\ \implies \hat{X}_t^\varepsilon &:= X_t^{N,\varepsilon} = \frac{1}{N+1} \sum_{i=1}^{N+1} \mathcal{X}_t^{i,\varepsilon} \quad \text{and} \quad \hat{P}_t^\varepsilon := P_t^{N,\varepsilon} = \frac{N+1}{N} \mathcal{P}_{\hat{\eta}_t^\varepsilon} \end{aligned} \quad (7.2)$$

Recall the unified representation (for both the **VEnKF** and the **DEnKF**) for the flow of the sample mean, sample covariance, and the sample error flow in equations (5.1) through to (5.4).

Now consider the modification of the state estimator (sample mean) update equation resulting from the (εT) -modification to the particle updates,

$$d\hat{X}_t^\varepsilon = (A - (\hat{P}_t^\varepsilon + \varepsilon T)S) \hat{X}_t^\varepsilon dt + (\hat{P}_t^\varepsilon + \varepsilon T) H' R_1^{-1} d\mathcal{Y}_t + \frac{1}{\sqrt{N+1}} \Sigma_{\kappa,\varepsilon}^{1/2}(\hat{P}_t^\varepsilon) d\mathcal{B}_t \quad (7.3)$$

with the mapping,

$$\Sigma_{\kappa,\varepsilon}(Q) := R + \kappa(Q + \varepsilon T)S(Q + \varepsilon T) \quad \text{with} \quad \kappa = \begin{cases} 1 & \text{in case (F1)} \\ 0 & \text{in case (F2)} \end{cases} \quad (7.4)$$

With $\hat{Z}_t^\varepsilon := (\hat{X}_t^\varepsilon - \mathcal{X}_t)$ we then also have that,

$$\begin{aligned} d\hat{Z}_t^\varepsilon &= (A - (\hat{P}_t^\varepsilon + \varepsilon T)S) \hat{Z}_t^\varepsilon dt + (\hat{P}_t^\varepsilon + \varepsilon T) H' R_1^{-1} d\mathcal{W}_t \\ &\quad - R^{1/2} d\mathcal{V}_t + \frac{1}{\sqrt{N+1}} \Sigma_{\kappa,\varepsilon}^{1/2}(\hat{P}_t^\varepsilon) d\mathcal{B}_t \\ &\stackrel{\text{law}}{=} (A - (\hat{P}_t^\varepsilon + \varepsilon T)S) \hat{Z}_t^\varepsilon dt + \Omega_{\kappa,\varepsilon}^{1/2}(\hat{P}_t^\varepsilon) d\hat{\mathcal{B}}_t \end{aligned} \quad (7.5)$$

with,

$$\Omega_{\kappa,\varepsilon} := \Sigma_{1,\varepsilon} + \frac{1}{N+1} \Sigma_{\kappa,\varepsilon} \quad (7.6)$$

In un-regularised ensemble Kalman filtering, we approximate ϕ_t by the sample covariance $\hat{\phi}_t$ since the dimension of $\hat{X}_t$ may be in the many millions; see [60]. However, when computing the sample covariance in high-dimensions, rank deficient estimation is common due to a lack of enough samples. Covariance inflation, leading to an approximation of the form $(\hat{\phi}_t^\varepsilon(Q) + \varepsilon T)$, in the update equation (e.g. (7.1)), is a common, simple means of addressing this rank deficiency [60]. The under-bias result in Theorem 5.2 or equation (5.16) can also motivate the use of some form of regularisation such as inflation.

Note that the perturbation in the resulting flow of $\hat{\phi}_t^\varepsilon(Q)$ comes from a (rather delicate) feedback loop adding εT to the covariance of the signal $\hat{X}_t^\varepsilon$ at each instant. The flow of $\hat{\phi}_t^\varepsilon(Q)$ is given by,

$$\begin{aligned} d\hat{P}_t^\varepsilon &= \left[\left(A - \frac{(1-\kappa)}{2} \varepsilon TS - \hat{P}_t^\varepsilon S \right) \hat{P}_t^\varepsilon + \hat{P}_t^\varepsilon \left(A - \frac{(1-\kappa)}{2} \varepsilon TS - \hat{P}_t^\varepsilon S \right)' + R + \kappa \varepsilon^2 T S T \right] dt \\ &\quad + \frac{2}{\sqrt{N}} \left[\hat{P}_t^{\varepsilon 1/2} d\mathcal{M}_t \Sigma_{\kappa,\varepsilon}^{1/2}(\hat{P}_t^\varepsilon) \right]_{\text{sym}} \end{aligned} \quad (7.7)$$

In the limit $N \rightarrow \infty$ we recover a perturbed, deterministic, Riccati equation that describes the flow of the limiting covariance. This perturbed Riccati equation is studied in [24, 20]. For any size $\|\varepsilon T\| < \infty$, the perturbed Riccati flow qualitatively retains all the stability properties of the nominal Riccati flow (e.g. (2.20), but with a different steady state value), and the size of the error between the two grows in a well-quantified continuous way.

In the limiting case, we have via [20, Theorem 2.1] that $\phi_t(Q) \leq \phi_t^\varepsilon(Q)$ in case (F1). In case (F2) we have that $\phi_t^\varepsilon(Q) \leq \phi_t(Q)$ in the limit $N \rightarrow \infty$.

For any $s \leq t$ and $Q \in \mathbb{S}_d^0$ we define the stochastic state-transition matrix,

$$\hat{\mathcal{E}}_{s,t}^\varepsilon(Q) := \exp \left[\oint_s^t \left(A - \varepsilon TS - \hat{\phi}_u^\varepsilon(Q) S \right) du \right] \Leftrightarrow \partial_t \hat{\mathcal{E}}_{s,t}^\varepsilon(Q) = \left(A - \varepsilon TS - \hat{\phi}_t^\varepsilon(Q) S \right) \hat{\mathcal{E}}_{s,t}^\varepsilon(Q) \quad (7.8)$$

Note that this semigroup $\hat{\mathcal{E}}_{s,t}^\varepsilon(Q)$ is associated with the evolution of the (inflation) regularised sample mean in (7.3) or the error flow (7.5) in both case (F1) and (F2). Unlike the un-regularised setting, this same semigroup is not directly related to the evolution of the sample covariance, in (7.7); for example, in case (F1) the semigroup associated with the evolution of the sample covariance is just $\hat{\mathcal{E}}_{s,t}(Q)$ as given in (5.8) and studied throughout the preceding section.

We can comment on the effect of inflation regularisation on the contraction properties of $\hat{\mathcal{E}}_{s,t}^\varepsilon(Q)$, as compared e.g. to $\hat{\mathcal{E}}_{s,t}(Q)$. Firstly, it is worth noting, given the contraction estimates in Section 5.2, that,

$$\mu((A - \varepsilon TS) - PS) \leq \mu(A - PS) \quad (7.9)$$

for any fixed matrix $P \in \mathbb{S}_d^0$ and $S \in \mathbb{S}_d^0$. Arguing as in (5.30), when $S \in \mathbb{S}_d^+$, then up to a change of basis we can always assume that $S = I$. We then have,

$$\mu(A) < \|\varepsilon T\| \implies \|\hat{\mathcal{E}}_{s,t}^\varepsilon(Q)\| \leq \exp(\mu(A - \varepsilon T)(t - s)) \xrightarrow{(t-s) \rightarrow \infty} 0 \quad (7.10)$$

which illustrates the added stabilising effects of εT in the extreme case in which $\hat{P}_t^\varepsilon S$ has no stabilising effect at all. Contrast this with (5.30). Then one interpretation of the preceding relationship is that εT extends the set of signal matrices $A \in \mathbb{M}_d$ for which one may immediately achieve stabilisation (regardless of the effect of $\hat{P}_t^\varepsilon S$). In practice, $\hat{P}_t^\varepsilon S$ will also act to stabilise the filter, see e.g. (5.56). Indeed, in the classical Kalman filtering setting (2.2), (2.3) with $\varepsilon = 0$, the time-varying matrix $(A - P_t S)$ is stabilising [16] for any $A \in \mathbb{M}_d$, even A unstable. In the **EnKF**, we know that $\hat{P}_t$ will fluctuate about P_t , e.g. see Theorem 5.4. Therefore, the stabilisation properties of $(A - \hat{P}_t S)$ are unclear; indeed the study of $\hat{\mathcal{E}}_{s,t}(Q)$ in the preceding Section 5.2 is concerned with precisely this issue. The above implies that the addition of εT can act to counter the negative effects of this fluctuation (and directly add a stabilising effect on the state estimation error).

Finally, we have $\phi_t(Q) \leq \phi_t^\varepsilon(Q)$ in case (F1) and $\phi_t^\varepsilon(Q) \leq \phi_t(Q)$ in case (F2). The semigroup associated with the error flow (7.5) in both cases is the same. The inequality $\phi_t(Q) \leq \phi_t^\varepsilon(Q)$ in case (F1) suggests that the diffusion fluctuation in (7.3) or (7.7) will increase. However, we conversely expect that with $S \in \mathbb{S}_d^+$ we have $\mu((A - \varepsilon TS) - \phi_t^\varepsilon(Q)S) \leq \mu((A - \varepsilon TS) - \phi_t(Q)S)$ and thus we gain a type of stabilising effect. Inflation in case (F1) is then a delicate balancing tradeoff between adding noise to the diffusion coefficients (which may kill the existence of sample covariance moments, for example), and adding a stabilising effect on the sample mean error flow. When $\varepsilon > 0$ is large enough we can achieve added stabilisation in case (F2), as compared to the non inflated case. This is not automatic as in case (F1) because $\phi_t^\varepsilon(Q) \leq \phi_t(Q)$. However, the fluctuations are (further) decreased with inflation in case (F2).

7.2 Discretisation Matters

The present article is primarily concerned with continuous-time filtering and **EnKF** models. In practice, the stochastic models and analysis used for the continuous-time **EnKF** are not applicable without an additional level of time-discretisation approximation. As alluded to earlier, the flow of the sample covariance for the **VEnKF** has quadratic drift coefficients, while the diffusion term has a polynomial growth of order 3/2. In this particular case, it follows by [74] that a basic Euler time-discretization may blow up, regardless of the boundedness properties of the diffusion.

In contrast with continuous-time models, discrete-time signal and observation models lead to the so-called discrete-time **EnKF**; e.g. see [34] and [126] respectively for the corresponding **VEnKF** and **DEnKF** methods (also referenced in discrete-time earlier). Convergence of the discrete-time **EnKF** models to their continuous-time counterparts (studied herein) with appropriate model time-step parameterisations is studied in [95, 96, 93].

The purely discrete-time **EnKF** is not defined by a single coupled diffusion process, but rather by a coupled two-step prediction-updating process (also known as forecast-analysis steps in the **EnKF** and data assimilation literature). Moreover, the Gaussian-nature of the diffusion models (e.g. the Riccati diffusion) arising in the analysis of continuous-time **EnKF** theory is also lost, and an inherent difficulty in discrete-time is the introduction of more sophisticated non-central chi-squared fluctuations.

We emphasize that the discrete-time **DEnKF** of Sakov and Oke [126] is not consistent; i.e. it does not converge to the optimal filter as the number of particles tends to infinity, even in the

linear-Gaussian case. The **VEnKF** is consistent in discrete-time, see [34]. In discrete-time, another class of **EnKF** methods, termed square-root **EnKF** methods [137, 96], are consistent [91, 94]. The discrete-time square-root and deterministic **EnKF** methods are closely related (see [126]) and in the continuous-time limit they converge to the same object [96, 93], i.e. the continuous-time **DEnKF** studied here.

Analysis of the discrete-time **VEnKF** [34] was studied in the linear-Gaussian setting in [51]. That article presents a rather complete analysis of the fluctuations and the long-time behaviour of **VEnKF** for one-dimensional models, including uniform estimates in the presence of transient and unstable latent signals.

7.3 Particle Filter Comparisons

We end this section with some theoretical comparisons between the discrete-time **VEnKF** [34, 51] and the particle filter (abbreviated **PF** in this section) [64, 45, 53, 58, 84, 85].

We remark that the **PF** and the **EnKF** are, in general, built on different prediction/correction mechanisms. The **EnKF** uses an empirical gain function to weight the observations with the predicted state estimate in a manner akin to the update/correction stage of the classical Kalman filter. This mechanism which involves “moving” the corrected state estimate has the potential to stabilise the state estimate about the, possibly unstable, latent signal. This idea has played a central focus in this article. As shown in Theorem 5.2, the empirical gain depends on an under-biased sampled covariance matrix which may fail to correct the effective unstable dimensions. The **EnKF** literature abounds with inconsistent but judicious ensemble transformations and regularisation methods like inflation/localisation-type procedures, aimed at addressing this issue. One basic inflation methodology is discussed in Section 7.1, see also [24], in the context of continuous-time **EnKF** methods where its action on the stabilisation properties of the **EnKF** are shown. Conversely, both the continuous and the discrete-time **PF** are based on genetic-type mutation-selection transitions: The basic discrete-time **PF** methods, e.g. see [64, 45, 84], evolve as a Markov chain on a product space. During the mutation transition, particles evolve independently according to the signal transition model. When an observation is delivered by the sensor, particles are selected with a probability proportional to their likelihoods. Importance sampling tricks can also be used to avoid degenerate mutations/predictions.

Next we provide a detailed comparison of the **PF** and **VEnKF** for one-dimensional linear-Gaussian models with a view toward their tracking capability in the event of an unstable latent signal. The one-dimensional discrete-time version of (2.1) has the following form,

$$\begin{aligned}\mathcal{X}_{t+1} &= A \mathcal{X}_t + R^{1/2} \mathcal{V}_{t+1} \\ \mathcal{Y}_t &= H \mathcal{X}_t + R_1^{1/2} \mathcal{W}_t\end{aligned}\tag{7.11}$$

where $t \in \mathbb{N}$ denotes the discrete time index and $(\mathcal{W}_t, \mathcal{V}_{t+1})$ is a sequence of 2-dimensional Gaussian random variables with zero-mean and unit variance. The initial condition of the signal $\mathcal{X}_0$ is Gaussian with mean and variance denoted by (X_0^-, P_0^-) (independent of $(\mathcal{W}_t, \mathcal{V}_{t+1})$), and (A, H, R, R_1) are the model parameters. Any non-zero values for (A, H, R, R_1) ensure that the model is (discrete-time) detectable and stabilisable.

The discrete-time version of the conditional nonlinear McKean-Vlasov-type diffusion process (F1) discussed in (3.2) is given by,

$$\begin{cases} \mathcal{X}_t = \mathcal{X}_t^- + \mathcal{G}_{\bar{\eta}_t^-} (\mathcal{Y}_t - (H \mathcal{X}_t^- + R_1^{1/2} \mathcal{W}_t)) & \text{with } \mathcal{G}_{\bar{\eta}_t^-} := H \mathcal{P}_{\bar{\eta}_t^-} / (H^2 \mathcal{P}_{\bar{\eta}_t^-} + R_1) \\ \mathcal{X}_{t+1}^- = A \mathcal{X}_t + R^{1/2} \mathcal{V}_{t+1} \end{cases}\tag{7.12}$$

In the above display, $\bar{\eta}_t^-$ denotes the conditional distribution of $\mathcal{X}_t^-$ given $\mathcal{Y}_t^- := (\mathcal{Y}_0, \dots, \mathcal{Y}_{t-1})$ and $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$ are independent copies of $(\mathcal{V}_t, \mathcal{W}_t, \mathcal{X}_0)$. Using a simple induction argument, it

is straightforward to show that,

$$\bar{\eta}_t^- = \text{Law}(\mathcal{X}_t^- | \mathcal{Y}_t^-) = \text{Law}(\mathcal{X}_t | \mathcal{Y}_t^-) \quad (7.13)$$

is Gaussian with mean X_t^- and variance P_t^- ; and,

$$\bar{\eta}_t = \text{Law}(\mathcal{X}_t | \mathcal{Y}_t) = \text{Law}(\mathcal{X}_t | \mathcal{Y}_t) \quad (7.14)$$

is also Gaussian with mean X_t and variance P_t . Moreover, the conditional means (X_t^-, X_t) and the variances (P_t^-, P_t) obey the Kalman filtering (update/correction and prediction) equations,

$$\begin{cases} X_t &= X_t^- + G_t (\mathcal{Y}_t - H X_t^-) \\ P_t &= (1 - G_t H) P_t^- \end{cases} \quad \text{and} \quad \begin{cases} X_{t+1}^- &= A X_t \\ P_{t+1}^- &= A^2 P_t + R, \end{cases} \quad (7.15)$$

where in this section G_t denotes the so-called Kalman gain parameter,

$$G_t := H P_t^- / (H^2 P_t^- + R_1) \implies 1 - G_t H = 1 / (1 + S P_t^-) \quad \text{with} \quad S := H^2 / R_1 \quad (7.16)$$

The particle approximation of the nonlinear Markov chain discussed above is given by an interacting particle system defined sequentially for any $1 \leq i \leq N + 1$ by the formulae,

$$\begin{cases} \mathcal{X}_t^i &= \mathcal{X}_t^{i-} + \hat{G}_t (\mathcal{Y}_t - (H \mathcal{X}_t^{i-} + R_1^{1/2} \mathcal{W}_t^i)) \quad \text{with} \quad \hat{G}_t := H \hat{P}_t^- / (H^2 \hat{P}_t^- + R_1) \\ \mathcal{X}_{t+1}^{i-} &= A \mathcal{X}_t^i + R^{1/2} \mathcal{V}_{t+1}^i \end{cases} \quad (7.17)$$

where $\hat{P}_t^-$ denotes for the normalised sample variance

$$\hat{P}_t^- := \frac{1}{N} \sum_{1 \leq i \leq N+1} (\mathcal{X}_t^{i-} - \hat{X}_t^-)^2 \quad \text{with} \quad \hat{X}_t^- := \frac{1}{N+1} \sum_{1 \leq i \leq N+1} \mathcal{X}_t^{i-} \quad (7.18)$$

The above model coincides with the discrete-time version of the **VEnKF** (i.e. in case **F1**) in (4.1), and follows from early results in [34]. The conditional mean $X_t = \mathbb{E}(\mathcal{X}_t | \mathcal{Y}_t)$ is approximated by the sample mean,

$$\hat{X}_t^{\text{EnKF}} := \frac{1}{N+1} \sum_{1 \leq i \leq N+1} \mathcal{X}_t^i \quad (7.19)$$

The next theorem provides uniform mean-error estimates.

Theorem 7.1. *For any non-zero values of the model parameters (A, H, R, R_1) , any $n \geq 1$, and for $N \geq 1$ sufficiently large, we have the uniform estimates*

$$\sup_{t \geq 0} \mathbb{E} \left[|\hat{X}_t^{\text{EnKF}} - X_t|^n \right]^{1/n} \leq c_n \frac{1}{\sqrt{N}} \quad (7.20)$$

We also have the conditional bias estimate,

$$\mathbb{E} \left[\left| \mathbb{E} \left[\hat{X}_t^{\text{EnKF}} | \mathcal{Y}_t \right] - X_t \right|^n \right]^{1/n} \leq c_n (P_0^-)^{1/N} \quad (7.21)$$

The proof of this discrete-time, one-dimensional, version of (6.16) is given in [51]; see also [21] for continuous-time analogues. Contrast this result also with the continuous-time multi-variate result in Theorem 5.9.

Particle filters, eg. see [64, 45, 49, 46, 53, 58, 57, 84, 85], are a popular method for approximate filtering in nonlinear state space models in relatively low dimensions. The genetic-type particle filter (also referred to as the bootstrap filter) is a Markov chain with a mutation and a selection transition,

$$(\xi_n^{i-})_{1 \leq i \leq N} \in \mathbb{R}^N \xrightarrow{\text{selection}} (\xi_n^i)_{1 \leq i \leq N} \in \mathbb{R}^N \xrightarrow{\text{mutation}} (\xi_{n+1}^{i-})_{1 \leq i \leq N} \quad (7.22)$$

The initial configuration $(\xi_0^{i-})_{1 \leq i \leq N}$ is defined by sampling N independent copies of $\mathcal{X}_0$. In its simplest form, the selection transition involves sampling N independent random variables $(\xi_t^i)_{1 \leq i \leq N}$ with the weighted distribution

$$\sum_{1 \leq i \leq N} \frac{e^{-(\mathcal{Y}_t - H\xi_t^{i-})^2/(2R_1)}}{\sum_{1 \leq j \leq N} e^{-(\mathcal{Y}_t - H\xi_t^{j-})^2/(2R_1)}} \delta_{\xi_t^{i-}} \quad (7.23)$$

The corresponding mutation transition coincides with prediction stage of the **VENKF** above; that is for any $1 \leq N$ we set

$$\xi_{t+1}^{i-} = A\xi_t^i + R^{1/2}\mathcal{V}_{t+1}^i. \quad (7.24)$$

In this context, the conditional means $X_t = \mathbb{E}(\mathcal{X}_t \mid \mathcal{Y}_t)$ are approximated by the sample means

$$\widehat{X}_t^{\text{PF}} := \frac{1}{N} \sum_{1 \leq i \leq N} \xi_t^i \quad (7.25)$$

A mature literature on the time-uniform contraction/convergence and fluctuation results for the particle filter exists, and a survey of this topic is beyond the scope of this article. However, typically the time-uniform particle filtering estimates discussed in the literature rely on mixing-type or certain contractive conditions on the mutation transitions, e.g. [49, 47, 114, 142, 48, 56]. In the case of linear-Gaussian state transition models, none of these conditions hold for general unstable transient signals.

A natural question: Can the particle filter track unstable (latent) signals, like the **VENKF** can (as per Theorem 7.1), in the simple linear-Gaussian setting of this section? Unfortunately, as we now show (and contrast with the **VENKF** and Theorem 7.1) the answer here is rather negative.

Following ideas of Mathieu Gerber (personal communication), whenever $A > 1$ we have

$$\xi_t^{i-} \geq A \inf_{1 \leq i \leq N} \xi_{t-1}^i + R^{1/2}\mathcal{V}_t^{*-} \quad \text{and} \quad \xi_t^i \geq \inf_{1 \leq i \leq N} \xi_t^{i-} \quad (7.26)$$

with

$$\mathcal{V}_t^{*-} := \inf_{1 \leq i \leq N} \mathcal{V}_t^i \quad (7.27)$$

This implies that,

$$\xi_t^{i-} \wedge \xi_t^i \geq X_t^{*-} := AX_{t-1}^{*-} + R^{1/2}\mathcal{V}_t^{*-} \quad \text{with} \quad X_0^{*-} = \inf_{1 \leq i \leq N} \xi_0^{i-} \quad (7.28)$$

Thus, for any given initial conditions $\xi_0^{i-} = x_0^i$ with

$$\inf_{1 \leq i \leq N} x_0^i \geq \epsilon + \frac{R^{1/2}}{A-1} \sqrt{2 \log(N)} \quad (7.29)$$

for some $\epsilon > 0$, we have

$$\inf_{1 \leq i \leq N} \mathbb{E}(\xi_t^i \mid \mathcal{Y}_t) \geq A^t \epsilon + A^t \frac{R^{1/2}}{A-1} \sqrt{2 \log(N)} + R^{1/2} \sum_{0 \leq s < t} A^s \mathbb{E}[\mathcal{V}_{t-s}^{*-}] \quad (7.30)$$

Recalling that $\mathbb{E}[\max_{1 \leq i \leq N} U_i] \leq \sqrt{2 \log(N)}$ for any sequence of N independent centered Gaussian random variables U_i with unit variance, we conclude that

$$\begin{aligned} \inf_{1 \leq i \leq N} \mathbb{E}[\xi_t^i \mid \mathcal{Y}_t] &\geq A^t \epsilon + A^t R^{1/2} \sqrt{2 \log(N)} \left(\frac{1}{A-1} - \sum_{1 \leq s \leq t} A^{-s} \right) \\ &= A^t \epsilon + \frac{R^{1/2}}{A-1} \sqrt{2 \log(N)} \end{aligned} \quad (7.31)$$

This yields the almost sure divergence result

$$\lim_{t \rightarrow \infty} \mathbb{E} \left[\widehat{X}_t^{\text{PF}} \mid \mathcal{Y}_t \right] = +\infty \quad (7.32)$$

This result is not restricted to proportional selection, but rather holds for any unbiased selection transition. Importantly, we emphasise that this result is true even when the unstable latent signal moves to $-\infty$; in which case there is very quickly a drastic divergence between the particle filtering estimate and the latent signal (and the optimal filter). This already indicates that the particle filter is not able to track unstable signals.

Theorem 7.2. *For any non-zero values of the model parameters (H, R, R_1) , and any $A > 1$ and $P_0^- > 0$, any $n \geq 1$, and for any $N \geq 1$ we find,*

$$\lim_{t \rightarrow \infty} \mathbb{E} \left[\left| \widehat{X}_t^{\text{PF}} - X_t \right|^n \right]^{1/n} = \infty = \lim_{t \rightarrow \infty} \mathbb{E} \left[\left| \widehat{X}_t^{\text{PF}} - \mathcal{X}_t \right|^n \right]^{1/n} \quad (7.33)$$

Before we proceed to the proof, we contrast the preceding result with the time-uniform bound and error control achievable with the EnKF detailed in Theorem 7.1 above. Theorem 7.2 states that there is no hope in stabilising the particle filtering estimate around an unstable and transient (latent) signal when using the same mutation-prediction (7.24) as the EnKF (and which is common in basic particle filtering implementations, cf. [64, 45, 53, 58, 84, 85]). Increasing the number of particles yields no benefits here. More specifically, and in contrast with the EnKF, the selection-correction stage of the above particle filter cannot compensate or correct for an unstable prediction-mutation.

Proof. Following the proof of (7.32), for any $A > 0$, note that

$$\xi_t^{i-} \geq A \inf_{1 \leq i \leq N} \xi_{t-1}^i + R^{1/2} \mathcal{V}_t^* \quad \text{and} \quad \xi_t^i \geq \inf_{1 \leq i \leq N} \xi_t^{i-} \quad (7.34)$$

with

$$\mathcal{V}_t^* := (-\mathcal{V}_t) \wedge \mathcal{V}_t^{*-} \quad (7.35)$$

Define,

$$X_t^* := A X_{t-1}^* + R^{1/2} \mathcal{V}_t^* \quad \text{with} \quad X_0^* := \inf_{1 \leq i \leq N} \xi_0^{i-} \quad (7.36)$$

Then it follows that,

$$\xi_t^{i-} \wedge \xi_t^i \geq X_t^* \quad (7.37)$$

We also then have,

$$\begin{aligned} X_t^* - \mathcal{X}_t &= A(X_{t-1}^* - \mathcal{X}_{t-1}) + R^{1/2} (\mathcal{V}_t^* + (-\mathcal{V}_t)) \\ &\geq A(X_{t-1}^* - \mathcal{X}_{t-1}) + 2R^{1/2} \mathcal{V}_t^* \end{aligned} \quad (7.38)$$

which implies that,

$$\frac{X_t^* - \mathcal{X}_t}{A^t} \geq (X_0^* - \mathcal{X}_0) + 2R^{1/2} \sum_{1 \leq k \leq t} \frac{\mathcal{V}_k^*}{A^k} \quad (7.39)$$

Thus, for any $A > 1$ and $\epsilon > 0$ on the event

$$\Omega_\epsilon := \left\{ \mathcal{X}_0 \leq 0 \quad \text{and} \quad X_0^* \geq \epsilon + \frac{2R^{1/2}}{A-1} \sqrt{2 \log(N+1)} \right\} \quad (7.40)$$

we readily check that

$$\mathbb{E} \left[\widehat{X}_t^{\text{PF}} - \mathcal{X}_t \mid X_0^*, \mathcal{X}_0 \right] \geq \mathbb{E} [X_t^* - \mathcal{X}_t \mid X_0^*, \mathcal{X}_0] \geq \epsilon A^t \xrightarrow{t \rightarrow \infty} \infty \quad (7.41)$$

Moreover, we have,

$$\begin{aligned}\mathbb{E} \left[\left| \hat{X}_t^{\text{PF}} - \mathcal{X}_t \right| \mid X_0^*, \mathcal{X}_0 \right] 1_{\Omega_\epsilon} &\geq \left| \mathbb{E} \left[\hat{X}_t^{\text{PF}} - \mathcal{X}_t \mid X_0^*, \mathcal{X}_0 \right] \right| 1_{\Omega_\epsilon} \\ &= \mathbb{E} \left[\hat{X}_t^{\text{PF}} - \mathcal{X}_t \mid X_0^*, \mathcal{X}_0 \right] 1_{\Omega_\epsilon} \\ &\geq \epsilon 1_{\Omega_\epsilon} A^t\end{aligned}\tag{7.42}$$

In discrete-time, the variance $P_t = \mathbb{E}[(X_t - \mathcal{X}_t)^2]$ is also uniformly bounded with respect to any time horizon, e.g. see [4, 92] (and similarly to (2.19) in continuous-time), and thus we have,

$$\epsilon \mathbb{P}[\Omega_\epsilon] A^t \leq \mathbb{E} \left[\left| \hat{X}_t^{\text{PF}} - \mathcal{X}_t \right| \right] \leq \mathbb{E} \left[\left| \hat{X}_t^{\text{PF}} - X_t \right| \right] + c\tag{7.43}$$

for some finite constant $c < \infty$.

Finally we confirm the non-zero probability,

$$P_0^- > 0 \implies \mathbb{P}[\Omega_\epsilon] = \mathbb{P}[\mathcal{X}_0 \leq 0] \mathbb{P} \left[\mathcal{X}_0 \geq \epsilon + \frac{2R^{1/2}}{A-1} \sqrt{2 \log(N+1)} \right]^N > 0\tag{7.44}$$

This ends the proof of the theorem. $\square$

To conclude this discussion, we note briefly that it is possible to stabilise the prediction-mutation step (about an unstable latent signal) at the expense of also changing the selection-correction stage (i.e. via importance sampling). For example, a mutation-prediction step sampled according to the so-called optimal proposal $\sim \text{Law}(x_t \mid x_{t-1}, \mathcal{Y}_t)$, see [58], or earlier in [46, Example 3], [47, Sections 2.4.3 and 12.6.6], is a stable option whenever $A/(1 + H^2 R/R_1) < 1$. However, in general nonlinear filtering problems these mutation transitions and the corresponding importance selection weights are intractable [58]. The terminology “optimal proposal”, see [58], is somehow confusing as this importance sampling strategy and the one discussed in (7.24) have the same Feynman-Kac-type mathematical structure and sampling according to this proposal doesn’t minimize the asymptotic variance. In this context, following [46, Section 4.2.2], we can use an auxiliary local particle approximation to sample $\text{Law}(x_t \mid x_{t-1}, \mathcal{Y}_t)$ and compute the corresponding importance weights. Given the topic of this article and this section in particular, we note an interesting approach in [115] employing an ensemble Kalman filter to define a proposal distribution (i.e. in the mutation step) that depends on the observation history [46, 115]. We do not explore this topic in further generality here.

8 Some Topics for Discussion

8.1 Comments on the Results Presented

In places, we switch between rather quantitative estimates to those more qualitative in nature. In part this is to simplify presentation, or when the details are (likely) not tight and thus perhaps of little quantitative interest. In some in places it is because we did not obtain more precise descriptions of the estimates involved. Refining these estimates may be of practical interest in some cases; e.g. when deriving estimates on the required number of particles $\mathbf{N}$ for stability of the sample covariance (or convergence to its invariant measure).

The results presented thus far consisted of constants, e.g. c , c_n , c_τ , etc, that depend on the model parameters (A, R, S) , but importantly not on the ensemble size $(\mathbf{N} + 1)$ or the time horizon $t \in [0, \infty[$. Due to the dependence on the model (e.g. (A, R, S)), these constants depend implicitly (via the matrix norms used) on the underlying signal dimension d . It would be of interest to pull this dependence out more explicitly depending on the matrix norm we are using, so as to quantify, at least in some general sense, the tradeoff between $\mathbf{N}$ and d . For example, in

Theorem 5.4 or Theorem 5.10 detailing the fluctuation of the sample covariance and sample mean about their limiting covariance and (Kalman-Bucy) state estimate values, it would be of interest to know how this fluctuation scales with dimension d , say e.g. with fixed N . Unfortunately, the proof tools used in the development of this work does not lend itself naturally to this analysis.

The matrix $S := H'R_1^{-1}H$ plays a critical role throughout with regards to obtaining time-uniform fluctuation and then subsequently stability/convergence results. In particular, the assumption that $S \in \mathbb{S}_d^+$ is strictly positive-definite, i.e. Assumption O, is needed in numerous places. This assumption amounts to a type of strong observability result; e.g. a requirement on the “fullness” of the observations and the size and rank of the observation matrix H . It is worth emphasising that this assumption appears in many technical articles discussing the performance properties of the ensemble Kalman filter; e.g. [82, 54, 138, 52, 43, 44]. Typically, the tools used in the proofs in [54, 22, 21, 19] are not sophisticated enough to accommodate zero eigenvalues of S . A basic example of this deficiency is in the proof of time-uniform moment boundedness of $\hat{P}_t$, stated in Theorem 5.3. In that proof, we resort to taking trace or eigenvalue-type reductions of the matrix-valued Riccati diffusion and studying a scalar comparison Riccati equation. This scalar reduction means that we must look at the minimum eigenvalue of S (because it appears with a minus sign in the Riccati equation) and thus we cannot allow this value to be zero (because we would lose this term completely in the scalar comparison). To obtain uniform-in-time bounds, one needs the stabilising effect of this non-zero S in the scalar comparison. See the proof in [19, Theorem 2.2] for this very transparent example. In this example, one may relax the condition on S to $S \in \mathbb{S}_d^0$ at the expense of time exponentially growing bounds. Related difficulties in allowing $S \in \mathbb{S}_d^0$ instead of $S \in \mathbb{S}_d^+$ arise in numerous other places (and as noted in other related works [82, 54, 138, 52, 43, 44]). One difficulty is related to stability of the (time-varying) matrix $(A - \hat{P}_t S)$ and the positive-definiteness properties of product $\hat{P}_t S$ as discussed subsequently.

We have focused significant effort on relaxing the assumption that the underlying signal is stable. Note that if A is stable, i.e. $\text{Abse}(A) < 0$, then the stability of $\mu(A - \hat{P}_t S)$, for some log-norm, may be trivially inherited whenever $S \in \mathbb{S}_d^+$ via a change of coordinates, see [54]. We see here again the use of $S \in \mathbb{S}_d^+$ as it pertains to the product $\hat{P}_t S$. If $S \in \mathbb{S}_d^0$ is only positive semi-definite, then one can construct counterexamples such that even if $\hat{\phi}_t = (\phi_t + \hat{\varphi}_t/\sqrt{N}) \in \mathbb{S}_d^+$ is positive definite, there exists flows $\hat{\varphi}_t$ such that $\mu(A - \hat{P}_t S) = \mu(A - \phi_t S - \hat{\varphi}_t S/\sqrt{N}) > 0$. The fluctuation term $\hat{\varphi}_t/\sqrt{N}$ *might* not interact well with the only positive semi-definite $S \in \mathbb{S}_d^0$. The assumption $\mu(A) < 0$ is made in [54] in the linear-Gaussian setting and follow also in, e.g., [82, 138, 52, 44] when reducing those studies to the linear-Gaussian setting.

If A is allowed to be unstable, then the asymptotic (time-varying) stability of $(A - P_t S)$ in the classical Kalman filter follows under so-called detectability (or observability) conditions [90, 116, 140]. Detectability intricately relates the relevant rank deficient directions in P_t and S in terms of the unstable directions in A (i.e. it basically ensures those directions of A that are unstable are observed (as captured by S) and non-zero weighted in the update Kalman gain via P_t). The rank of the sample covariance $\hat{P}_t$ is at most $N \geq 1$. If $N < d$, then $\hat{P}_t$ is almost surely rank deficient and thus has zero eigenvalues in some directions. In general, we cannot control the directions in which the random, sub-rank, $\hat{P}_t$ has zero eigenvalues (e.g. to play nicely with S in the sense of detectability). If A is unstable in those directions, the filter is consequently unstable in those directions. Thus, there is a basic, unavoidable, but also transparent tradeoff in requiring either stability of $\text{Abse}(A) < 0$ or sufficiently large ensemble sizes $N \geq d$ in the derivation of uniform-in-time stability results for the EnKF.

In the stability results stated in this work, we emphasised unstable models A but required sufficiently large ensemble sizes $N \geq d$. Nevertheless, most stability results stated in this work with the hypothesis that “ N is sufficiently large” may be restated with this condition replaced with “ $N \geq 1$ and $\text{Abse}(A) < 0$ ”, and such results hold time-uniformly over infinite time horizons. In [22, 21, 19] the details on “ N is sufficiently large” are given more explicitly. Note some

results that do not consider or rely on the long time stability behaviour of the samples, e.g. the fluctuation size of the sample covariance about its true value, hold with $N \geq 1$ and any matrix A , e.g. this is true for the **DEnKF** in Theorem 5.4.

We remark that the assumption that the true Kalman-Bucy filter is stable in the sense $\mu(A - P_\infty S) < 0$ is used in a number of the fluctuation (on the sample mean) and long-time behavioural results given in the cited, prior, work [54, 22, 21, 19, 17]. It was originally believed by the authors that this condition was a stronger assumption than the more natural condition $\text{Absc}(A - P_\infty S) < 0$. The latter follows from the very natural model assumptions of detectability and stabilisability, see (2.11) and the discussion following that equation, e.g. [92, Theorems 9.12, 9.15]. However, the particular logarithmic norm $\mu(\cdot)$ used throughout the prior work in [54, 22, 21, 19, 17] is unimportant, i.e. the matrix norm defining the log-norm can be chosen arbitrarily. From [132, Theorem 5], it is now known that if $\text{Absc}(A - P_\infty S) < 0$, then there exists a particular log-norm $\mu(\cdot)$ such that $\mu(A - P_\infty S) < 0$. Thus, the conditions antecedent in much of the prior work, in which we ask for $\mu(A - P_\infty S) < 0$ can be replaced with just asking for classical detectability and stabilisability conditions. (In prior work it was discussed and claimed by the authors that asking for $\mu(A - P_\infty S) < 0$ may be viewed as asking for a type of strong observability and controllability. This may be true if one specifies first the log-norm of interest. But if one does not care which log-norm is used, we can significantly relax the setting and simply ask for detectability/stabilisability which leads to $\text{Absc}(A - P_\infty S) < 0$). Much of the analysis, as already discussed, requires $S \in \mathbb{S}_d^+$ which automatically implies detectability (in fact a much stronger condition than observability).

It is worth noting again that all moment boundedness and fluctuation results stated in this work hold with any $N \geq 1$ and without further assumptions if one replaces the constants $c, c_n, c_n(Q), c_n(z, Q) \dots$ with functions that depend on (and grow with) the time horizon $t \geq 0$.

8.2 Bridging the Gap to Nonlinear Ensemble Filtering

The focus of this article is ensemble filtering in the linear-Gaussian (continuous-time) setting. The results surveyed herein portray a rather detailed theory of fluctuation and stability/contraction results in that case. In practice, the ensemble Kalman filtering methodology is applied in high-dimensional, nonlinear state-space models [60, 61]. The evolution equations for each ensemble member in the case of nonlinear state-space models are given in (4.10).

In [41] and in [145, 133] the novel idea of a McKean-Vlasov-type diffusion which has conditional distribution equal to the true Bayesian filter is studied, see also [117]. The mean-field approximation of this McKean-Vlasov-type diffusion in [145, 133], termed the *feedback particle filter*, resembles somewhat superficially the ensemble filters in (4.10). However, the analogue of the gain function in (4.10) in the feedback particle filter of [145, 133] is derived as the solution of a certain Poisson-type partial differential equation. In the linear-Gaussian case, the filter of [145] coincides with the **DEnKF**.

In the nonlinear model setting, the ensemble filters in (4.10) are not derived as sampled versions of an equation whose (conditional) distribution is equal to the Bayesian filter. That is, these filters are not derived as sampled versions of the McKean-Vlasov-type diffusion in [145, 133]. Conversely, in the limit ($N \rightarrow \infty$) the ensemble filters in (4.10) do not converge to an object with distribution equal to the optimal Bayes filter. In fact, the object these filters converge to has not been rigorously established in general and its properties, as compared to the true Bayesian filter, remain an open topic. Thus, in the nonlinear setting, the ensemble filters discussed in this work, see (4.10), may be viewed as approximations of the feedback particle filter in [145, 133] only in some very weak sense (despite any superficial resemblance to the contrary). Indeed, the gain function approximation in (4.10) is likely a very poor approximation of the solution of the Poisson-type partial differential equation in [145]; except of course in linear-Gaussian models. Rather, we may argue, as we have earlier in this article, that the ensemble filters in (4.10) should be viewed in the context of so-called observer theory, and related not

to Bayesian filtering but rather to the more general topic of (dynamic) state estimation [4, 10]. The goal of state estimation in this context is to design an observer that tracks in some suitable (typically point-wise) sense the underlying signal and perhaps provides some usable measure of uncertainty on this estimate. The goal is not to develop an approximation (at each time) of the true conditional (Bayesian) distribution of the signal given the observations. The latter contains significantly more information than is perhaps needed in many practical applications. Nevertheless, we also argue that the filtering ideas in [145, 133], and suitable approximations thereof, are in need of further investigation.

In [52] a class of so-called ensemble extended Kalman filters (**En-EKF**) is developed that is based on a type of particle approximation of the linearisation-based extended Kalman filter, see [4]. This ensemble filter is interesting because the sample mean is shown to converge (in $N \rightarrow \infty$) to the extended Kalman filter state estimate. This extended Kalman state estimator has been widely studied in nonlinear filtering and control theory [4, 10, 124, 25, 81], and may be viewed more as a type of nonlinear state estimator rather than a Bayesian filter [10, 124].

When considering nonlinear signal models, the long time behavioural analysis of various **EnKF** methods in [82, 138, 52] assumes a strong type of stability property on the signal (which in the linear case would reduce to assuming that A is Hurwitz stable in our model (2.1)). This stability assumption on the true signal is precisely what we aim to relax in our work; albeit limited in our study to linear models. Filter stability without assumptions on the stability of the true signal will ultimately require some control of the fluctuation properties of the sampled observer, e.g. see the discussion in the preceding section on this topic (in the linear-Gaussian model setting). This fluctuation analysis is lacking somewhat in the nonlinear model setting. It is complicated in that case by the absence of any closed-form evolution equations for the sample mean and sample covariance.

Viewing, or even designing, an ensemble filter (or its sample mean for example) as a (dynamic) state estimator (or observer) may have some benefits. In particular, stability may be a larger design consideration if starting from this viewpoint rather than seeking Bayesian probabilistic properties. It may be possible to then also exploit the properties of existing nonlinear state estimators which have traditionally been rigorously analysed, e.g. [10, 124, 25, 81].

This is exemplified in the (**En-EKF**) in [52] that converges to the extended Kalman filter in the limit $N \rightarrow \infty$. The stability of the extended Kalman filter as a nonlinear observer has been widely studied, e.g. see [10, 124, 25, 81]. Although strong signal stability assumptions are taken in [52], it would be natural to consider the (**En-EKF**) in [52] without the underlying signal stability assumption and look at developing the fluctuation type analysis considered herein in the linear-Gaussian setting. We may then also exploit the stability analysis that already exists [10, 124, 25, 81] for the limiting extended Kalman state estimator. This is analogous in many ways to the stability properties and observability/controllability properties used herein in the linear-Gaussian setting.

Inflation is used in [82, 139] in the nonlinear model setting to aid in stability. This is similar to the study considered herein on stability under inflation in linear-Gaussian models. It seems natural that added inflation acts to stabilise the various ensemble filters. In the context of the preceding discussion, inflation-based state estimators may also be viewed in the context of stable nonlinear observers, rather than heuristic adaptations of approximate Bayesian filters.

Finally, we remark that the transport-based ensemble filter **DEnTF**, see case (NF3) in (4.10), is studied in [43] in a particular nonlinear setting. Non-asymptotic (i.e. finite sample) uniform-in-time accuracy and stability of the **DEnTF** is studied in [43] under the assumptions of small observation noise, and a square observation matrix, or in other words with linear observations and a change of coordinate so that $H = I$. Note this latter assumption is made also in [82, 138, 52, 44] which otherwise consider certain classes of nonlinear signals and different **EnKF** variants. Thus, this strong (and linear) observability assumption seems key to analysis in the ensemble filtering literature even when moving away from the linear signal model.

8.3 Other Related Literature

The focus of this article is a detailed fluctuation and contraction analysis of the relevant ensemble Kalman filtering terms (e.g. the exponential semigroups, sample mean, and sample covariance) in the linear-Gaussian setting. There is considerable work on the periphery of this rather specific topic and analysis. A broad overview of the literature on filtering is not possible, but we note below some topics and literature for further study.

For example, the introduction of this article was focused primarily on (specific) related literature in ensemble Kalman approximation methods, and so-called data assimilation. The topic of filter stability in the case of the true nonlinear filter (e.g. as given by the Kallianpur-Striebel formula [77]) has been studied widely; see a broad but necessarily incomplete snapshot in [89, 112, 8, 14, 31, 11, 140, 141]. Known results in general suggest that sufficiently informative observation processes and/or sufficiently contractive/ergodic latent signals translate to contractive stability of the nonlinear filter. Emphasis on the continuous-time, linear-Gaussian model setting is studied in [3, 112, 140, 16, 18] where deeply understood observability and controllability model conditions explicitly formalise the properties leading to filter stability. See also [28] for a contraction analysis of a low-rank Kalman-Bucy filter particularly relevant in the application domain of this article.

The fluctuation and stability of other approximations schemes in nonlinear filtering have also been studied; e.g. see [97, 32, 47, 39, 114, 68, 142, 144, 56, 15] for a snapshot of some of these methods. A detailed discussion of these approximation methods and their fluctuation and stability properties is beyond the scope of the ensemble Kalman-type methods studied here. However, the monograph [47] provides a detailed study of the fluctuation and contractive properties of numerous particle filtering methods. We note in passing that in general when studying the stability of filtering approximation schemes it is rather common to assume the latent signal is stable/ergodic and/or the observation sequence is stationary, e.g. see [32, 47, 142, 56, 15]. In this article, and in prior work [54, 22, 21, 19], we relax those assumptions and primarily rely on a strong form of observability (i.e. a strong form of observational informativeness). The results presented here *do not* rely on *any* form of latent signal stability in general. This latter fact distinguishes this work, but also the ensemble Kalman approximation method in its ability to handle totally unstable latent signals. The latter point was illustrated above via comparisons with the bootstrap particle filter in one-dimensional unstable linear-Gaussian models.

In the introduction we list some of the seminal data assimilation and ensemble Kalman methodology articles. Here, we give a by no means complete list of some methodological approaches that in some sense can be considered cousins of

Different particle filtering methods based on evolving particles according to an ordinary differential equation that bridges the prior or predictive distribution with the posterior have been developed, e.g. see [42, 33]. These so-called Bayesian homotopy methods aim to introduce nonlinearities arising from the Bayes update in a tempered fashion. See also [120] for an introduction with connections to ensemble Kalman filtering. Applications of the multilevel Monte Carlo method [63] to ensemble Kalman filtering have been considered in, e.g., [69, 38]. The continuous-time linear-Gaussian case similarly formulated as in this article is considered in detail in [38]. A related extension in multi-index ensemble Kalman filtering was proposed in [70]. The problem of unbiased ensemble Kalman filtering has been considered in [2] with an emphasis and analysis closely related to the formulation considered herein.

The ensemble Kalman filter has been applied to the problem of (log) normalisation constant estimation for continuous-time filtering problems; e.g. see [40, 125]. More generally, ensemble Kalman methods for inverse problems have also been considered in the literature [75, 37] with some related analysis [128, 129]. Particle filtering and ensemble filtering methods have also been applied in optimisation [146]. See these references for further details on the respective topics.

It was noted in the preceding section that certain McKean-Vlasov equations can be derived [145] that superficially resemble the ensemble filters in (4.10), but with gain functions derived as

the solution of certain Poisson-type partial differential equations. In general, the (conditional) law of these nonlinear diffusions is equal to the filtering distribution. Applying mean-field particle approximations in the linear-Gaussian case, the filter of [145] coincides with the **DEnKF**. In the general nonlinear setting, other gain function approximations can be employed such as discussed in [135, 134].

Finally, we remark in passing that different models involving backward matrix Riccati diffusions arise in linear-quadratic optimal control problems with random coefficients; see e.g. [27, 73, 86]. Another class of random Riccati equations, different from the Riccati (matrix quadratic) diffusion equations studied herein arises in network control and filtering with random observation losses; see e.g. [131, 136]. The details of these works are beyond the scope of the forward-in-time Riccati diffusions considered herein. The forward-in-time Riccati diffusion of (5.2) is nevertheless of interest on its own, as with $\kappa = 0$ it is a prototypical model of a matrix stochastic differential equation with a quadratic drift term. In the deterministic setting, it is worth noting that the contraction of the Riccati equation, e.g. (2.3), (2.20), can naturally be studied with different metrics, e.g. in the Riemannian space of positive definite matrices, see e.g. [29, 101, 28, 102], and such contraction results may be of interest and/or practical value in the stochastic setting of the Riccati diffusion describing the flow of sample covariances.

8.4 Some Open Problems

The stochastic analysis and stability of the **EnKF** models considered in this article are rather well understood in the linear-Gaussian, continuous-time, setting *even with unstable latent signal processes*. However, there are still some worthy gaps. In particular, we may contrast the very strong and complete picture in the one-dimensional setting (following from [21] and presented above) with the more general and relevant multivariate setting.

For example, the $\mathbb{L}_n$ contraction estimates on the Riccati diffusion flow describing the sample covariance in (6.27) are only available in the scalar case. We know, e.g. see Theorem 5.5, that the Markov semigroup is contractive and that the law of the sample covariance converges to an invariant measure in the multi-dimensional setting, analogously to the convergence of the deterministic Riccati equation to an equilibrium state. However, multi-variate versions of the contraction estimates (6.27) are unknown.

Moreover, multi-variate contraction estimates for the **EnKF** mean as in (6.28) are known only on certain finite time intervals (with a finite sample size) and a general multi-variate filtering contraction estimate as in (6.28) under mild assumptions would be of interest. There is also a disparity in available results in the multi-variate setting between the **VEnKF** and **DEnKF**.

In both (6.27) and (6.28) in the scalar case we recover in the limit with the ensemble size the deterministic optimal rates of the classical Kalman-Bucy filter. Extensions of these rate estimates to the multi-variate setting would also be of interest.

Finally in reference to the scalar setting, we can explicitly state in closed form the invariant measure of the sample covariance, e.g. see (6.19) for the **VEnKF** and (6.21) for the **DEnKF**, and Figure 3. It is unlikely that closed-form expressions are possible to derive in the multivariate setting, however, it would be of further interest to confirm various properties such as the heavy-tailed nature of the stationary measure in the case of the **VEnKF**. Such properties as discussed earlier have practical consequences such as the potential to lead to so-called catastrophic divergence and numerical instability.

While not given explicitly, we may crudely introduce the signal dimension into the constants of the results presented in this work. However, a detailed study of the dimension as it pertains to stochastic fluctuation properties of the relevant sample covariance and subsequently ensemble mean is warranted since dimensionality versus computational expense is a primary driver of these methods in practice. Work in this direction with various covariance inflation mechanisms was considered in [24] in the linear-Gaussian setting and more generally in [44] for the transport-inspired ensemble Kalman-Bucy filter **DEnTF**.

The main assumptions employed throughout are Assumptions O and C. The Assumption C can be relaxed to just stabilisability, and is thus completely in line with classical Kalman filtering analysis. It would be of interest to relax Assumption O to observability or even just detectability in line with classical Kalman-Bucy filtering. We have discussed in the preceding sections various issues surrounding the need for this stronger observability-type Assumption O, and difficulties with its possible relaxation.

This article focused on the continuous-time linear-Gaussian model case, and the preceding open problems also fall under this setting. Direct extensions of the results presented in this article to the discrete-time linear-Gaussian model setting would be of interest and some work has been published in this case, particularly in the scalar setting, e.g. see [51] and Theorem 7.1 and the earlier discussion. Finally, without considering specifics and re-discussing the relevant literature, extensions of these results to the nonlinear model setting is of great practical interest, particularly under testable and natural model assumptions (e.g. observability-type assumptions) also accommodating unstable or transient latent signal processes (with the latter relaxation being a major driver of the results presented herein).

References

- [1] J.I. Allen, M. Eknes and G. Evensen. An ensemble Kalman filter with a complex marine ecosystem model: Hindcasting phytoplankton in the Cretan Sea. *Annales Geophysicae*. vol. 21. pp. 399–411 (2003).
- [2] M. Alvarez, N.K. Chada, and A. Jasra. Unbiased Estimation of the Vanilla and Deterministic Ensemble Kalman-Bucy Filters. arXiv e-print, [arXiv:2208.03947](https://arxiv.org/abs/2208.03947) (2022).
- [3] B.D.O. Anderson. Stability properties of Kalman-Bucy filters. *Journal of the Franklin Institute*. vol. 291, no. 2. pp. 137–144 (1971).
- [4] B.D.O. Anderson and J.B. Moore. *Optimal Filtering*. Dover Publications (1979).
- [5] J.L. Anderson. An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*. vol. 129, no. 12. pp. 2884–2903 (2001).
- [6] J.L. Anderson. A local least squares framework for ensemble filtering. *Monthly Weather Review*. vol. 131, no. 4. pp 634–642 (2003).
- [7] J.L. Anderson and S.L. Anderson. A Monte Carlo Implementation of the Nonlinear Filtering Problem to Produce Ensemble Assimilations and Forecasts. *Monthly Weather Review*. vol. 127, no. 12. pp. 2741–2758 (1999).
- [8] R. Atar. Exponential stability for nonlinear filtering of diffusion processes in a noncompact domain. *Annals of Probability*. pp. 1552–1574 (1998).
- [9] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*. Springer (2009).
- [10] J.S. Baras, A. Bensoussan, and M.R. James. Dynamic observers as asymptotic limits of recursive filters: Special cases. *SIAM Journal on Applied Mathematics*. vol. 48, no. 5. pp. 1147–1158 (1988).
- [11] P. Baxendale, P. Chigansky, and R. Liptser. Asymptotic stability of the Wonham filter: ergodic and nonergodic signals. *SIAM Journal on Control and Optimization*. vol. 43, no. 2. pp. 643–669 (2004).
- [12] V.E. Beneš. Exact finite-dimensional filters for certain diffusions with nonlinear drift. *Stochastics*. vol. 5, no. 1-2. pp. 65–92 (1981).

- [13] K. Bergemann and S. Reich. An ensemble Kalman-Bucy filter for continuous data assimilation. *Meteorologische Zeitschrift*. vol. 21, no. 3. pp. 213–219 (2012).
- [14] A.G. Bhatt, A. Budhiraja, and R.L. Karandikar. Markov property and ergodicity of the nonlinear filter. *SIAM Journal on Control and Optimization*. vol. 39, no. 3. pp. 928–949 (2000).
- [15] A.N. Bishop and E.V. Bonilla. Recurrent Neural Networks and Universal Approximation of Bayesian Filters. arXiv e-print, [arXiv:2211.00335](#) (2022).
- [16] A.N. Bishop and P. Del Moral. On the Stability of Kalman-Bucy Diffusion Processes. *SIAM Journal on Control and Optimization*. vol. 55, no. 6. pp. 4015–4047 (2017); arxiv e-print [arXiv:1610.04686](#) updated.
- [17] A.N. Bishop and P. Del Moral. Stability Properties of Systems of Linear Stochastic Differential Equations with Random Coefficients. *SIAM Journal on Control and Optimization*. vol. 57, no. 2. pp. 1023–1042 (2019); arXiv e-print, [arXiv:1804.09349](#).
- [18] A.N. Bishop and P. Del Moral. An explicit Floquet-type representation of Riccati aperiodic exponential semigroups. *International Journal of Control*. vol. 94, no. 1. pp. 258–266 (2021); arXiv e-print, [arXiv:1805.02127](#) (2018).
- [19] A.N. Bishop and P. Del Moral. On the stability of matrix-valued Riccati diffusions. *Electronic Journal of Probability*. vol. 24, paper no. 24 (2019); arXiv e-print, [arXiv:1808.00235](#).
- [20] A.N. Bishop and P. Del Moral. On the robustness of Riccati flows to complete model misspecification. *Journal of the Franklin Institute*. vol. 355, no. 15. pp. 7178–7200 (2018).
- [21] A.N. Bishop, P. Del Moral, K. Kamatani, and B. Remillard. On one-dimensional Riccati diffusions. *Annals of Applied Probability*. vol. 29, no. 2. pp. 1127–1187 (2019); arXiv e-print, [arXiv:1711.10065](#).
- [22] A.N. Bishop, P. Del Moral, and A. Niclas. A perturbation analysis of stochastic matrix Riccati diffusions. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*. vol. 56, no. 2. pp. 884–916 (2020); arXiv e-print, [arXiv:1709.05071](#).
- [23] A.N. Bishop, P. Del Moral and A. Niclas. An introduction to Wishart matrix moments. *Foundations and Trends in Machine Learning*. vol. 11, no. 2. pp. 97–218 (2018); arXiv e-print, [arXiv:1710.10864](#).
- [24] A.N. Bishop, P. Del Moral and S. Pathiraja. Perturbations and Projections of Kalman-Bucy Semigroups. *Stochastic Processes and their Applications*. vol. 128, no. 9. pp. 2857–2904. (2018).
- [25] A.N. Bishop and P. Jensfelt. A stochastically stable solution to the problem of robocentric mapping. In *Proceedings of the 2009 IEEE International Conference on Robotics and Automation*, Kobe, Japan (May, 2009).
- [26] C.H. Bishop, B.J. Etherton, and S.J. Majumdar. Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Monthly Weather Review* vol. 129, no. 3. pp. 420–436 (2001).
- [27] J.M. Bismut. Linear Quadratic Optimal Stochastic Control with Random Coefficients. *SIAM Journal on Control and Optimization*. vol. 14, no. 3. pp. 419–444 (1976).
- [28] S. Bonnabel and R. Sepulchre. The Geometry of Low-Rank Kalman Filters. *Matrix Information Geometry* (eds: F. Nielsen, R. Bhatia). pages 53–68. Springer (2013). See also: arXiv e-print, [arXiv:1203.4049](#) (2012).

- [29] P. Bougerol. Kalman filtering with random coefficients and contractions. *SIAM Journal on Control and Optimization*. vol. 31, no. 4. pp. 942–959 (1993).
- [30] R.S. Bucy. Global Theory of the Riccati Equation. *Journal of Computer and System Sciences*. vol. 1. pp. 349–361 (1967).
- [31] A. Budhiraja. Asymptotic stability, ergodicity and other asymptotic properties of the nonlinear filter. *Annales de l’IHP Probabilites et Statistiques*. vol. 39, no. 6. pp. 919–941 (2003).
- [32] A. Budhiraja and H.J. Kushner. Approximation and limit results for nonlinear filters over an infinite time interval. *SIAM Journal on Control and Optimization*. vol. 37, no. 6. pp. 1946–1979 (1999).
- [33] P. Bunch and S. Godsill. Approximations of the optimal importance density using Gaussian particle flow importance sampling. *Journal of the American Statistical Association*. vol. 111, no. 514. pp. 748–762 (2016).
- [34] G. Burgers, P.J. van Leeuwen, and G. Evensen. Analysis Scheme in the Ensemble Kalman Filter. *Monthly Weather Review*, vol. 126, no. 6. pp. 1719–1724 (1998).
- [35] F.M. Callier and J.L. Willems. Criterion for the Convergence of the Solution of the Riccati Differential Equation. *IEEE Transactions on Automatic Control*. vol. 26, no. 6. pp. 1232–1242 (1981).
- [36] F.M. Callier and J. Winkin. Convergence of the Time-Invariant Riccati Differential Equation towards Its Strong Solution for Stabilizable Systems. *Journal of Mathematical Analysis and Applications*. vol. 192, no. 1. pp. 230–257 (1995).
- [37] N.K. Chada, M.A. Iglesias, L. Roininen, and A.M. Stuart. Parameterizations for ensemble Kalman inversion. *Inverse Problems*. vol. 34, no. 5. (2018).
- [38] N.K. Chada, A. Jasra, and F. Yu. Multilevel ensemble Kalman-Bucy filters. *SIAM/ASA Journal on Uncertainty Quantification*. vol. 10, no. 2. pp. 584–618 (2022).
- [39] N. Chopin. Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics*. vol. 32, no. 6. pp. 2385–2411 (2004).
- [40] D. Crisan, P. Del Moral, A. Jasra, and H. Ruzayqat. Log-normalization constant estimation using the ensemble Kalman-Bucy filter with application to high-dimensional models. *Advances in Applied Probability*. vol. 54, no. 4. pp. 1139–1163 (2022).
- [41] D. Crisan and J. Xiong. Approximate McKean–Vlasov representations for a class of SPDEs. *Stochastics*. vol 82, no. 1. pp 53–68 (2010).
- [42] F. Daum, J. Huang and A. Noushin. Exact particle flow for nonlinear filters. In *Proceedings of Signal Processing, Sensor Fusion, and Target Recognition XIX*, vol. 7697, of the International Society for Optics and Photonics (SPIE). pp. 92–110 (2010).
- [43] J. de Wiljes, S. Reich, and W. Stannat. Long-Time Stability and Accuracy of the Ensemble Kalman-Bucy Filter for Fully Observed Processes and Small Measurement Noise. *SIAM Journal on Applied Dynamical Systems*. vol. 17, no. 2. pp. 1152–1181 (2018).
- [44] J. de Wiljes and X.T. Tong. Analysis of a localised nonlinear Ensemble Kalman Bucy Filter with complete and accurate observations. *Nonlinearity*. vol. 33, no. 9. pp. 4752 (2020); arXiv e-print, [arXiv:1908.10580](https://arxiv.org/abs/1908.10580) (2019).
- [45] P. Del Moral. Non Linear Filtering: Interacting Particle Solution. *Markov Processes and Related Fields*. vol. 2, no. 4, pp. 555–580 (1996).

- [46] P. Del Moral. Measure-valued processes and interacting particle systems. Application to nonlinear filtering problems. *The Annals of Applied Probability*. vol. 8, no. 2, pp. 438–495 (1998).
- [47] P. Del Moral. *Feynman-Kac Formulae*. Springer (2004).
- [48] P. Del Moral. Mean field simulation for Monte Carlo integration. *Monographs on Statistics and Applied Probability* 126, no. 26 (2013).
- [49] P. Del Moral, A. Guionnet. On the stability of measure valued processes with applications to filtering. *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics*. vol. 329, no. 5. pp. 429–434 (1999).
- [50] P. Del Moral, A. Guionnet. On the stability of interacting processes with applications to filtering and genetic algorithms. *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*. vol. 37, no. 2, pp. 155–194 (2001).
- [51] P. Del Moral, E. Horton. A theoretical analysis of one-dimensional discrete generation ensemble Kalman particle filters. *arXiv e-print*, [arXiv:2107.01855](https://arxiv.org/abs/2107.01855) (2021).
- [52] P. Del Moral, A. Kurtzmann, and J. Tugaut. On the Stability and the Uniform Propagation of Chaos of a Class of Extended Ensemble Kalman–Bucy Filters. *SIAM Journal on Control and Optimization*. vol. 55, no.1. pp. 119–155 (2017).
- [53] P. Del Moral, L. Miclo. Branching and interacting particle systems approximations of Feynman-Kac formulae with applications to non-linear filtering. *Séminaire de Probabilités XXXIV*, pp. 1–145 (2000).
- [54] P. Del Moral and J. Tugaut. On the stability and the uniform propagation of chaos properties of ensemble Kalman-Bucy filters. *Annals of Applied Probability*. vol. 28, no. 2. pp 790–850 (2018).
- [55] J.L. Doob. *Stochastic Processes*. John Wiley & Sons, New York (1953).
- [56] R. Douc, E. Moulines, and J. Olsson. Long-term stability of sequential Monte Carlo methods under verifiable conditions. *The Annals of Applied Probability*. vol. 24, no. 5. pp. 1767–1802 (2014).
- [57] A. Doucet, N. de Freitas and N.J. Gordon (editors). *Sequential Monte Carlo Methods in Practice*. Springer (2001).
- [58] A. Doucet, S. Godsill and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*. vol. 10, no. 3. pp. 197–208 (2000).
- [59] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*. vol. 99, no. C5. pp. 10143–10162 (1994).
- [60] G. Evensen. The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation. *Ocean Dynamics*. vol. 53, no. 4. pp. 343–367 (2003).
- [61] G. Evensen. *Data Assimilation: The Ensemble Kalman Filter*. Springer Science & Business Media, 2nd edition, (2009).
- [62] G. Evensen, J. Hove, H.C. Meisingset, E. Reiso, K.S. Seim. Using the EnKF for assisted history matching of a North Sea Reservoir Model. In *Proceedings of the 2007 SPE Reservoir Simulation Symposium*, Houston, Texas (February, 2007).

- [63] M.B. Giles. Multilevel monte carlo path simulation. *Operations Research*. vol. 56, no. 3. pp. 607–617 (2008).
- [64] N. Gordon, J. Salmond and A. Smith. A novel approach to non-linear/non-Gaussian Bayesian state estimation. *IEE Proceedings on Radar and Signal Processing*. vol. 140, no. 2. pp. 107–113 (1993).
- [65] G.A. Gottwald and A.J. Majda. A mechanism for catastrophic filter divergence in data assimilation for sparse observation networks. *Nonlinear Processes in Geophysics*. vol. 20, no. 5. pp. 705–712 (2013).
- [66] T.M. Hamill, J.S. Whitaker, and C. Snyder. Distance-Dependent Filtering of Background Error Covariance Estimates in an Ensemble Kalman Filter. *Monthly Weather Review*. vol. 129, no. 11. pp. 2776–2790 (2001).
- [67] J. Harlim and A.J. Majda. Catastrophic Filter Divergence in Filtering Nonlinear Dissipative Systems. *Communications in Mathematical Sciences*. vol. 8, no. 1. pp. 27–43 (2010).
- [68] K. Heine and D. Crisan. Uniform approximations of discrete-time filters. *Advances in Applied Probability*. vol. 40, no. 4. pp. 979–1001 (2008).
- [69] H. Hoel, K.J.H. Law, and R. Tempone. Multilevel ensemble Kalman filtering. *SIAM Journal on Numerical Analysis*. vol. 54, no. 3. pp. 1813–1839 (2016).
- [70] H. Hoel, G. Shaimerdenova, and R. Tempone. Multi-index ensemble Kalman filtering. *Journal of Computational Physics*. vol. 470. (2022).
- [71] P.L. Houtekamer and H.L. Mitchell. Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*. vol. 126, no. 3. pp. 796–811 (1998).
- [72] P.L. Houtekamer and H.L. Mitchell. A Sequential Ensemble Kalman Filter for Atmospheric Data Assimilation. *Monthly Weather Review*. vol. 129, no. 1. pp. 123–137 (2001).
- [73] Y. Hu and X.Y. Zhou. Indefinite stochastic Riccati equations. *SIAM Journal on Control Optimization*. vol. 42, no. 1. pp. 123–137 (2003).
- [74] M. Hutzenthaler, A. Jentzen and P.E. Kloeden. Strong and weak divergence in finite time of Euler’s method for stochastic differential equations with non-globally Lipschitz continuous coefficients. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. vol. 467, no. 2130. pp. 1563–1576 (2011).
- [75] M.A. Iglesias, K.J.H. Law, and A.M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*. vol. 29, no. 4. (2013).
- [76] C.J. Johns and J. Mandel. A two-stage ensemble Kalman filter for smooth data assimilation. *Environmental and Ecological Statistics*. vol. 15, no. 1. pp. 101–110 (2008).
- [77] G. Kallianpur and C. Striebel. Stochastic differential equations occurring in the estimation of continuous parameter stochastic processes. *Theory of Probability & Its Applications*. vol. 14, no. 4. pp. 567–594 (1969).
- [78] R.E. Kalman and R.S. Bucy. New Results in Linear Filtering and Prediction Theory. *Journal of Basic Engineering*. vol. 83, no. 1. pp. 95–108 (1961).
- [79] E. Kalnay. *Atmospheric Modelling, Data Assimilation and Predictability*. Cambridge University Press (2003).
- [80] I. Karatzas and S.E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer (1996).

- [81] T. Karvonen, S. Bonnabel, E. Moulines and S. Särkkä. On stability of a class of filters for non-linear stochastic systems. arXiv e-print, [arXiv:1809.05667](https://arxiv.org/abs/1809.05667) (2018).
- [82] D. Kelly, K.H.J. Law, and A.M. Stuart. Well-posedness and accuracy of the ensemble Kalman filter in discrete and continuous time. *Nonlinearity*. vol. 27, no. 10. pp. 2579–2603 (2014).
- [83] D. Kelly, A.J. Majda, and X.T. Tong. Concrete ensemble Kalman filters with rigorous catastrophic filter divergence. *Proceedings of the National Academy of Sciences of the United States of America*. vol 112, no. 34. pp. 10589–10594 (2015).
- [84] G. Kitagawa. A Monte Carlo Filtering and Smoothing Method for Non-Gaussian Nonlinear State Space Models. *Proceedings of the 2nd U.S.-Japan Joint Seminar on Statistical Time Series Analysis*: pp. 110–131 (1993).
- [85] G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*. vol. 5, no.1, pp 1–25 (1996).
- [86] M. Kohlmann and S. Tang. Multidimensional backward stochastic Riccati equations and applications. *SIAM Journal on Control Optimization*. vol. 41, no. 6. pp. 1696–1721 (2003).
- [87] G.M. Krause. Bounds for the Variation of Matrix Eigenvalues and Polynomial Roots. *Linear Algebra and its Applications*. vol. 208-209. pp. 73–82 (1994).
- [88] V. Kucera. A Contribution to Matrix Quadratic Equations. *IEEE Transactions on Automatic Control*. vol. 17, no. 3. pp. 344–347 (1972).
- [89] H. Kunita. Asymptotic behavior of the nonlinear filtering errors of Markov processes. *Journal of Multivariate Analysis*. vol. 1, no. 4. pp. 365–393 (1971).
- [90] H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems*. Wiley-Interscience (1972).
- [91] E. Kwiatkowski and J. Mandel. Convergence of the Square Root Ensemble Kalman Filter in the Large Ensemble Limit. *SIAM/ASA Journal on Uncertainty Quantification*. vol. 3, no. 1. pp. 1–17 (2015).
- [92] P. Lancaster and L. Rodman. *Algebraic Riccati Equations*. Oxford University Press (1995).
- [93] T. Lange. Derivation of ensemble Kalman-Bucy filters with unbounded nonlinear coefficients. *Nonlinearity*. vol. 35, no. 2. (2021).
- [94] T. Lange and W. Stannat. Mean field limit of Ensemble Square Root Filters—discrete and continuous time. *Foundations of Data Science*. vol. 3, no. 3. pp. 563–588 (2021).
- [95] T. Lange and W. Stannat. On the continuous time limit of the ensemble Kalman filter. *Mathematics of Computation*. vol. 90, no. 327. pp. 233–265 (2021).
- [96] T. Lange and W. Stannat. On the continuous time limit of ensemble square root filters. *Communications in Mathematical Sciences*. vol. 19, no. 7. (2021).
- [97] B.F. La Scala, R.R. Bitmead, and M.R. James. Conditions for stability of the extended Kalman filter and their application to the frequency tracking problem. *Mathematics of Control, Signals and Systems*. vol. 8, no. 1. (1995).
- [98] K.J.H. Law, A.M. Stuart and K. Zygalakis. *Data Assimilation: A Mathematical Introduction*. Springer (2015).

- [99] K.J.H. Law, H. Tembine and R. Tempone. Deterministic Mean-Field Ensemble Kalman Filtering. *SIAM Journal on Scientific Computing*. vol. 38, no. 3. pp. A1251-A1279 (2016).
- [100] F. Le Gland, V. Monbet and V.D. Tran. Large sample asymptotics for the ensemble Kalman filter. Chapter 22 in *The Oxford Handbook of Nonlinear Filtering*. pp. 598–631 (2011).
- [101] H. Lee, and Y. Lim. Invariant metrics, contractions and nonlinear matrix equations. *Nonlinearity*. vol. 21, no. 4. pp. 857 (2008).
- [102] B.C. Levy, and M. Zorzi. A contraction analysis of the convergence of risk-sensitive filters. *SIAM Journal on Control and Optimization*. vol. 54, no. 4. pp. 2154–2173 (2016).
- [103] K.A. Lisaeter, J. Rosanova and G. Evensen. Assimilation of ice concentration in a coupled ice-ocean model using the Ensemble Kalman Filter. *Ocean Dynamics*. vol. 53, no. 4. pp. 368–388 (2003).
- [104] D.M. Livings, S.L. Dance and N.K. Nichols. Unbiased Ensemble Square Root Filters. *Physica D: Nonlinear Phenomena*. vol. 237, no. 8. pp. 1021–1028 (2008).
- [105] A.J. Majda and J. Harlim. *Filtering Complex Turbulent Systems*. Cambridge University Press (2012).
- [106] A.J. Majda and X.T. Tong. Performance of Ensemble Kalman Filters in Large Dimensions. *Communications in Mathematical Sciences*. vol. 71, no. 5. pp. 892–937 (2018).
- [107] J. Mandel, L. Cobb, and J.D. Beezley. On the convergence of the ensemble Kalman filter. *Applications of Mathematics*. vol. 56, no. 6. pp. 533–541 (2011).
- [108] H.P. McKean. A class of Markov processes associated with nonlinear parabolic equations. *Proceedings of the National Academy of Sciences*. vol. 56, no. 6. pp.1907–1911 (1966).
- [109] H.L. Mitchell, P.L. Houtekamer and G. Pellerin. Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Monthly Weather Review*. vol. 130, no. 11. pp. 2791–2808 (2002).
- [110] B.P. Molinari. The time-invariant linear-quadratic optimal control problem. *Automatica*. vol. 13, no. 4. pp. 347–357 (1977).
- [111] G. Naevdal, L.M. Johnsen, S.I. Aanonsen and E.H. Vefring. Reservoir monitoring and continuous model updating using ensemble Kalman filter. In *Proceedings of the 2003 SPE Annual Technical Conference and Exhibition*, Denver, Colorado (October, 2003).
- [112] D. Ocone and E. Pardoux. Asymptotic stability of the optimal filter with respect to its initial condition. *SIAM Journal on Control and Optimization*. vol. 34, no. 1. pp. 226–243 (1996).
- [113] E. Ott, B.R. Hunt, I. Szunyogh, A.V. Zimin, E.J. Kostelich, M. Corazza, E. Kalnay, D. Patil, and J.A. Yorke. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A*. vol. 56, no. 5. pp. 415–428 (2004).
- [114] N. Oudjane and S. Rubenthaler. Stability and uniform particle approximation of nonlinear filters in case of non ergodic signals. *Stochastic Analysis and Applications*. vol. 23, no. 3. pp. 421–448 (2005).
- [115] N. Papadakis, E. Mémin, A. Cuzol, and N. Gengembre. Data assimilation with the weighted ensemble Kalman filter. *Tellus A: Dynamic Meteorology and Oceanography*. vol. 62, no. 5. pp. 673–697 (2010).

- [116] P. Park and T. Kailath. Convergence of the DRE solution to the ARE strong solution. *IEEE Transactions on Automatic Control*. vol. 42, no. 4. 573–578 (1997).
- [117] S. Pathiraja, S. Reich, and W. Stannat. McKean-Vlasov SDEs in nonlinear filtering. *SIAM Journal on Control and Optimization*. vol. 59, no. 6. pp. 4188–4215 (2021).
- [118] M-A. Poubelle, I.R. Petersen, M.R. Gevers, and R.R. Bitmead. A Miscellany of Results on an Equation of Count J. F. Riccati. *IEEE Transactions on Automatic Control*. vol. 31, no. 7. pp. 651–654 (1986).
- [119] P. Rebeschini and R. Van Handel. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*. vol. 25, no. 5. pp. 2809–2866 (2015).
- [120] S. Reich. Data assimilation: A dynamic homotopy-based coupling approach. arXiv e-print, [arXiv:2209.05279](https://arxiv.org/abs/2209.05279) (2022).
- [121] S. Reich and C.J. Cotter. Ensemble filter techniques for intermittent data assimilation. *Large Scale Inverse Problems: Computational Methods and Applications in the Earth Sciences* (eds: M. Cullen, M.A. Freitag, S. Kindermann, R. Scheichl), pages 91–134. De Gruyter Publishers (2013). See also: arXiv e-print, [arXiv:1208.6572](https://arxiv.org/abs/1208.6572) (2012).
- [122] S. Reich and C.J. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press (2015)
- [123] S. Reich and S. Shin. On the consistency of ensemble transform filter formulations. *Journal of Computational Dynamics*. vol. 1, no. 1. pp. 177–189 (2014).
- [124] K. Reif, S. Gunther, E. Yaz, and R. Unbehauen. Stochastic stability of the continuous-time extended Kalman filter. *IEE Proceedings – Control Theory and Applications*. vol. 147, no. 1. pp. 45–52 (2000).
- [125] H. Ruzayqat, N.K. Chada, and A. Jasra. Multilevel estimation of normalization constants using ensemble Kalman-Bucy filters. *Statistics and Computing*. vol. 32, no. 3. pp. 1–25 (2022).
- [126] P. Sakov and P.R. Oke. A deterministic formulation of the ensemble Kalman filter: an alternative to ensemble square root filters. *Tellus A*. vol. 60, no. 2. pp. 361–371 (2008).
- [127] P. Sakov and P.R. Oke. Implications of the Form of the Ensemble Transformation in the Ensemble Square Root Filters. *Monthly Weather Review*. vol. 136, no. 3. pp. 1042–1053 (2008).
- [128] C. Schillings and A.M. Stuart. Convergence analysis of ensemble Kalman inversion: the linear, noisy case. *Applicable Analysis*. vol. 97, no. 1. pp. 107–123 (2017).
- [129] C. Schillings and A.M. Stuart. Analysis of the ensemble Kalman filter for inverse problems. *SIAM Journal on Numerical Analysis*. vol. 55, no. 3. pp. 1264–1290 (2017).
- [130] A. Seiler, G. Evensen, J.-A. Skjervheim, J. Hove, and J.G. Vab. Using the EnKF for history matching and uncertainty quantification of complex reservoir models. In *Large-Scale Inverse Problems and Quantification of Uncertainty* (eds: L. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. Mallick, Y. Marzouk, L. Tenorio, B. van Bloemen Waanders and K. Willcox). pp. 247–271. John Wiley & Sons, New York (2010).
- [131] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M.I. Jordan and S.S. Sastry. Kalman Filtering with Intermittent Observations. *IEEE Transactions on Automatic Control*. vol. 49, no. 9. pp. 1453–1464 (2004).

- [132] T. Ström. On logarithmic norms. *SIAM Journal on Numerical Analysis*. vol. 12, no. 5. pp. 741–753 (1975).
- [133] A. Taghvaei and P.G. Mehta. An optimal transport formulation of the linear feedback particle filter. In *Proc. of the 2016 American Control Conference (ACC)*, Boston, USA (July, 2016).
- [134] A. Taghvaei and P.G. Mehta. A Survey of Feedback Particle Filter and related Controlled Interacting Particle Systems (CIPS). *arXiv e-print*, [arXiv:2301.00935](https://arxiv.org/abs/2301.00935) (2023).
- [135] A. Taghvaei, P.G. Mehta, and S.P. Meyn. Diffusion map-based algorithm for gain function approximation in the feedback particle filter. *SIAM/ASA Journal on Uncertainty Quantification*. vol. 8, no. 3. pp. 1090–1117 (2020).
- [136] A. Tanwani and O. Yufereva. Error covariance bounds for suboptimal filters with Lipschitzian drift and Poisson-sampled measurements. *Automatica*. vol. 122. (2020).
- [137] M.K. Tippett, J.L. Anderson, C.H. Bishop, T.M. Hamill, and J.S. Whitaker. Ensemble square root filters. *Monthly Weather Review*. vol. 131, no. 7. pp. 1485–1490 (2003).
- [138] X.T. Tong, A.J. Majda, and D. Kelly. Nonlinear stability and ergodicity of ensemble based Kalman filters. *Nonlinearity*. vol. 29, no. 2. pp 657–691 (2016).
- [139] X.T. Tong, A.J. Majda, and D. Kelly. Nonlinear stability of the ensemble Kalman filter with adaptive covariance inflation. *Communications in Mathematical Sciences*. vol. 14, no. 5. pp. 1283–1313 (2016).
- [140] R. Van Handel. Observability and nonlinear filtering. *Probability Theory and Related Fields*. vol. 145, no. 1-2. pp. 35–74 (2009).
- [141] R. Van Handel. The stability of conditional Markov processes and Markov chains in random environments. *The Annals of Probability*. vol. 37, no. 5. 1876–1925 (2009).
- [142] R. Van Handel. Uniform time average consistency of Monte Carlo particle filters. *Stochastic Processes and their Applications*. vol. 119, no. 11. pp. 3835–3861 (2009).
- [143] J.S. Whitaker and T.M. Hamill. Ensemble data assimilation without perturbed observations. *Monthly Weather Review*. vol. 130, no. 7. pp. 1913–1924 (2002).
- [144] N. Whiteley. Stability properties of some particle filters. *The Annals of Applied Probability*. vol. 23, no. 6. pp. 2500–2537 (2013).
- [145] T. Yang, R.S. Laugesen, P.G. Mehta, and S.P. Meyn. Multivariable feedback particle filter. *Automatica*. vol. 71. pp. 10–23 (2016).
- [146] C. Zhang, A. Taghvaei, and P.G. Mehta. A mean-field optimal control formulation for global optimization. *IEEE Transactions on Automatic Control*. vol. 64, no. 1. pp. 282–289 (2018).