Available online at www.sciencedirect.com

ScienceDirect

Stochastic Processes and their Applications 162 (2023) 481–546

stochastic processes and their applications

www.elsevier.com/locate/spa

Online parameter estimation for the McKean–Vlasov stochastic differential equation

Louis Sharrock[a,b,*], Nikolas Kantas[c], Panos Parpas[c], Grigorios A. Pavliotis[c]

[a] *Department of Mathematics and Statistics, Lancaster University, LA1 4YR, UK*
[b] *School of Mathematics, University of Bristol, BS8 1UG, UK*
[c] *Department of Mathematics, Imperial College London, SW7 2AZ, UK*

Received 21 January 2022; received in revised form 20 April 2023; accepted 3 May 2023
Available online 9 May 2023

Abstract

We analyse the problem of online parameter estimation for a stochastic McKean–Vlasov equation, and the associated system of weakly interacting particles. We propose an online estimator for the parameters of the McKean–Vlasov SDE, or the interacting particle system, which is based on a continuous-time stochastic gradient ascent scheme with respect to the asymptotic log-likelihood of the interacting particle system. We characterise the asymptotic behaviour of this estimator in the limit as $t \to \infty$, and also in the joint limit as $t \to \infty$ and $N \to \infty$. In these two cases, we obtain almost sure or $\mathbb{L}_1$ convergence to the stationary points of a limiting contrast function, under suitable conditions which guarantee ergodicity and uniform-in-time propagation of chaos. We also establish, under the additional condition of global strong concavity, $\mathbb{L}_2$ convergence to the unique maximiser of the asymptotic log-likelihood of the McKean–Vlasov SDE, with an asymptotic convergence rate which depends on the learning rate, the number of observations, and the dimension of the non-linear process. Our theoretical results are supported by two numerical examples, a linear mean field model and a stochastic opinion dynamics model.
© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

MSC: 60F05; 60F25; 60H10; 62F12

Keywords: McKean–Vlasov equation; Maximum likelihood; Parameter estimation; Stochastic gradient descent

1. Introduction

In this paper, we consider a family of McKean–Vlasov stochastic differential equations (SDEs) on $\mathbb{R}^d$, parametrised by $\theta \in \mathbb{R}^p$, of the form

$$\mathrm{d}x_t^\theta = B(\theta, x_t^\theta, \mu_t^\theta)\mathrm{d}t + \sigma(x_t^\theta)\mathrm{d}w_t, \quad t \geq 0, \tag{1}$$

* Corresponding author at: Department of Mathematics and Statistics, Lancaster University, LA1 4YR, UK.
E-mail address: l.sharrock@lancaster.ac.uk (L. Sharrock).

https://doi.org/10.1016/j.spa.2023.05.002
0304-4149/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

where $B : \mathbb{R}^p \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d$, $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ are Borel measurable functions, $(w_t)_{t\geq 0}$ is a $\mathbb{R}^d$-valued standard Brownian motion, and $\mu_t^\theta = \mathcal{L}(x_t^\theta)$ denotes the law of x_t^θ. We assume that $x_0 \in \mathbb{R}^d$, or that x_0 is a $\mathbb{R}^d$-valued random variable with law μ_0, independent of $(w_t)_{t\geq 0}$. This equation is non-linear in the sense of McKean [74,75,98]; in particular, the drift coefficient depends on the law of the solution, in addition to the solution itself. We will restrict our attention to the case in which the dependence on the law only enters linearly in the drift, namely, that

$$B(\theta, x, \mu) = b(\theta, x) + \int_{\mathbb{R}^d} \phi(\theta, x, y)\mu(\mathrm{d}y),$$

for some Borel measurable functions $b : \mathbb{R}^p \times \mathbb{R}^d \to \mathbb{R}^d$ and $\phi : \mathbb{R}^p \times \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$. This choice of dynamics, while not the most general possible, is sufficiently broad for many applications of interest. Moreover, it includes the case in which b and ϕ both have gradient forms, that is, $b(\theta, x) = -\nabla V_\theta(x)$ and $\phi(\theta, x, y) = -\nabla W_\theta(x-y)$, in which case V_θ and W_θ are referred to as the confinement potential and the interaction potential, respectively (e.g., [37,72]).

The McKean–Vlasov SDE arises naturally as the mean field limit ($N \to \infty$) of the weakly interacting particle system (IPS)

$$\mathrm{d}x_t^{\theta,i,N} = B(\theta, x_t^{\theta,i,N}, \mu_t^{\theta,N})\mathrm{d}t + \sigma(x_t^{\theta,i,N})\mathrm{d}w_t^i \ , \quad i = 1, \ldots, N, \tag{2}$$

where $(w_t^i)_{t\geq 0}$ are N independent $\mathbb{R}^d$-valued independent standard Brownian motions, x_0^i are a family of i.i.d. $\mathbb{R}^d$-valued random variables with common law μ_0, independent of $(w_t^i)_{t\geq 0}$, and $\mu_t^{\theta,N} = \frac{1}{N}\sum_{i=1}^N \delta_{x_t^{\theta,i,N}}$ is the empirical law of the interacting particles. In particular, under relatively weak assumptions, it is well known that the empirical law $\mu_t^{\theta,N} \to \mu_t^\theta$ weakly as $N \to \infty$ (e.g., [82]). This phenomenon is commonly referred to as the *propagation of chaos* [98].

The McKean–Vlasov SDE also has a natural connection to a non-linear, non-local partial differential equation on the space of probability measures (e.g., [26]). In particular, under some regularity conditions on b and ϕ, one can show that $\mu_t^\theta = \mathcal{L}(x_t^\theta)$ is absolutely continuous with respect to the Lebesgue measure for all $t \geq 0$ [75,99] and its density, which we will denote by u_t^θ, satisfies a non-linear partial differential equation of the form

$$\begin{aligned}\partial_t u_t^\theta(x) = &-\sum_{k=1}^d \partial_{x_k}\left[\left[b_k(\theta, x) + \int_{\mathbb{R}^d} \phi_k(\theta, x, y)u_t^\theta(y)\mathrm{d}y\right]u_t^\theta(x)\right] \\ &+ \frac{1}{2}\sum_{k,k'=1}^d \partial^2_{x_k,x_{k'}}\left[a_{k,k'}(x)u_t^\theta(x)\right],\end{aligned} \tag{3}$$

where $a(x) = \sigma(x)\sigma^T(x)$. In the particular case that $b(\theta, x) = -\nabla V_\theta(x)$ and $\phi(\theta, x, y) = -\nabla W_\theta(x-y)$, this is commonly referred to as the granular media equation or the McKean–Vlasov–Fokker–Planck equation (e.g., [6,26]).

1.1. Literature review

The systematic study of McKean–Vlasov SDEs was first initiated by McKean [74] in the 1960s, inspired by Kac's programme in Kinetic Theory [56]. We refer to [40,77,98] for some other classical references. In the last two decades, the study of non-linear diffusions has continued to receive considerable attention, with extensive results on well-posedness (e.g., [27,52]), existence and uniqueness (e.g., [5,40,48,55,64,79,103]), ergodicity (e.g., [4,11,18,25,26,33,38,49,72,100]), and propagation of chaos (e.g., [6,17,37,72,73]). This

has no doubt been motivated, at least in part, by the increasing number of applications for McKean–Vlasov SDEs, including in statistical physics [7], multi-agent systems [6], mean-field games [20–22], stochastic control [15], filtering [32], mathematical biology (including neuroscience [2] and structured models of population dynamics [16,80]), epidemic dynamics [3], social sciences (including opinion dynamics [28,46] and cooperative behaviours [19]), financial mathematics [45], Bayesian inference [66], and the analysis of neural networks [51,76,89,95].

Despite the recent renewed interest in the study of McKean–Vlasov SDEs, however, the problem of parameter estimation for this class of equations has received relatively little attention. This is contrast to the wealth of literature on parameter inference in linear (i.e., not measure-dependent) diffusion processes (e.g., [13,59,65]). Recently, Wen et al. [105] established the asymptotic consistency and asymptotic normality of the (offline) maximum likelihood estimator (MLE) for a broad class of McKean–Vlasov SDEs, based on continuous observation of $(x_t)_{t\in[0,T]}$. These results have since been extended by Liu et al. to the path-dependent case [67]. We also mention the work of Catalot and Laredo [42–44], who have studied parametric inference for a particular class of nonlinear self-stabilising SDEs using an approximate log-likelihood function, again based on continuous observation of the non-linear diffusion process, and established asymptotic properties (consistency, normality, convergence rates) of the resulting estimators in several asymptotic regimes (e.g., small noise and long time limit). Several authors have also studied parameter estimation in the corresponding IPS as the number of particles $N \to \infty$ [1,10,29,45,57,71,78,86]. On a related topic, Gomes et al. [47] have considered parameter estimation for a McKean–Vlasov PDE, based on independent realisations of the associated non-linear SDE, in the context of models for pedestrian dynamics.

In a slightly different framework, Maestra and Hoffmann [35] consider non-parametric estimation of the drift-term in a McKean–Vlasov SDE, and the solution of the corresponding non-linear Fokker–Planck equation, based on continuous observation of the associated IPS over a fixed time horizon, namely $(x_t^{i,N})_{t\in[0,T]}^{i=1,\dots,N}$, in the limit as $N \to \infty$. The authors obtain adaptive estimators based on the solution map of the Fokker–Planck equation, and prove their optimality in a minimax sense. Moreover, in the case that $b(x) = -\nabla V(x)$ and $\phi(x, y) = -\nabla W(x - y)$, the authors derive an estimator of the interaction potential, and establish its consistency. We also refer to [31,62,69,70,107] for some other recent contributions on non-parametric inference for IPSs.

Despite these recent contributions, there are no existing works which tackle the problem of online parameter estimation for McKean–Vlasov SDEs. The main purpose of this paper is to address this gap. There is significant motivation for this approach. In comparison to classical (i.e., offline) methods, which process observations in a batch fashion, online methods perform inference in real time, can track changes in parameters over time, are more computationally efficient, and have significantly smaller storage requirements. Even for standard diffusion processes, literature on online parameter estimation is somewhat sparse, with some notable recent exceptions [9,94,96,97]. The problem of recursive estimation in continuous-time stochastic processes was first rigorously analysed by Levanony et al. [63]. These authors propose an online MLE which, irrespective of initial conditions, is consistent and asymptotically efficient. This problem has more recently been revisited by Sirignano and Spiliopoulos [94,96], who propose an online statistical learning algorithm - 'stochastic gradient descent in continuous time' - for the estimation of the parameters in a fully observed ergodic diffusion process. These authors establish the a.s. convergence of this estimator to the stationary points of a suitably defined objective function [94], and, under additional assumptions, also obtain an $\mathbb{L}^p$ convergence rate and a central limit theorem [96]. These results have since also been extended to partially observed diffusion processes [97], and jump–diffusion processes [9].

1.2. Contributions

The main contributions of this paper relate to both the methodology and the theory of online parameter estimation for the McKean–Vlasov SDE (1). Regarding methodology:

- We discuss how one can formulate an appropriate approximation to the true likelihood function in this problem, under various modelling assumptions.
- We distinguish between cases in which the data consists of multiple independent samples from the McKean–Vlasov SDE (Case I), or multiple interacting trajectories from the IPS (Case II).

In each case, we perform a rigorous asymptotic analysis of the online MLE, both in the limit as $t \to \infty$, and in the joint limit as $t \to \infty$ and $N \to \infty$. Our main theoretical contributions can be summarised as follows:

- We propose implementable online estimators for the parameters of the McKean–Vlasov SDE, which evolve according to continuous-time stochastic gradient ascent algorithms with respect to the asymptotic log-likelihood function of the corresponding IPS.
- We establish that these estimators converge, either almost surely (in the limit as $t \to \infty$) or in $\mathbb{L}^1$ (in the limit as $t \to \infty$ and $N \to \infty$) to the stationary points of certain contrast functions. These results hold under assumptions which guarantee ergodicity and, in the double asymptotic framework, uniform-in-time propagation of chaos.
- We establish, under the additional condition of global strong concavity, that these estimators converge in $\mathbb{L}^2$ to the true parameter θ_0, i.e., the unique maximiser of the asymptotic log-likelihood of the McKean–Vlasov SDE. We also provide explicit convergence rates, which depend on the dimension of the non-linear process, the number of observations, and the learning rate.

Finally, we provide numerical examples to illustrate the application of the proposed methodology to two cases of interest, namely, a linear mean-field model and a stochastic opinion dynamics model. It is worth emphasising that, given the connection between the McKean–Vlasov SDE (1) and the non-linear, non-local PDE (3), the results of this paper are also applicable when one is primarily interested in parameter estimation for the non-linear PDE (3).

1.3. Paper organisation

The remainder of this paper is organised as follows. In Section 2, we formulate the estimation problem, and propose a recursive estimator for the McKean–Vlasov SDE. In Section 3, we state our assumptions and our main results. In Section 4, we provide the proofs of these results. In Section 5, we provide several numerical examples illustrating the performance of the proposed algorithm. Finally, in Section 6, we provide some concluding remarks.

1.4. Additional notation

We will assume throughout this paper that solutions $(x_t^\theta)_{t\geq 0}$ of (1) are defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P}_\theta)$, equipped with filtration $(\mathcal{F}_t)_{t\geq 0}$, and denote the corresponding expectation by $\mathbb{E}_\theta$. In addition, if $(x_t^\theta)_{t\geq 0}$ is a solution of (1) with $x_0 = x \in \mathbb{R}^d$, we will occasionally make explicit the dependence on the initial condition by writing $x_t^\theta := x_t^{\theta,x}$,

$\mu_t^\theta := \mu_t^{\theta,x}$, and $\mathbb{E}_\theta := \mathbb{E}_{\theta,x}$. We will use $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$ to denote, respectively, the Euclidean inner product and the corresponding norm on $\mathbb{R}^d$. We write $\mathcal{P}(\mathbb{R}^d)$ for the collection of all probability measures on $\mathbb{R}^d$. Meanwhile, for $p > 0$, we will write $\mathcal{P}_p(\mathbb{R}^d)$ for the collection of all probability measures on $\mathbb{R}^d$ with finite pth moment. In a slight abuse of notation, we will sometimes write $\mu(\|\cdot\|^p)$ for the pth moment of μ; that is, $\mu(\|\cdot\|^p) = \int_{\mathbb{R}^d} \|x\|^p \mu(\mathrm{d}x)$. Finally, for $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, we will write $\mathbb{W}_p(\mu, \nu)$ to denote the Wasserstein distance between μ and ν, viz

$$\mathbb{W}_p(\mu, \nu) = \inf_{\pi \in \Pi(\mu,\nu)} \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p \pi(\mathrm{d}x, \mathrm{d}y) \right]^{\frac{1}{\max\{1,p\}}},$$

where $\Pi(\mu, \nu)$ denotes the set of all couplings of μ and ν. That is, if $\pi \in \Pi(\mu, \nu)$, then $\pi(A \times \mathbb{R}^d) = \mu(A)$ and $\pi(\mathbb{R}^d \times A) = \nu(A)$ for all $A \in \mathcal{B}(\mathbb{R}^d)$.

2. Parameter estimation for the McKean–Vlasov SDE

We will assume, throughout this paper, that there exists a true, static parameter $\theta_0 \in \mathbb{R}^p$ which generates observations $(x_t)_{t\geq 0} := (x_t^{\theta_0})_{t\geq 0}$ of the McKean–Vlasov SDE (1), with corresponding law $(\mu_t)_{t\geq 0} := (\mu_t^{\theta_0})_{t\geq 0}$. Thus, in our notation, we will suppress the dependence of the observed path, and its law, on the true parameter θ_0. We will make the same assumption when instead we observe trajectories of the IPS (2), in which case the observations are given by $(x_t^{i,N})_{t\geq 0}^{i=1,\ldots,N} = (x_t^{\theta_0,i,N})_{t\geq 0}^{i=1,\ldots,N}$, with corresponding empirical law $(\mu_t^N)_{t\geq 0} := (\mu_t^{\theta_0,N})_{t\geq 0}$.

2.1. The likelihood function

Let $\mathbb{P}_\theta^t$ denote the probability measure induced by a path $(x_s^\theta)_{s\in[0,t]}$ of the McKean–Vlasov SDE (1). Under certain regularity conditions, to be specified below, $\mathbb{P}_\theta^t$ is absolutely continuous with respect to $\mathbb{P}_{\theta_0}^t$ for all $\theta \in \mathbb{R}^p$ (e.g. [65, Theorem 7.19], [67,105]). One can thus use the Girsanov formula to obtain a likelihood function as

$$\mathcal{L}_t(\theta) = \log \frac{\mathrm{d}\mathbb{P}_\theta^t}{\mathrm{d}\mathbb{P}_{\theta_0}^t} = \int_0^t \left\langle \left[B(\theta, x_s, \mu_s^\theta) - B(\theta_0, x_s, \mu_s)\right], (\sigma(x_s)\sigma^T(x_s))^{-1} \mathrm{d}x_s \right\rangle - \frac{1}{2}\int_0^t \left[\left\|\sigma^{-1}(x_s)B(\theta, x_s, \mu_s^\theta)\right\|^2 - \left\|\sigma^{-1}(x_s)B(\theta_0, x_s, \mu_s)\right\|^2 \right] \mathrm{d}s.$$

In the case that the diffusion coefficient σ depends on the unknown parameter θ, the measures $\{\mathbb{P}_\theta^t\}$ are, in general, mutually singular (e.g., [60, Section 1.3]). We thus adopt the standard condition of parameter independence for the diffusion coefficient, and for convenience set $\sigma = 1$ (e.g., [13,63,105]). For linear diffusion processes (in the sense of McKean), there are several approaches which can be applied when σ does depend on an unknown parameter, including those based on a quasi log-likelihood function [50,53], or on minimising a least squares type function for the diffusion coefficient [94]. In principle, the methods introduced in this paper can be extended to either of these cases.

In order to proceed, it will be convenient to define the functions $G : \mathbb{R}^p \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d$ and $L : \mathbb{R}^p \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ according to $G(\theta, x, \mu, \mu') := B(\theta, x, \mu) - B(\theta_0, x, \mu')$ and $L(\theta, x, \mu, \mu') := -\frac{1}{2}\|G(\theta, x, \mu, \mu')\|^2$. We can then write the

log-likelihood function as

$$\mathcal{L}_t(\theta) = \int_0^t \underbrace{-\tfrac{1}{2}\left\| B(\theta, x_s, \mu_s^\theta) - B(\theta_0, x_s, \mu_s)\right\|^2}_{:=L(\theta, x_s, \mu_s^\theta, \mu_s)} \mathrm{d}s + \int_0^t \Big\langle \underbrace{B(\theta, x_s, \mu_s^\theta) - B(\theta_0, x_s, \mu_s)}_{:=G(\theta, x_s, \mu_s^\theta, \mu_s)}, \mathrm{d}w_s \Big\rangle. \tag{4}$$

In general, while we may observe a sample path $(x_t)_{t\geq 0}$ of the McKean–Vlasov SDE, typically we will not have direct access to its law $(\mu_t)_{t\geq 0}$. As such, it is generally not possible to compute the likelihood function $\mathcal{L}_t(\theta)$ in (4) directly. On this basis, even if one is interested in fitting data to the McKean–Vlasov SDE, it will typically be necessary to approximate the corresponding likelihood function. In order to make such an approximation, we will assume that we can simultaneously observe multiple continuous sample paths. Indeed, this is much more typical of the data that we observe in practice. Within this framework, there are now two main possibilities, as we outline below.

2.1.1. Case I

The first possibility is to assume the observed paths are N independent instances $(x_t^i)_{t\geq 0}^{i=1,\dots,N}$ of the McKean–Vlasov SDE (1). In this case, we can approximate the true log-likelihood function $\mathcal{L}_t(\theta)$ using

$$\begin{aligned}\mathcal{L}_t^{[N]}(\theta) &:= \frac{1}{N}\sum_{i=1}^N \mathcal{L}_t^{[i,N]}(\theta) \\ &= \frac{1}{N}\sum_{i=1}^N \left[\int_0^t L(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]})\mathrm{d}s + \int_0^t \left\langle G(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]}), \mathrm{d}w_s^i\right\rangle\right],\end{aligned} \tag{5}$$

where $\mu_t^{[N]} = \frac{1}{N}\sum_{i=1}^N \delta_{x_t^i}$ denotes the empirical measure of the sample paths. In this approximation, the functions $\mathcal{L}_t^{[i,N]}(\theta)$, $i = 1, \dots, N$, can be viewed as N 'Monte Carlo esque' approximations of $\mathcal{L}_t(\theta)$, obtained by replacing both μ_t^θ and μ_t in (4) by the empirical law $\mu_t^{[N]}$. The approximation $\mathcal{L}_t^{[N]}(\theta)$ then follows by independence. Alternatively, this approximation can be obtained by substituting the observations of the McKean–Vlasov SDE $(x_t^i)_{t\geq 0}^{i=1,\dots,N}$ into the log-likelihood function of the corresponding IPS (see below).

2.1.2. Case II

The second possibility is to assume that the observed paths correspond to the trajectories of N particles $(x_t^{i,N})_{t\geq 0}^{i=1,\dots,N}$ from the IPS (2). In this case, we can replace $\mathcal{L}_t(\theta)$ by the Girsanov log-likelihood function for the IPS, which is given by (e.g., [10,29,57,71])

$$\begin{aligned}\mathcal{L}_t^N(\theta) &:= \frac{1}{N}\sum_{i=1}^N \mathcal{L}_t^{i,N}(\theta) \\ &= \frac{1}{N}\sum_{i=1}^N \left[\int_0^t L(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N)\mathrm{d}s + \int_0^t \left\langle G(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N), \mathrm{d}w_s^i\right\rangle\right],\end{aligned} \tag{6}$$

where $\mu_t^N = \frac{1}{N}\sum_{j=1}^N \delta_{x_t^{j,N}}$ denotes the empirical measure of the IPS, and we have included $\frac{1}{N}$ as a normalisation factor. The connection between the functions $\mathcal{L}_t^N(\theta)$ and $\mathcal{L}_t^{[N]}(\theta)$ is now abundantly clear. In particular, they are identical as functions of the data.

Table 1
Parameter estimation: Summary of different cases.

Case	Data-generating model	Observation(s)	Likelihood function	
			Approximate	*True*
Case I	MVSDE (1)	$(x_t^i)_{t\geq 0}^{i=1,\dots,N}$	$\mathcal{L}_t^{[N]}(\theta)$ in (5)	$\mathcal{L}_t(\theta)$ in (4)
Case II	IPS (2)	$(x_t^{i,N})_{t\geq 0}^{i=1,\dots,N}$	$\mathcal{L}_t^{N}(\theta)$ in (6)	$\mathcal{L}_t^{N}(\theta)$ in (6)

2.1.3. Discussion

These two cases are summarised in Table 1. Case I is of particular interest if one believes that the true data-generating model is the McKean–Vlasov SDE (1). In particular, the function $\mathcal{L}_t^{[N]}(\theta)$ in (5) should be regarded as an approximation of the true log-likelihood $\mathcal{L}_t(\theta)$ in (4), which can be computed under more realistic modelling assumptions. Meanwhile, Case II is of particular interest if one believes that the true data-generating model is the IPS, for a fixed and finite number of particles. Indeed, in this scenario, the function $\mathcal{L}_t^{N}(\theta)$ in (6) can and should be regarded as the true likelihood function. Nonetheless, it will be interesting to consider the behaviour of this function in the mean-field limit, in order to analyse any differences that arise with Case I.

In the limit as $N \to \infty$, standard propagation-of-chaos results (e.g., [72]) show that the dynamics of the observations in Cases I and II will coincide. In our results, we will establish rigorously that this also holds for the different implied 'likelihood' functions, $\mathcal{L}_t^{N}(\theta)$ and $\mathcal{L}_t^{[N]}(\theta)$. This should not be a surprise given the similarities between these two functions.

2.2. Online parameter estimation

In the online setting, our objective is to estimate the true parameter θ_0 in real time, using the continuous stream of observations. In this case, a natural objective function is the asymptotic (or average) log-likelihood function $\tilde{\mathcal{L}}(\theta)$ of the McKean–Vlasov SDE, viz

$$\tilde{\mathcal{L}}(\theta) = \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t(\theta) \overset{a.s.}{=} \int_{\mathbb{R}^d} L(\theta, x, \mu_\infty^\theta, \mu_\infty)\mu_\infty(\mathrm{d}x), \tag{7}$$

where μ_∞^θ and $\mu_\infty = \mu_\infty^{\theta_0}$ denote the unique invariant measures of $(x_t^\theta)_{t\geq 0}$ and $(x_t)_{t\geq 0} := (x_t^{\theta_0})_{t\geq 0}$, respectively, which exist under appropriate conditions on the McKean–Vlasov SDE (see Section 3). In the spirit of [9,94,97], one can recursively maximise this objective function by defining an estimator $(\theta_t)_{t\geq 0}$ which follows the gradient of the integrand of the log-likelihood $\mathcal{L}_t(\cdot)$ in (4), evaluated with the current parameter estimate. This represents a stochastic estimate for the direction of steepest ascent of the asymptotic log-likelihood $\tilde{\mathcal{L}}(\cdot)$ in (7). In our case, initialised at $\theta_{\text{init}} \in \mathbb{R}^p$, this yields the update equation

$$\mathrm{d}\theta_t = \gamma_t\Big(\underbrace{\nabla_\theta L(\theta_t, x_t, \mu_t^{\theta_t}, \mu_t)\mathrm{d}t}_{\text{(noisy) ascent term}} + \underbrace{\nabla_\theta B(\theta_t, x_t, \mu_t^{\theta_t})\mathrm{d}w_t}_{\text{noise term}}\Big), \tag{8}$$

where $\gamma_t : \mathbb{R}_+ \to \mathbb{R}_+^p$ is a positive, non-increasing function known as the *learning rate* and, in a slight abuse of notation, $(\mu_t^{\theta_t})_{t\geq 0}$ denotes the law of $(x_t^{\theta_t})_{t\geq 0}$, the solution of (1) integrated with the online parameter estimate. We note that one also arrives at this estimator by considering a 'least-squares' type objective for the drift function (see [94]). This evolution equation does

indeed represent a continuous-time stochastic gradient ascent scheme on the asymptotic log-likelihood function. To see this, we can rewrite the parameter update equation (8) in the form

$$\mathrm{d}\theta_t = \gamma_t\Bigg(\underbrace{\nabla_\theta\tilde{\mathcal{L}}(\theta_t)\mathrm{d}t}_{\text{(true) ascent term}} + \underbrace{(\nabla_\theta L(\theta_t, x_t, \mu_t^{\theta_t}, \mu_t) - \nabla_\theta\tilde{\mathcal{L}}(\theta_t))\mathrm{d}t}_{\text{fluctuations term}} + \underbrace{\nabla_\theta B(\theta_t, x_t, \mu_t^{\theta_t})\mathrm{d}w_t}_{\text{noise term}}\Bigg). \quad (9)$$

The first term in this decomposition represents the true ascent direction $\nabla_\theta\tilde{\mathcal{L}}(\theta_t)$, the second term the deviation between the stochastic gradient ascent direction $\nabla_\theta L(\theta_t, x_t, \mu_t^{\theta_t}, \mu_t)$ and the true gradient ascent direction $\nabla_\theta\tilde{\mathcal{L}}(\theta_t)$, while the third term is a zero-mean noise term. Heuristically, we might expect that, provided the learning rate decreases (sufficiently quickly), the true ascent term will dominate the fluctuations term and the noise term for sufficiently large t. If this is the case, we could then reasonably expect θ_t to converge to a local maximum of $\tilde{\mathcal{L}}(\theta)$.

Unfortunately, there are several challenges associated with implementing the 'ideal' online estimator (8). In particular, this estimator depends on $(\mu_t^{\theta_t})_{t\geq 0}$, the law of the solution of the McKean–Vlasov SDE (1), integrated with the online parameter estimate. We will thus seek approximate, easy-to-implement estimators based on the approximations to the log-likelihood given in (5) and (6). Following our earlier discussion, we consider two cases.

2.2.1. Case I

We begin with the case in which the N sample paths $(x_t^i)_{t\geq 0}^{i=1,\dots,N}$ correspond to independent replicates of the McKean–Vlasov SDE (1). In this case, it is natural to consider an estimator $(\theta_t^{[i,N]})_{t\geq 0}$ which evolves according to the gradient of the integrand of the function $\mathcal{L}_t^{[i,N]}(\theta)$ in (5). This yields the update equation

$$\mathrm{d}\theta_t^{[i,N]} = \gamma_t\left[\nabla_\theta L(\theta_t^{[i,N]}, x_t^i, \mu_t^{[N]}, \mu_t^{[N]})\mathrm{d}t + \nabla_\theta B(\theta_t^{[i,N]}, x_t^i, \mu_t^{[N]})\mathrm{d}w_t^i\right]. \quad (10)$$

Alternatively, one could average over all of the trajectories, which is equivalent to defining an estimator $(\theta_t^{[N]})_{t\geq 0}$ which evolves according to the gradient of the integrand of the function $\mathcal{L}_t^{[N]}(\theta)$ in (5), viz

$$\mathrm{d}\theta_t^{[N]} = \gamma_t\frac{1}{N}\sum_{i=1}^{N}\left[\nabla_\theta L(\theta_t^{[N]}, x_t^i, \mu_t^{[N]}, \mu_t^{[N]})\mathrm{d}t + \nabla_\theta B(\theta_t^{[N]}, x_t^i, \mu_t^{[N]})\mathrm{d}w_t^i\right]. \quad (11)$$

The advantage of (10) is that the computation can be performed locally at each particle, following a message passing step for retrieving the empirical law $\mu_t^{[N]}$. It is thus convenient for a distributed implementation. On the other hand, (11) will typically be more accurate, as we will later demonstrate (see, in particular, the bounds in Theorem 1.2*). These two schemes can be seen as stochastic gradient ascent algorithms for maximising the average of the approximate 'likelihood' function of the ith realisation of the McKean–Vlasov SDE or the approximate 'likelihood' function of N realisations of the McKean–Vlasov SDE, as defined by (5). That is, for maximising

$$\tilde{\mathcal{L}}^{[i,N]}(\theta) = \lim_{t\to\infty}\frac{1}{t}\mathcal{L}_t^{[i,N]}(\theta) \quad \text{or} \quad \tilde{\mathcal{L}}^{[N]}(\theta) = \lim_{t\to\infty}\frac{1}{t}\mathcal{L}_t^{[N]}(\theta). \quad (12)$$

Under suitable conditions on the learning rate, we can expect, based on a decomposition analogous to (9), that the online estimates defined in (10) and (11) will converge to local maxima of $\tilde{\mathcal{L}}^{[i,N]}(\theta)$ and $\tilde{\mathcal{L}}^{[N]}(\theta)$ as $t\to\infty$. In Theorems 1.1 and 1.1*, we show rigorously

that this is indeed the case. Meanwhile, by the law of large numbers, we expect that $\tilde{\mathcal{L}}^{[i,N]}(\theta)$ and $\tilde{\mathcal{L}}^{[N]}(\theta)$ will converge to a limiting contrast function $\underline{\tilde{\mathcal{L}}}(\theta)$ as $N \to \infty$. In the joint limit as $t \to \infty$ and $N \to \infty$ it thus seems reasonable to hypothesise that the estimates (10) and (11) will in fact converge to stationary points of $\underline{\tilde{\mathcal{L}}}(\theta)$. In Theorem 1.1*, we establish rigorously that this is the case, and provide an explicit characterisation of the limiting contrast function as

$$\underline{\tilde{\mathcal{L}}}(\theta) := \lim_{t,N\to\infty} \frac{1}{t}\mathcal{L}_t^{[N]}(\theta) = \int_{\mathbb{R}^d} L(\theta, x, \mu_\infty, \mu_\infty)\mu_\infty(\mathrm{d}x). \tag{13}$$

It is worth noting that this function is distinct from $\tilde{\mathcal{L}}(\theta)$, the asymptotic log-likelihood of the McKean–Vlasov SDE, as defined in (7). Importantly, however, these functions are both maximised at the true parameter $\theta = \theta_0$.

2.2.2. Case II

We now consider the case in which the observations $(x_t^{i,N})_{t\geq 0}^{i=1,\dots,N}$ correspond to the trajectories of the IPS (2). In this case, we can use essentially the same update equations. This should not be surprising, given that the 'approximations' of the log-likelihood of the McKean–Vlasov SDE given in (5) and (6) were identical up to specification of the data. To be explicit, we now have

$$\mathrm{d}\theta_t^{i,N} = \gamma_t \left[\nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N, \mu_t^N)\mathrm{d}t + \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\right] \tag{14}$$

and

$$\mathrm{d}\theta_t^N = \gamma_t \frac{1}{N}\sum_{i=1}^N \left[\nabla_\theta L(\theta_t^N, x_t^{i,N}, \mu_t^N, \mu_t^N)\mathrm{d}t + \nabla_\theta B(\theta_t^N, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\right]. \tag{15}$$

These two schemes now represent stochastic gradient ascent algorithms for maximising the average log-likelihood of the ith particle in the IPS, or the average log-likelihood of all of the particles in the IPS, respectively. That is, for maximising the functions

$$\tilde{\mathcal{L}}^{i,N}(\theta) = \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t^{i,N}(\theta) \quad \text{or} \quad \tilde{\mathcal{L}}^{N}(\theta) = \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t^{N}(\theta). \tag{16}$$

Based on similar arguments to before, we expect that the estimates defined in (14) and (15) will converge to local maxima of $\tilde{\mathcal{L}}^{i,N}(\theta)$ and $\tilde{\mathcal{L}}^{N}(\theta)$ as $t \to \infty$. We establish this rigorously in Theorems 2.1–2.2. Meanwhile, due to uniform-in-time propagation of chaos, we also expect $\tilde{\mathcal{L}}^{i,N}(\theta)$ and $\tilde{\mathcal{L}}^{N}(\theta)$ to converge to $\underline{\tilde{\mathcal{L}}}(\theta)$ as $N \to \infty$, the same limiting function as obtained in Case I. Thus, in the long-time and large-particle regime, we should obtain convergence results very similar to those obtained before. This is the subject of Theorems 2.1*–2.2*.

3. Main results

We are now ready to state our assumptions. We begin with the following integrability assumption on the initial condition.

Assumption A.1. The initial law satisfies $\mu_0 \in \mathcal{P}_k(\mathbb{R}^d)$ for all $k \in \mathbb{N}$.

This condition, alongside the conditions below, guarantees that the solutions of the McKean–Vlasov SDE and the IPS have bounded moments of all orders (see Proposition 1), and so do their invariant measures (see Proposition 2).

Regarding the drift function $B(\theta_0, \cdot, \cdot) : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d$ we will work under the following rather classical assumptions.

Assumption B.1. The function $b(\theta_0, \cdot) : \mathbb{R}^d \to \mathbb{R}^d$ has the following properties.

(i) $b(\theta_0, \cdot)$ is locally Lipschitz. That is, for all $x, x' \in \mathbb{R}^d$ such that $\|x\|, \|x'\| \leq R$, there exists $0 < L_{\theta_0,1} < \infty$ such that

$$\|b(\theta_0, x) - b(\theta_0, x')\| \leq L_{\theta_0,1}\|x - x'\|.$$

(ii) $b(\theta_0, \cdot)$ is 'monotonic'. That is, for all $x, x' \in \mathbb{R}^d$, there exists $\alpha_{\theta_0} > 0$ such that

$$\langle x - x', b(\theta_0, x) - b(\theta_0, x')\rangle \leq -\alpha_{\theta_0}\|x - x'\|^2.$$

Assumption B.2. The function $\phi(\theta_0, \cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ has the following property.

(i) $\phi(\theta_0, \cdot, \cdot)$ is globally Lipschitz. In particular, there exists $0 < 2L_{\theta_0,2} < \alpha_{\theta_0}$ such that, for all $x, y, x', y' \in \mathbb{R}^d$,

$$\|\phi(\theta_0, x, y) - \phi(\theta_0, x', y')\| \leq L_{\theta_0,2}(\|x - x'\| + \|y - y'\|).$$

These two conditions are used to establish existence and uniqueness of the strong solution to the McKean–Vlasov SDE (e.g., [77, Theorem 2.2]), uniform moment bounds, uniform-in-time propagation of chaos, and the existence of, and exponential convergence to, a unique invariant measure (see Appendix A). In the literature on non-linear diffusions, it is typical to consider the case in which $b(\theta, x) = -\nabla V_\theta(x)$ for some confinement potential V, and $\phi(\theta, x, y) = -\nabla W_\theta(x - y)$ for some interaction potential W. In this case, classical conditions which can be used in place of Conditions B.1–B.2 (e.g., convexity) may be found in [72].

Let us remark briefly upon some weaker conditions under which these results still hold, and therefore under which the main results of our paper will also still hold (albeit with some additional technical overhead). In the case that there is no confinement potential (i.e., $V \equiv 0$), and the interaction potential is uniformly convex with gradient that is locally Lipschitz with polynomial growth, Malrieu established uniform-in-time propagation of chaos and exponential convergence to equilibrium [73]. Cattiaux et al. [26] later established the same results in the case that the interaction potential is degenerately convex. Meanwhile, in [24,25], the authors establish exponential convergence to equilibrium under the strict convexity condition $\text{Hess}(V + 2W) \geq \beta I_d$, for some $\beta > 0$.

In the case that $V + 2W$ is not convex, far fewer results are available; indeed, without additional conditions on V and W, even the existence of a unique stationary distribution is not guaranteed (see, e.g., [49,101]). This being said, Bolley et al. [11] proved uniform exponential convergence to equilibrium in both degenerately convex, and weakly non-convex cases. More recently, [37,38] have established uniform-in-time propagation of chaos and exponential convergence to equilibrium in the non-convex case, provided the confinement potential V is strictly convex outside a ball, and the interaction potential is globally Lipschitz with sufficiently small Lipschitz constant. For some other relevant recent contributions, see also [23,34,61,68].

In addition to Assumptions B.1–B.2, we will also impose the following regularity conditions.

Assumption C.1. The functions $b : \mathbb{R}^p \times \mathbb{R}^d \to \mathbb{R}^d$ and $\phi : \mathbb{R}^p \times \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ have the following properties.

(i) $b(\theta, \cdot) \in C(\mathbb{R}^d)$, $\phi(\theta, \cdot, \cdot) \in C(\mathbb{R}^d, \mathbb{R}^d)$ for all $\theta \in \mathbb{R}^p$; and $b(\cdot, x)$ and $\phi(\cdot, x, y) \in \mathcal{C}^3(\mathbb{R}^p)$ for all $x, y \in \mathbb{R}^d$.

(ii) For $k = 0, 1, 2, 3$, there exist finite constants $q, K < \infty$ such that, for all $\theta \in \mathbb{R}^p$ and for all $x, y \in \mathbb{R}^d$,

$$\left\| \nabla_\theta^k b(\theta, x) - \nabla_\theta^k b(\theta, x') \right\| \leq K[\|x - x'\|][1 + \|x\|^q + \|x'\|^q] \tag{17}$$

$$\left\| \nabla_\theta^k \phi(\theta, x, y) - \nabla_\theta^k \phi(\theta, x', y') \right\| \leq K[\|x - x'\| + \|y - y'\|][1 + \|x\|^q + \|x'\|^q + \|y\|^q + \|y'\|^q]. \tag{18}$$

These conditions, variations of which have also appeared in, e.g., [9,94,96], are required in order to control the ergodic behaviour of the solution of the McKean–Vlasov SDE, and the associated IPS, which is central to establishing convergence of the online MLE. In particular, they ensure that fluctuation terms of the form $\int_0^t \gamma_s(\nabla_\theta L(\theta_s, x_s, \mu_s^\theta, \mu_s) - \nabla_\theta \tilde{\mathcal{L}}(\theta_s))\mathrm{d}s$, as in (9), tend to zero sufficiently quickly as $t \to \infty$. Using an approach which is now well established in the analysis of stochastic approximation algorithms in continuous time [9,90–92,94,96,97], we control such terms by rewriting them in terms of the solutions of some related Poisson equations. Condition C.1 ensures that these solutions are unique, and that they grow at most polynomially in a suitable sense (see Lemma 17 in Appendix B). In practice, the assumption that there exist universal constants q, K such that $b(\theta, x)$ and $\phi(\theta, x, y)$ satisfy (17)–(18) can be a little restrictive. However, as in [96], it is possible to establish our results under a weaker version of Condition C.1(ii), in which $b(\theta, x)$ and $\phi(\theta, x, y)$ are allowed to grow linearly with respect to θ, without significant complication.

We now proceed with some additional assumptions which will be required in order to establish our $\mathbb{L}^2$ convergence results (Theorems 1.2 and 1.2* and Theorems 2.2 and 2.2*).

Assumption D.1. There exist a positive constant $R < \infty$, and an almost everywhere positive function $\kappa : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, such that, for all $\|\theta\| \geq R$,

$$\langle \nabla_\theta L(\theta, x, \mu, \mu'), \theta \rangle \leq -\kappa(x, \mu, \mu')\|\theta\|^2.$$

Assumption D.2. Define the function $\tau : \mathbb{R}^p \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ according to

$$\tau(\theta, x, \mu) = \left\langle \nabla_\theta B(\theta, x, \mu) \nabla_\theta B^T(\theta, x, \mu) \frac{\theta}{\|\theta\|}, \frac{\theta}{\|\theta\|} \right\rangle^{\frac{1}{2}}.$$

Then, there exists $0 < q, K < \infty$ such that, for all $\theta, \theta' \in \mathbb{R}^p$, for all $x, y \in \mathbb{R}^d$,

$$|\tau(\theta, x, \mu) - \tau(\theta', x, \mu)|| \leq K\|\theta - \theta'\|(1 + \|x\|^q + \|\mu(\|\cdot\|^2)\|^{\frac{q}{2}}).$$

These two conditions ensure, via the comparison theorem (e.g., [54,106]), that the online parameter estimates generated by the McKean–Vlasov SDE and the IPS, namely $(\theta_t)_{t \geq 0}$ and $(\theta_t^{i,N})_{t \geq 0}$, have uniformly bounded moments (see Lemma 18). We refer to [58] for some more general conditions under which this result still holds. The first condition relates to the drift terms in the two parameter update equations, and can be seen as a recurrence condition; the second condition relates to the diffusion terms, and can be seen as an extension of Condition B.2. This condition was introduced in [94], and has since also appeared in [9]. Instead of these conditions, one could instead use a projection or truncation device to ensure that the online parameter estimates remain bounded (e.g., [30]). This is fairly common within the stochastic approximation literature (e.g., [12]).

In order to establish consistency, we will also require strong concavity for at least one of the asymptotic objective functions.

Assumption E.1. The function $\tilde{\mathcal{L}}^{[i,N]}(\theta)$, or equivalently $\tilde{\mathcal{L}}^{i,N}(\theta)$, is strongly concave with constant $\eta > 0$, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$.

Assumption E.2. The function $\underline{\tilde{\mathcal{L}}}(\theta)$ is strongly concave with constant $\eta > 0$.

The first condition implies that $\tilde{\mathcal{L}}^{i,N}(\theta)$ and $\tilde{\mathcal{L}}^{[i,N]}(\theta)$ have unique maximisers for each fixed $N \in \mathbb{N}$, and will be used when we consider asymptotics in the long-time regime, for a fixed number of particles. The second condition implies that $\underline{\tilde{\mathcal{L}}}(\theta)$ has a unique maximiser, and will be relevant when we consider the double asymptotic framework as $t \to \infty$ and $N \to \infty$. It is worth emphasising that neither Assumption E.1, nor Assumption E.2 is redundant. In particular, one of these conditions may hold, while the other does not. For example, in the linear mean field model studied in Section 5.1, we will see that Assumption E.1 is satisfied, while Assumption E.2 is not (see Appendix C).

It remains only to specify our conditions on the learning rate.

Assumption F.1. The learning rate $\gamma_t : \mathbb{R}_+ \to \mathbb{R}_+$ is a positive, non-increasing function such that $\int_0^\infty \gamma_t \mathrm{d}t = \infty$, $\int_0^\infty \gamma_t^2 \mathrm{d}t < \infty$, $\int_0^\infty \gamma_t' \mathrm{d}t < \infty$. Moreover, there exists $p > 0$ such that $\lim_{t\to\infty} \gamma_t^2 t^{2p+\frac{1}{2}} = 0$.

Assumption F.2. Let $\Phi_{s,t} = \exp(-2\eta \int_s^t \gamma_u \mathrm{d}u)$, for the constant η defined in Condition E.2. The learning rate $\gamma_t : \mathbb{R}_+ \to \mathbb{R}_+$ satisfies $\int_0^t \gamma_s^2 \Phi_{s,t} \mathrm{d}s = O(\gamma_t)$, $\int_0^t \gamma_t' \Phi_{s,t} \mathrm{d}s = O(\gamma_t)$, $\int_0^t \gamma_s \Phi_{s,t} \mathrm{d}s = O(1)$, $\int_0^t \gamma_s \Phi_{s,t} e^{-\lambda s} \mathrm{d}s = O(\gamma_t)$, and $\Phi_{1,t} = O(\gamma_t)$.

The first of these conditions represents the continuous-time analogue of the standard step-size condition used in the convergence analysis of stochastic approximation algorithms in discrete time (e.g., [88,94]). The second condition, first introduced in [96], is specific to stochastic gradient descent in continuous time, and is only required in order to establish our $\mathbb{L}^2$ convergence results (Theorems 1.2 and 1.2* and Theorems 2.2 and 2.2*). A standard choice of learning rate which satisfies both of these conditions is $\gamma_t = C_\gamma (C_{\theta,0} + t)^{-1}$, where $C_\gamma, C_{\theta,0} > 0$ are positive constants such that $C_\gamma \eta > 1$.

3.1. Main results

We are now ready to state our main results. These results are summarised in Table 2.

3.1.1. Case I

We begin, once more, with the case in which we observe N independent sample paths $(x_t^i)_{t\geq 0}^{i=1,\ldots,N}$ from the McKean–Vlasov SDE (1). We thus generate online parameter estimates according to (10) or (11). We will consider two asymptotic regimes.

(i) $t \to \infty$ and $N \in \mathbb{N}$. We begin by considering asymptotics only in the long-time regime, with the number of sample paths fixed and finite. In this case, we can establish a.s. convergence of (10) and (11) to the stationary points of the contrast functions $\tilde{\mathcal{L}}^{[i,N]}(\theta)$ and $\tilde{\mathcal{L}}^{[N]}(\theta)$, respectively. Under additional assumptions on these functions, including strong concavity (Condition E.1), we also obtain $\mathbb{L}_2$ convergence to their unique maximisers, with a convergence rate which depends on the learning rate $(\gamma_t)_{t\geq 0}$.

Table 2
Online parameter estimation: Summary of main results.

Case	Limit(s)	Theorems	Parameter estimates	Objective function	Convergence rate
Case I	$t \to \infty,\ N \in \mathbb{N}$	1.1–1.2	$\theta_t^{[i,N]}$ from (10)	$\tilde{\mathcal{L}}^{[i,N]}(\theta)$	$(K_{\theta_0,1} + K_{\theta_0,2})\gamma_t$
			$\theta_t^{[N]}$ from (11)	$\tilde{\mathcal{L}}^{[N]}(\theta)$	$(K_{\theta_0,1} + \frac{K_{\theta_0,2}}{N})\gamma_t$
Case I	$t \to \infty,\ N \to \infty$	1.1*–1.2*	$\theta_t^{[i,N]}$ from (10)	$\underline{\tilde{\mathcal{L}}}(\theta)$	$\left(K_{\theta_0,1} + K_{\theta_0,2}\right)\gamma_t + K_{\theta_0,3}\alpha(N)$
			$\theta_t^{[N]}$ from (11)		$(K_{\theta_0,1} + \frac{K_{\theta_0,2}}{N})\gamma_t + K_{\theta_0,3}\alpha(N)$
Case II	$t \to \infty,\ N \in \mathbb{N}$	2.1–2.2	$\theta_t^{i,N}$ from (14)	$\tilde{\mathcal{L}}^{i,N}(\theta)$	$(K^\dagger_{\theta_0,1} + K^\dagger_{\theta_0,2})\gamma_t$
			θ_t^{N} from (15)	$\tilde{\mathcal{L}}^{N}(\theta)$	$(K^\dagger_{\theta_0,1} + \frac{K^\dagger_{\theta_0,2}}{N})\gamma_t$
Case II	$t \to \infty,\ N \to \infty$	2.1*–2.2*	$\theta_t^{i,N}$ from (14)	$\underline{\tilde{\mathcal{L}}}(\theta)$	$(K^\dagger_{\theta_0,1} + K^\dagger_{\theta_0,2})\gamma_t + K^\dagger_{\theta_0,3}\alpha(N) + \frac{K^\dagger_{\theta_0,4}}{N^{1/2}}$
			θ_t^{N} from (15)		$(K^\dagger_{\theta_0,1} + \frac{K^\dagger_{\theta_0,2}}{N})\gamma_t + K^\dagger_{\theta_0,3}\alpha(N) + \frac{K^\dagger_{\theta_0,4}}{N^{1/2}}$

Theorem 1.1. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$, it holds a.s. that*

$$\lim_{t\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}^{[i,N]}(\theta_t^{[i,N]})\| = \lim_{t\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}^{[N]}(\theta_t^{[N]})\| = 0.$$

Theorem 1.2. *Assume that Conditions A.1, B.1–B.2, C.1, D.1–D.2, E.1, and F.1–F.2 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$, and for sufficiently large t, there exist positive constants $K_{\theta_0,1}, K_{\theta_0,2}$, such that*

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{[i,N]} - \theta_0\|^2\right] \leq \left(K_{\theta_0,1} + K_{\theta_0,2}\right)\gamma_t,$$
$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{[N]} - \theta_0\|^2\right] \leq (K_{\theta_0,1} + \frac{K_{\theta_0,2}}{N})\gamma_t.$$

(ii) $t \to \infty$ and $N \to \infty$. We next consider joint asymptotics in the many observations and the long-time regime. In this case, we establish $\mathbb{L}^1$ convergence of (10) and (11) to the stationary points of $\underline{\tilde{\mathcal{L}}}(\theta)$, the contrast function defined in (13). Under the additional assumption of strong log-concavity for this contrast function (Condition E.2), among other technical assumptions, we then obtain convergence to the unique maximiser of $\tilde{\mathcal{L}}(\theta)$, the asymptotic log-likelihood of the McKean–Vlasov SDE, with a convergence rate which depends on the learning rate $(\gamma_t)_{t\geq 0}$, the number of sample paths N, and the dimension of the non-linear process d.

It is worth emphasising that, even though we now take $N \to \infty$, the first of these two results does not guarantee convergence to the stationary points of $\tilde{\mathcal{L}}(\theta)$, the true asymptotic log-likelihood function of the McKean–Vlasov SDE. This should be seen as the penalty for replacing $(\mu_t^{\theta_t})_{t\geq 0}$ in the exact (but intractable) update equation (8), by $(\mu_t^{[N]})_{t\geq 0}$ in the approximate (but implementable) update equations (10) and (11). In particular, the former depends on the path of the online parameter estimates $(\theta_t)_{t\geq 0}$, while the latter depends only on the true data-generating parameter θ_0.

Theorem 1.1*. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Then, in $\mathbb{L}^1$, it holds that*

$$\lim_{t,N\to\infty} \|\nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^{[i,N]})\| = \lim_{t,N\to\infty} \|\nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^{[N]})\| = 0.$$

Suppose, in addition, that $\Theta_0 = \{\theta \in \mathbb{R}^p : \nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta) = 0\} = \{\theta_0\}$. Then, in $\mathbb{L}^1$, it also holds that

$$\lim_{t,N\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^{[i,N]})\| = \lim_{t,N\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^{[N]})\| = 0.$$

Theorem 1.2*. *Assume that Conditions A.1, B.1–B.2, C.1, D.1–D.2, E.2, and F.1–F.2 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \dots, N$, and for sufficiently large t, there exist positive constants $K_{\theta_0,1}$, $K_{\theta_0,2}$, $K_{\theta_0,3}$, such that*

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{[i,N]} - \theta_0\|^2\right] \leq \left(K_{\theta_0,1} + K_{\theta_0,2}\right)\gamma_t + K_{\theta_0,3}\alpha(N),$$
$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{[N]} - \theta_0\|^2\right] \leq (K_{\theta_0,1} + \frac{K_{\theta_0,2}}{N})\gamma_t + K_{\theta_0,3}\alpha(N),$$

where $\alpha : \mathbb{N} \to \mathbb{R}_+$ is a function defined according to

$$\alpha(N) = \begin{cases} N^{-\frac{1}{4}} & \text{if} \quad d = 1, \\ N^{-\frac{1}{4}} \log(1+N)^{\frac{1}{2}} & \text{if} \quad d = 2, \\ N^{-\frac{1}{2d}} & \text{if} \quad d \geq 3. \end{cases} \tag{19}$$

3.1.2. Case II

We now turn our attention to the case in which we observe N particles $(x_t^i)_{t\geq 0}^{i=1,\dots,N}$ from the IPS (2), and generate online parameter estimates according to (14) or (15). Once again, we will consider two asymptotic regimes.

(i) $N \in \mathbb{N}$ and $t \to \infty$. We first consider asymptotics in the long-time regime, for a fixed number of particles. This is perhaps the most natural asymptotic regime for Case II, in which implicitly we have assumed that the true data-generating model is a finite system of interacting particles. Thus, arguably, our objective should simply be to maximise the 'partial' asymptotic log-likelihood of the IPS, $\tilde{\mathcal{L}}^{i,N}(\theta)$, or the 'complete' asymptotic log-likelihood of the IPS, $\tilde{\mathcal{L}}^{N}(\theta)$, for finite N, since this corresponds to the true log-likelihood of the data-generating process.
We begin by establishing a.s. convergence of (14) or (15) to the stationary points of these two functions. Assuming also strong concavity for these asymptotic log-likelihood functions (Condition E.1), we then obtain $\mathbb{L}_2$ convergence to their unique maximisers, with an $\mathbb{L}_2$ convergence rate which depends on the learning rate $(\gamma_t)_{t\geq 0}$.

Theorem 2.1. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \dots, N$, it holds a.s. that*

$$\lim_{t\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| = \lim_{t\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}^{N}(\theta_t^{N})\| = 0.$$

Theorem 2.2. *Assume that Conditions A.1, B.1–B.2, C.1, D.1–D.2, E.1, and F.1–F.2 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \dots, N$, and for sufficiently large t, there exist positive constants $K^\dagger_{\theta_0,1}$, $K^\dagger_{\theta_0,2}$, such that c*

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{i,N} - \theta_0\|^2\right] \leq (K^\dagger_{\theta_0,1} + K^\dagger_{\theta_0,2})\gamma_t,$$

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^N - \theta_0\|^2\right] \leq (K^\dagger_{\theta_0,1} + \frac{K^\dagger_{\theta_0,2}}{N})\gamma_t.$$

(ii) $N \to \infty$ and $t \to \infty$. Finally, for completeness, we consider joint asymptotics in the large particle and the long-time regime. Given that the assumed form of the data-generating mechanism in this case (the IPS) and in the previous case (the McKean–Vlasov SDE) coincide as $N \to \infty$ and $t \to \infty$, it is natural to analyse whether the same is true for the resulting parameter estimates.
Similar to before, we first establish $\mathbb{L}^1$ convergence of (14) and (15) to the stationary points of the contrast function $\underline{\tilde{\mathcal{L}}}(\theta)$. Under the additional assumption of strong concavity for this contrast function (Condition E.2), we then obtain $\mathbb{L}^2$ convergence to the unique maximiser of $\tilde{\mathcal{L}}(\theta)$, with a convergence rate which once again depends on the learning rate $(\gamma_t)_{t\geq 0}$, the number of particles N, and the dimension of each of the particles d.

Theorem 2.1*. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Then, in $\mathbb{L}^1$, it holds that*

$$\lim_{t,N\to\infty} \|\nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^{i,N})\| = \lim_{t,N\to\infty} \|\nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^{N})\| = 0.$$

Suppose, in addition, that $\Theta_0 = \{\theta \in \mathbb{R}^p : \nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta) = 0\} = \{\theta_0\}$. Then, in $\mathbb{L}^1$, it also holds that

$$\lim_{t,N\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^{i,N})\| = \lim_{t,N\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^{N})\| = 0.$$

Theorem 2.2*. *Assume that Conditions A.1, B.1–B.2, C.1, D.1–D.2, E.2, and F.1–F.2 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$, and for sufficiently large t, there exist positive constants $K^\dagger_{\theta_0,1}$, $K^\dagger_{\theta_0,2}$, $K^\dagger_{\theta_0,3}$, $K^\dagger_{\theta_0,4}$, such that*

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{i,N} - \theta_0\|^2\right] \leq (K^\dagger_{\theta_0,1} + K^\dagger_{\theta_0,2})\gamma_t + K^\dagger_{\theta_0,3}\alpha(N) + \frac{K^\dagger_{\theta_0,4}}{N^{\frac{1}{2}}},$$

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{N} - \theta_0\|^2\right] \leq (K^\dagger_{\theta_0,1} + \frac{K^\dagger_{\theta_0,2}}{N})\gamma_t + K^\dagger_{\theta_0,3}\alpha(N) + \frac{K^\dagger_{\theta_0,4}}{N^{\frac{1}{2}}},$$

where $\alpha : \mathbb{N} \to \mathbb{R}_+$ is the function defined in (19).

3.1.3. Discussion

Let us briefly compare the results obtained in Case I (Theorems 1.1–1.2 and Theorems 1.1*–1.2*) and Case II (Theorems 2.1–2.2 and Theorems 2.1*–2.2*). As noted previously, the online parameter estimates in both of these cases follow the same update equations; the only difference is in the assumed form of the data-generating model. We thus expect that the results obtained in these two cases will be similar, if not identical. In Theorems 1.1*–1.2* and Theorems 2.1*–2.2*, this is indeed seen to be the case.

In particular, Theorems 1.1* and 2.1* establish that, regardless of the data-generating mechanism, the online parameter estimates generated via (10)–(11) and (14)–(15) converge to the stationary points of $\underline{\tilde{\mathcal{L}}}(\theta)$ as $N \to \infty$ and $t \to \infty$. Meanwhile, Theorems 1.2* and 2.2* establish $\mathbb{L}^2$ convergence to the true parameter θ_0. It is worth noting that there is a difference between the rates established in Theorems 1.2* and 2.2*. In particular, in Case II (Theorem 2.2*) there is an additional $O(N^{-\frac{1}{2}})$ term. We can interpret this term as a penalty for the mismatch between the asymptotic likelihood function implied by the data-generating model in Case II, namely the IPS (2), and the asymptotic likelihood function of the limiting McKean–Vlasov SDE (1).

4. Proof of main results

In this section, we provide proofs of our main results. We will focus on Theorems 2.1* and 2.2*, whose proofs can easily be adapted to obtain the other results.

4.1. Additional notation

As in Section 2.1, we will write $(x_t^i)_{t\geq 0}$ for a solution of the McKean–Vlasov SDE (1), where the Brownian motion $(w_t)_{t\geq 0}$ is replaced by $(w_t^i)_{t\geq 0}$. We will still write $(\mu_t)_{t\geq 0}$ for the law of this solution. We also now write

$$\mathcal{L}_t^i(\theta) = \int_0^t L(\theta, x_s^i, \mu_s^\theta, \mu_s)\mathrm{d}s + \int_0^t \langle G(\theta, x_s^i, \mu_s^\theta, \mu_s), \mathrm{d}w_s^i\rangle,$$
$$\underline{\mathcal{L}}_t^i(\theta) = \int_0^t L(\theta, x_s^i, \mu_s, \mu_s)\mathrm{d}s + \int_0^t \langle G(\theta, x_s^i, \mu_s, \mu_s), \mathrm{d}w_s^i\rangle. \tag{20}$$

In addition, we now define $\hat{x}_t^N = (x_t^{1,N}, \ldots, x_t^{N,N})^T \in (\mathbb{R}^d)^N$, the process consisting of the concatenation of the N solutions of the IPS (2). This process is the solution of the following SDE on $(\mathbb{R}^d)^N$

$$\mathrm{d}\hat{x}_t^N = \hat{B}(\theta, \hat{x}_t^N)\mathrm{d}t + \mathrm{d}\hat{w}_t^N, \tag{21}$$

where $\hat{w}_t^N$ is a $(\mathbb{R}^d)^N$-valued Brownian motion, and the function $\hat{B}(\theta, \cdot) : (\mathbb{R}^d)^N \to (\mathbb{R}^d)^N$ is of the form $\hat{B}(\theta, \hat{x}^N) = (\hat{B}^{1,N}(\theta, \hat{x}^N), \ldots, \hat{B}^{N,N}(\theta, \hat{x}^N))^T$, where, for $i = 1, \ldots, N$, $\hat{B}^{i,N}(\theta, \cdot) : (\mathbb{R}^d)^N \to \mathbb{R}^d$ is defined according to

$$\hat{B}^{i,N}(\theta, \hat{x}^N) = b(\theta, x^{i,N}) + \frac{1}{N}\sum_{j=1}^N \phi(\theta, x^{i,N}, x^{j,N}).$$

It will also be useful to define the functions $\hat{G}^{i,N}(\theta, \cdot) : (\mathbb{R}^d)^N \to \mathbb{R}^d$ and $\hat{L}^{i,N}(\theta, \cdot) : (\mathbb{R}^d)^N \to \mathbb{R}$ according to $\hat{G}^{i,N}(\theta, \hat{x}^N) = \hat{B}^{i,N}(\theta, \hat{x}^N) - \hat{B}^{i,N}(\theta_0, \hat{x}^N)$ and $\hat{L}^{i,N}(\theta, \hat{x}) = -\frac{1}{2}\|\hat{G}^{i,N}(\theta, \hat{x}^N)\|^2$. We will also write $\hat{\mu}_t^N = \mathcal{L}(\hat{x}_t^N)$ to denote the law of $\hat{x}_t^N = (x_t^{1,N}, \ldots, x_t^{N,N})$.

It is worth noting that, in the bounds which appear in subsequent section, we will use $0 < K < \infty$ or $0 < K_{\theta_0} < \infty$ to denote fixed constants which do not depend on the number of particles N, the time t, or any other important parameters. In particular, we will use K when the constant is independent of the true parameter θ_0 (e.g., the constants in Assumption C.1), and K_{θ_0} when the constant depends on θ_0 (e.g., the constants arising in Propositions 1–4). The constant K or K_{θ_0} may change from line to line, but it will always be denoted by the same symbol.

4.2. Proof of Theorems 1.1, 1.1*, 2.1 and 2.1*

Proof. We will focus on the proof of Theorems 2.1 and 2.1*, later demonstrating how the same approach can be used to obtain Theorems 1.1 and 1.1*. The proof of these results will rely on a series of lemmas, which will be proved subsequently.

Using the triangle inequality, we can decompose $\tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})$, the approximate asymptotic log-likelihood of the McKean–Vlasov SDE evaluated at the online parameter estimates generated by the IPS, as

$$\|\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})\| \leq \underbrace{\|\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N}) - \tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^i(\theta_t^{i,N})\|}_{\to 0 \text{ as } t\to\infty \quad \forall N\in\mathbb{N} \text{ by Lemma 1}} + \underbrace{\|\tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^i(\theta_t^{i,N}) - \tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta_t^{i,N})\|}_{\to 0 \text{ as } N\to\infty \quad \forall t\in\mathbb{R}_+ \text{ by Lemma 4}}$$

$$+ \underbrace{\|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta_t^{i,N}) - \tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{i,N}(\theta_t^{i,N})\|}_{\to 0 \text{ as } N\to\infty \quad \forall t \in \mathbb{R}_+ \text{ by Lemma 3}} + \underbrace{\|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{i,N}(\theta_t^{i,N}) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\|}_{\to 0 \text{ as } t\to\infty \quad \forall N \in \mathbb{N} \text{ by Lemma 2}}$$

$$+ \underbrace{\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\|}_{\to 0 \text{ as } t\to\infty \quad \forall N \in \mathbb{N} \text{ by Lemma 5}}, \tag{22}$$

or, almost identically, as

$$\|\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^N)\| \leq \underbrace{\|\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^N) - \tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^i(\theta_t^N)\|}_{\to 0 \text{ as } t\to\infty \quad \forall N \in \mathbb{N} \text{ by Lemma 1}} + \underbrace{\|\tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^i(\theta_t^N) - \tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[N]}(\theta_t^N)\|}_{\to 0 \text{ as } N\to\infty \quad \forall t \in \mathbb{R}_+ \text{ by Lemma 4}}$$

$$+ \underbrace{\|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[N]}(\theta_t^N) - \tfrac{1}{t}\nabla_\theta \mathcal{L}_t^N(\theta_t^N)\|}_{\to 0 \text{ as } N\to\infty \quad \forall t \in \mathbb{R}_+ \text{ by Lemma 3}} + \underbrace{\|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^N(\theta_t^N) - \nabla_\theta \tilde{\mathcal{L}}^N(\theta_t^N)\|}_{\to 0 \text{ as } t\to\infty \quad \forall N \in \mathbb{N} \text{ by Lemma 2}}$$

$$+ \underbrace{\|\nabla_\theta \tilde{\mathcal{L}}^N(\theta_t^N)\|}_{\to 0 \text{ as } t\to\infty \quad \forall N \in \mathbb{N} \text{ by Lemma 5}}, \tag{23}$$

where $\tilde{\underline{\mathcal{L}}}(\theta)$ is defined in (13), $\tilde{\mathcal{L}}^{[i,N]}(\theta)$ and $\tilde{\mathcal{L}}^{[N]}(\theta)$ are defined in (12), and $\tilde{\mathcal{L}}^{i,N}(\theta)$ and $\tilde{\mathcal{L}}^N(\theta)$ are defined in (16).

In both (22) and (23), the final limit holds a.s. This proves Theorem 2.1. Meanwhile, all of the limits hold in $\mathbb{L}^1$. This proves the first statement in Theorem 2.1*. Now suppose that $\Theta_0 = \{\theta : \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta) = 0\} = \{\theta_0\}$. Then, from the previous result, we have $\theta_t^{i,N} \to \theta_0$ and $\theta_t^N \to \theta_0$ in $\mathbb{L}^1$. By continuity, it follows immediately that $\|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^{i,N})\| \to \|\nabla_\theta \tilde{\mathcal{L}}(\theta_0)\| = 0$ and $\|\nabla_\theta \tilde{\mathcal{L}}(\theta_t^N)\| \to \|\nabla_\theta \tilde{\mathcal{L}}(\theta_0)\| = 0$ in $\mathbb{L}^1$, using the definition of $\tilde{\mathcal{L}}(\theta)$ in (7). This proves the second statement in Theorem 2.1*.

Theorems 1.1 and 1.1* are proved in the same fashion. In this case, in (22) and (23), we must now replace $\theta_t^{i,N}$ by $\theta_t^{[i,N]}$, $\mathcal{L}_t^{i,N}(\cdot)$ by $\mathcal{L}_t^{[i,N]}(\cdot)$, and $\tilde{\mathcal{L}}^{i,N}(\cdot)$ by $\tilde{\mathcal{L}}^{[i,N]}(\cdot)$ in (22), and θ_t^N by $\theta_t^{[N]}$, $\mathcal{L}_t^N(\cdot)$ by $\mathcal{L}_t^{[N]}(\cdot)$, and $\tilde{\mathcal{L}}^N(\cdot)$ by $\tilde{\mathcal{L}}^{[N]}(\cdot)$ in (23). The third term in (22) and (23) now vanishes entirely. The limits for the remaining terms in (22) and (23), established in Lemmas 1, 2, 4, and 5 all still hold. Thus, arguing as above, we obtain the claimed results. □

Before we proceed to the proofs of Lemmas 1–5, let us provide a brief high level overview of these results. The proofs of these Lemmas will themselves rely on several auxiliary results, whose full statements and proofs can be found in Appendices A and B.

(i) In Lemmas 1 and 2, we establish the existence of $\tilde{\underline{\mathcal{L}}}(\theta)$, $\tilde{\mathcal{L}}^{i,N}(\theta)$, and $\tilde{\mathcal{L}}^N(\theta)$, as well as their derivatives. We provide explicit expressions for these functions in terms of the unique invariant measure of the McKean–Vlasov SDE or the IPS, and prove appropriate convergence results as $t \to \infty$ (both a.s. and in $\mathbb{L}^1$).

(iii) In Lemma 3, we prove that, for all $t \geq 0$, the gradient of the particle based approximation of the log-likelihood of the McKean–Vlasov SDE, as defined in (5), converges to the gradient of the log-likelihood of the IPS as $N \to \infty$ (in $\mathbb{L}^1$). We also provide $\mathbb{L}^1$ convergence rates. The proof of this result relies on classical uniform-in-time propagation of chaos results [17,72].

(iv) In Lemma 4, we prove that, for all $t \geq 0$, the gradient of the particle based approximation of the log-likelihood of the McKean–Vlasov SDE in (5) converges to the gradient of the 'approximate' log-likelihood of the McKean–Vlasov SDE in (20) as $N \to \infty$ (in $\mathbb{L}^1$). We also provide $\mathbb{L}^1$ convergence rates. The proof relies on existing results on the convergence of the empirical law of the McKean–Vlasov SDE [39].

(vi) In Lemma 5, we establish that, for all $N \in \mathbb{N}$, the gradient of the asymptotic log-likelihood of the IPS, evaluated at the relevant online parameter updates, converges to zero as $t \to \infty$ (both a.s. and in $\mathbb{L}^1$). This result can be seen as a generalisation of [94, Theorem 2.4].

Lemma 1. *Assume that Conditions A.1, B.1–B.2, and C.1 hold. Then the processes* $\frac{1}{t}\nabla_\theta^k \underline{\mathcal{L}}_t^i(\theta)$, $k = 0, 1, 2$, *converge, both a.s. and in* $\mathbb{L}^1$, *to the functions*

$$\nabla_\theta^k \tilde{\underline{\mathcal{L}}}(\theta) = \int_{\mathbb{R}^d} \nabla_\theta^k L(\theta, x, \mu_\infty, \mu_\infty)\mu_\infty(\mathrm{d}x).$$

Proof. We will prove Lemma 1 for the function $\underline{\mathcal{L}}_t^i(\cdot)$, in the case $k = 0$. The results for $k = 1, 2$, are proved similarly.[1] Let us begin by recalling the definition of $\frac{1}{t}\mathcal{L}_t^i(\theta)$, viz

$$\frac{1}{t}\underline{\mathcal{L}}_t^i(\theta) = \underbrace{\frac{1}{t}\int_0^t L(\theta, x_s^i, \mu_s, \mu_s)\mathrm{d}s}_{I_1^N(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t \langle G(\theta, x_s^i, \mu_s, \mu_s), \mathrm{d}w_s^i\rangle}_{I_2^N(\theta,t)}.$$

We will first consider the first term on the RHS. We will characterise the asymptotic behaviour of this term via the following decomposition

$$\underbrace{\frac{1}{t}\int_0^t L(\theta, x_s^i, \mu_s, \mu_s)\mathrm{d}s}_{I_1^N(\theta,t)} = \underbrace{\frac{1}{t}\int_0^t \left[L(\theta, x_s^i, \mu_s, \mu_s) - L(\theta, x_s^i, \mu_\infty, \mu_\infty)\right]\mathrm{d}s}_{I_{1,1}^N(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t L(\theta, x_s^i, \mu_\infty, \mu_\infty)\mathrm{d}s}_{I_{1,2}^N(\theta,t)}, \tag{24}$$

where μ_∞ is the unique invariant measure of $(x_t)_{t\geq 0}$, which exists as a consequence of Proposition 3 (see Appendix A). We begin with the observation that, as $t \to \infty$, we have

$$\begin{aligned}&\frac{1}{t}\int_0^t [L(\theta, x_s^i, \mu_s, \mu_s) - L(\theta, x_s^i, \mu_\infty, \mu_\infty)]\mathrm{d}s \xrightarrow{\text{a.s.}} 0 \quad \text{and}\\ &\frac{1}{t}\int_0^t L(\theta, x_s^i, \mu_\infty, \mu_\infty)\mathrm{d}s \xrightarrow[\mathbb{L}^1]{\text{a.s.}} \tilde{\underline{\mathcal{L}}}(\theta),\end{aligned} \tag{25}$$

the former by Proposition 3 (see Appendix A,) and the latter by an appropriate version of the ergodic theorem (e.g., [87, Chapter X]). Let us now demonstrate that $I_{1,1}^N(\theta, t)$ also converges to zero in $\mathbb{L}^1$. Using lity Lemma 14, we have that

$$\begin{aligned}&\left\|L(\theta, x_s^i, \mu_s, \mu_s) - L(\theta, x_s^i, \mu_\infty, \mu_\infty)\right\|\\ &\quad\leq K\mathbb{W}_2(\mu_s, \mu_\infty)\left[1 + \|x_s^i\|^q + \mu_s(\|\cdot\|^q) + \mu_\infty(\|\cdot\|^q)\right].\end{aligned}$$

Thus, using also Proposition 3 (the exponential contractivity of the McKean–Vlasov SDE), Condition A.1 (integrability of the initial condition), Proposition 1 (moment bounds for the McKean–Vlasov SDE), and Proposition 2 (moment bounds for the invariant measure of the

[1] Regarding the results for $k = 1, 2$, we note that the processes $\frac{1}{t}\nabla_\theta^k \underline{\mathcal{L}}_t^i(\theta)$, and hence also $\nabla_\theta^k \tilde{\underline{\mathcal{L}}}(\theta)$, exist due to Condition C.1. With this established, the proof when $k = 1, 2$ is almost identical to the proof when $k = 0$.

McKean–Vlasov SDE), it follows that

$$\begin{aligned}\left\| L(\theta, x_s^i, \mu_s, \mu_s) - L(\theta, x_s^i, \mu_\infty, \mu_\infty) \right\| &\leq K \mathbb{W}_2(\mu_0, \mu_\infty) \left[1 + \|x_s^i\|^q \right. \\ &\quad \left. + \mu_s(\|\cdot\|^q) + \mu_\infty(\|\cdot\|^q) \right] e^{-\lambda_{\theta_0} s} \\ &\leq K_{\theta_0} \left[1 + \|x_s^i\|^q \right] e^{-\lambda_{\theta_0} s},\end{aligned}$$

where $\lambda_{\theta_0} = \alpha_{\theta_0} - 2L_{\theta_0,2} > 0$ (from Proposition 3), and $\alpha_{\theta_0}, \lambda_{\theta_0} > 0$ are the constants defined in Condition B.2. Thus, making use once more of Proposition 1, and allowing the value of the constant K_{θ_0} to increase between inequalities, that

$$\mathbb{E}_{\theta_0}[|I_{1,1}^N(t)|] \leq \frac{1}{t} \int_0^t K_{\theta_0} \left(1 + \mathbb{E}_{\theta_0}\left[\|x_s^i\|^q\right]\right) e^{-\lambda_{\theta_0} s} \mathrm{d}s \leq \frac{K_{\theta_0}}{t} \int_0^t e^{-\lambda_{\theta_0} s} \mathrm{d}s \leq \frac{K_{\theta_0}(1 - e^{-\lambda_{\theta_0} t})}{\lambda_{\theta_0} t},$$

so that the convergence of $I_{1,1}^N(\theta, t)$ to zero does also hold in $\mathbb{L}^1$. We thus have, substituting (25) into (24), that $I_1^N(\theta, t) \to \tilde{\mathcal{L}}(\theta)$, both a.s. and in $\mathbb{L}^1$.

We now turn our attention $I_2^N(\theta, t)$. Using the Itô's isometry, Condition C.1(ii) (the polynomial growth of B and therefore also G), Proposition 1 (the bounded moments of the McKean–Vlasov SDE), and Lemma 6 (the asymptotic growth rate of the moments of the McKean–Vlasov SDE), we have

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\left| \int_0^t \langle G(\theta, x_s^i, \mu_s, \mu_s), \mathrm{d}w_s^i \rangle \right|^2 \right] &= \mathbb{E}_{\theta_0}\left[\int_0^t \|G(\theta, x_s^i, \mu_s, \mu_s)\|^2 \mathrm{d}s \right] \\ &\leq \mathbb{E}_{\theta_0}\left[\int_0^t K \left(1 + \|x_s^i\|^q + \mu_s(\|\cdot\|^q)\right) \mathrm{d}s \right] \\ &\leq K_{\theta_0} t \left[1 + \sqrt{t} \right] \end{aligned} \tag{26}$$

where, in the final inequality, we have once more made explicit the dependence of the constant K_{θ_0} on the true parameter θ_0 (due to Proposition 1 and Lemma 6). It follows that for each $\theta \in \mathbb{R}^p$,

$$\mathbb{E}_{\theta_0}\left[\| \frac{1}{t} \langle G(\theta, x_s, \mu_s), \mathrm{d}w_s \rangle \|^2 \right] \leq \frac{K_{\theta_0}(1 + \sqrt{t})}{t},$$

and so this term converges in $\mathbb{L}^2$ (and hence $\mathbb{L}^1$) to zero. It remains only to demonstrate a.s. convergence of this term to zero. To do so, consider the local martingale

$$\begin{aligned}M_t &= \int_0^t \frac{1}{s} \langle G(\theta, x_s^i, \mu_s^\theta, \mu_s), \mathrm{d}w_s^i \rangle \\ &= \frac{1}{t} \int_0^t \langle G(\theta, x_s^i, \mu_s^\theta, \mu_s), \mathrm{d}w_s^i \rangle + \int_0^t \frac{1}{s^2} \left[\int_0^s \langle G(\theta, x_u^i, \mu_u^\theta, \mu_u), \mathrm{d}w_u^i \rangle \right] \mathrm{d}s,\end{aligned}$$

where the second equality follows from Itô's Lemma. Using the Itô isometry, Condition C.1(ii) (which implies the polynomial growth of G), and Proposition 1 (the bounded moments of the McKean–Vlasov SDE), and arguing as above, we have

$$\begin{aligned}\sup_{t>0} \mathbb{E}_{\theta_0}\left[|M_t|^2\right] &= \mathbb{E}_{\theta_0}\left[\int_0^\infty \frac{1}{s^2} \mathbb{E}_{\theta_0}\left[\|G(\theta, x_s^i, \mu_s^\theta, \mu_s)\|^2 \right] \mathrm{d}s \right] \\ &\leq K_{\theta_0}\left[\int_0^t \frac{1}{s^2} \left(1 + \mathbb{E}_{\theta_0}\left[\|x_s^i\|^q\right]\right) \mathrm{d}s \right] < \infty. \end{aligned} \tag{27}$$

By Doob's martingale convergence theorem [36], there thus exists a finite random variable M_∞ such that $M_t \to M_\infty$ a.s. It follows immediately that $\frac{1}{t}\int_0^t \langle G(\theta, x_s^i, \mu_s^\theta, \mu_s), \mathrm{d}w_s^i\rangle$ also converges to zero a.s., as claimed. Putting everything together, we thus have that $\frac{1}{t}\mathcal{L}_t^i(\theta)$ converges to $\tilde{\mathcal{L}}(\theta)$ both a.s. and in $\mathbb{L}^1$. □

Lemma 2. *Assume that Conditions A.1, B.1–B.2, and C.1 hold. Then, for all $N \in \mathbb{N}$, the processes $\frac{1}{t}\nabla_\theta^k \mathcal{L}_t^{i,N}(\theta)$ and $\frac{1}{t}\nabla_\theta^k \mathcal{L}_t^N(\theta)$, $m = 0, 1, 2$, converge, both a.s. and in $\mathbb{L}^1$, to the functions*

$$\nabla_\theta^k \tilde{\mathcal{L}}^{i,N}(\theta) = \int_{(\mathbb{R}^d)^N} \nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{x}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) \quad , \quad \nabla_\theta^k \tilde{\mathcal{L}}^N(\theta) = \frac{1}{N}\sum_{i=1}^N \nabla_\theta^k \tilde{\mathcal{L}}^{i,N}(\theta).$$

Proof. We will begin by proving that the statement holds for the function $\mathcal{L}_t^{i,N}(\theta)$. The proof, in this case, is very similar to the proof of Lemma 1, with some simplifications. We will provide a sketch of the proof, signposting differences with the previous proof where necessary. As previously, we will only consider the case $k = 0$, with the results for $k = 1, 2$ proved analogously. We begin by recalling the definition of the function $\frac{1}{t}\mathcal{L}_t^{i,N}(\theta)$ from (6), which we now write in the form

$$\frac{1}{t}\mathcal{L}_t^{i,N}(\theta) = \frac{1}{t}\int_0^t \hat{L}^{i,N}(\theta, \hat{x}_s^N)\mathrm{d}s + \frac{1}{t}\int_0^t \langle \hat{G}^{i,N}(\theta, \hat{x}_s^N), \mathrm{d}w_s^i\rangle. \tag{28}$$

We begin with the first term on the RHS. By Proposition 3, the IPS admits a unique invariant measure $\hat{\mu}_\infty^N \in \mathcal{P}((\mathbb{R}^d)^N)$. Thus, for all $N \in \mathbb{N}$, by the ergodic theorem (e.g., [87, Chapter X]) we have that as $t \to \infty$,

$$\frac{1}{t}\int_0^t \hat{L}^{i,N}\left(\theta, \hat{x}_s^N\right)\mathrm{d}s \xrightarrow[\mathbb{L}^1]{\text{a.s.}} \int_{(\mathbb{R}^d)^N} \hat{L}^{i,N}\left(\theta, \hat{x}^N\right)\hat{\mu}_\infty(\mathrm{d}\hat{x}^N) = \tilde{\mathcal{L}}^{i,N}(\theta).$$

It remains to bound the second term on the RHS of (28). We show that this term converges to zero a.s. and in $\mathbb{L}^1$, and satisfies the required convergence rate, using essentially identical arguments to those used in the proof of Lemma 1, c.f. (26)–(27). This concludes the proof. We now turn our attention to the function $\mathcal{L}_t^N(\theta)$. The proof of the statements regarding this function now follows easily. In particular, using the definition of $\mathcal{L}_t^N(\theta)$, c.f. (6), and the results above, we have

$$\frac{1}{t}\mathcal{L}_t^N(\theta) = \frac{1}{t}\left[\frac{1}{N}\sum_{i=1}^N \mathcal{L}_t^{i,N}(\theta)\right] = \frac{1}{N}\sum_{i=1}^N\left[\frac{1}{t}\mathcal{L}_t^{i,N}(\theta)\right] \xrightarrow[\mathbb{L}^1]{\text{a.s.}} \frac{1}{N}\sum_{i=1}^N \tilde{\mathcal{L}}^{i,N}(\theta) = \tilde{\mathcal{L}}^N(\theta). \quad \square$$

Lemma 3. *Assume that Conditions A.1, B.1–B.2, and C.1 hold. Then, for all $\theta \in \mathbb{R}^p$, for all $t > 0$, for all $i = 1, \ldots, N$, we have, in $\mathbb{L}^1$, that*

$$\lim_{N\to\infty} \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{i,N}(\theta)\| = \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta)\| \quad \text{and} \quad \lim_{N\to\infty} \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{N}(\theta)\| = \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[N]}(\theta)\|.$$

In addition, there exists a positive constant K_{θ_0} such that, for all $\theta \in \mathbb{R}^p$, for all $t > 0$, for all $N \in \mathbb{N}$, and for all $i = 1, \ldots, N$,

$$\mathbb{E}_{\theta_0}\left[\left\|\frac{1}{t}\nabla_\theta \mathcal{L}_t^{i,N}(\theta) - \frac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta)\right\|\right] \leq \frac{K_{\theta_0}}{\sqrt{N}}\left(1 + \frac{1}{\sqrt{t}}\right),$$

and this bound also holds if $\mathcal{L}_t^{i,N}(\cdot)$ and $\mathcal{L}_t^{[i,N]}$ are replaced by $\mathcal{L}_t^N(\cdot)$ and $\mathcal{L}_t^{[N]}(\cdot)$.

Proof. We begin by proving that the two statements hold for $\mathcal{L}_t^{i,N}(\theta)$. First recall that

$$\frac{1}{t}\nabla_\theta \mathcal{L}_t^{i,N}(\theta) = \underbrace{\frac{1}{t}\int_0^t \nabla_\theta L(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N)\mathrm{d}s}_{I_1^{i,N}(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t \left\langle \nabla_\theta G(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N), \mathrm{d}w_s^i \right\rangle}_{I_2^{i,N}(\theta,t)},$$

$$\frac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta) = \underbrace{\frac{1}{t}\int_0^t \nabla_\theta L(\theta, x_s^{i}, \mu_s^{[N]}, \mu_s^{[N]})\mathrm{d}s}_{I_1^{[i,N]}(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t \left\langle \nabla_\theta G(\theta, x_s^{i}, \mu_s^{[N]}, \mu_s^{[N]}), \mathrm{d}w_s^i \right\rangle}_{I_2^{[i,N]}(\theta,t)}.$$

We will begin by seeking a bound for $\mathbb{E}_{\theta_0}[\|I_2^{i,N}(\theta,t) - I_2^{[i,N]}(\theta,t)\|]$. Using Assumption C.1, Proposition 1 (bounded moments of the McKean–Vlasov SDE and the IPS), and Proposition 4 (uniform-in-time propagation-of-chaos), there exists a constant $K_{\theta_0} > 0$, such that

$$\begin{aligned}&\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\|\nabla_\theta G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta G(\theta, x_t^i, \mu_t^{[N]}, \mu_t^{[N]})\|^2\right] \\ &\quad\leq K_{\theta_0}\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\|x_t^i - x_t^{i,N}\|^2 + \mathbb{W}_2^2(\mu_t^N, \mu_t^{[N]})\right] \leq \frac{K_{\theta_0}}{N}. \end{aligned} \tag{29}$$

It follows, using also the Itô isometry and Fubini's Theorem, that for all $\theta \in \mathbb{R}^p$, and for all $t > 0$,

$$\begin{aligned}&\mathbb{E}_{\theta_0}\left[\|I_2^{i,N}(\theta,t) - I_2^{[i,N]}(\theta,t)\|^2\right] \\ &\quad\leq \frac{K_{\theta_0}}{t^2}\int_0^t \mathbb{E}_{\theta_0}\left[\|\nabla_\theta G(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N) - \nabla_\theta G(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]})\|^2\right]\mathrm{d}s \leq \frac{K_{\theta_0}}{Nt}, \end{aligned} \tag{30}$$

and so, by the Hölder inequality,

$$\mathbb{E}_{\theta_0}\left[\|I_2^{i,N}(\theta,t) - I_2^{[i,N]}(\theta,t)\|\right] \leq \frac{K_{\theta_0}}{\sqrt{Nt}}. \tag{31}$$

In much the same fashion, we can obtain a bound for $\mathbb{E}_{\theta_0}\|I_1^{i,N}(\theta,t) - I_1^{[i,N]}(\theta,t)\|$. We begin by noting that

$$\begin{aligned}&\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\left\|\nabla_\theta L(\theta, x_s^{i,N}, \mu_s^N, \mu_s^N) - \nabla_\theta L(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]})\right\|\right] \\ &= \sup_{t\geq 0}\mathbb{E}_{\theta_0}\Big[\Big\| \underbrace{\nabla_\theta G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N)}_{Y_N}\underbrace{G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N)}_{Z_N} \\ &\qquad - \underbrace{\nabla_\theta G(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})}_{Y_{[N]}}\underbrace{G(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})}_{Z_{[N]}}\Big\|\Big] \\ &\leq \sup_{t\geq 0}\left[\mathbb{E}_{\theta_0}\left[\|Z_N\|^2\right]^{\frac{1}{2}}\mathbb{E}_{\theta_0}\left[\|Y_N - Y_{[N]}\|^2\right]^{\frac{1}{2}} + \mathbb{E}_{\theta_0}\left[\|Y_{[N]}\|^2\right]^{\frac{1}{2}}\mathbb{E}_{\theta_0}\left[\|Z_N - Z_{[N]}\|^2\right]^{\frac{1}{2}}\right], \end{aligned} \tag{32}$$

where the final line follows from the triangle inequality and the Cauchy–Schwarz inequality. Now, using Condition C.1 (the polynomial growth of G) and Proposition 1 (bounded moments of the McKean–Vlasov SDE and the IPS) for the two left-most inequalities, and arguing as in (29) for the two right-most inequalities, there exist constants $K_{\theta_0,1}, K_{\theta_0,2}, K_{\theta_0,3}, K_{\theta_0,4} > 0$ (all

independent of θ), such that

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\| \underbrace{G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N)}_{Z_N} \|^2\right] \leq K_{\theta_0,1}, \tag{33}$$

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\| \underbrace{\nabla_\theta G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N)}_{Y_N} - \underbrace{\nabla_\theta G(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})}_{Y_{[N]}} \|^2 \leq \frac{K_{\theta_0,2}}{N}, \tag{34}$$

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\| \underbrace{\nabla_\theta G(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})}_{Y_{[N]}} \|^2\right] \leq K_{\theta_0,3}, \tag{35}$$

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\| \underbrace{G(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N)}_{Z_N} - \underbrace{G(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})}_{Z_{[N]}} \|^2 \leq \frac{K_{\theta_0,4}}{N}. \tag{36}$$

Substituting (33)–(36) into (32), we finally arrive at

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\left\|\nabla_\theta L(\theta, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta L(\theta, x_t^{i}, \mu_t^{[N]}, \mu_t^{[N]})\right\|\right]$$
$$\leq K_{\theta_0,1}^{\frac{1}{2}}\left[\frac{K_{\theta_0,2}}{N}\right]^{\frac{1}{2}} + K_{\theta_0,3}^{\frac{1}{2}}\left[\frac{K_{\theta_0,4}}{N}\right]^{\frac{1}{2}} \leq \frac{K_{\theta_0}}{\sqrt{N}}$$

and so, arguing similarly to (30), that for all $t > 0$,

$$\mathbb{E}_{\theta_0}\left[\|I_1^{i,N}(\theta, t) - I_1^{[i,N]}(\theta, t)\|\right] \leq \frac{K_{\theta_0}}{\sqrt{N}}. \tag{37}$$

Combining inequalities (31) and (37), and making use of the triangle inequality one final time, we have that the desired result. This establishes convergence in $\mathbb{L}^1$ as $N \to \infty$, for all $t > 0$. It remains only to establish that the statements of the lemma also hold for $\mathcal{L}_t^N(\theta)$, which now follows immediately. We omit the calculations, which are essentially identical to those used at the end of the proof of Lemma 2. □

Lemma 4. *Assume that Conditions A.1, B.1–B.2, and C.1 hold. Then, for all $\theta \in \mathbb{R}^p$, for all $t \geq 0$, for all $i = 1, \ldots, N$, we have, in $\mathbb{L}^1$, that*

$$\lim_{N\to\infty} \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta)\| = \|\tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^{i}(\theta)\| \quad \text{and} \quad \lim_{N\to\infty} \|\tfrac{1}{t}\nabla_\theta \mathcal{L}_t^{[N]}(\theta)\| = \|\tfrac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^{i}(\theta)\|.$$

In addition, there exists a positive constant K_{θ_0} such that, for all $\theta \in \mathbb{R}^p$, for all $t > 0$, for all $N \in \mathbb{N}$, and for all $i = 1, \ldots, N$,

$$\mathbb{E}_{\theta_0}\left[\left\|\frac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^{i}(\theta) - \frac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta)\right\|\right] \leq K_{\theta_0}\alpha(N)\left(1 + \frac{1}{\sqrt{t}}\right),$$

where the function $\alpha : \mathbb{N} \to \mathbb{R}_+$ is defined according to

$$\alpha(N) = \begin{cases} N^{-\frac{1}{4}} & \text{if} \quad d = 1, \\ N^{-\frac{1}{4}}\log(1+N)^{\frac{1}{2}} & \text{if} \quad d = 2, \\ N^{-\frac{1}{2d}} & \text{if} \quad d \geq 3. \end{cases}$$

Moreover, this bound also holds if $\mathcal{L}_t^{[i,N]}(\cdot)$ is replaced by $\mathcal{L}_t^{[N]}(\cdot)$.

Proof. The proof of this result is very similar to the proof of the previous result. We will just sketch the main details. Recall that

$$\frac{1}{t}\nabla_\theta \mathcal{L}_t^{[i,N]}(\theta) = \underbrace{\frac{1}{t}\int_0^t \nabla_\theta L(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]})\mathrm{d}s}_{I_1^{[i,N]}(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t \left\langle \nabla_\theta G(\theta, x_s^i, \mu_s^{[N]}, \mu_s^{[N]}), \mathrm{d}w_s^i \right\rangle}_{I_2^{[i,N]}(\theta,t)},$$

$$\frac{1}{t}\nabla_\theta \underline{\mathcal{L}}_t^{i}(\theta) = \underbrace{\frac{1}{t}\int_0^t \nabla_\theta L(\theta, x_s^i, \mu_s, \mu_s)\mathrm{d}s}_{I_1^{i}(\theta,t)} + \underbrace{\frac{1}{t}\int_0^t \left\langle \nabla_\theta G(\theta, x_s^i, \mu_s, \mu_s), \mathrm{d}w_s^i \right\rangle}_{I_2^{i}(\theta,t)}.$$

Using Assumption C.1 and Proposition 1 (bounded moments of the McKean–Vlasov SDE), there exists a constant $K_{\theta_0} > 0$ (independent of θ), such that

$$\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\|\nabla_\theta G(\theta, x_t^i, \mu_t^{[N]}, \mu_t^{[N]}) - \nabla_\theta G(\theta, x_t^i, \mu_t, \mu_t)\|^2\right] \leq K_{\theta_0}\sup_{t\geq 0}\mathbb{E}_{\theta_0}\left[\mathbb{W}_2^2(\mu_t^{[N]}, \mu_t)\right] \leq K_{\theta_0}\alpha(N)^2,$$

where the final inequality follows from Theorem 1 in [39] (the rate of convergence of the empirical law of the McKean–Vlasov SDE to the true law of the McKean–Vlasov SDE). It follows, arguing as in Lemma 3, that

$$\mathbb{E}_{\theta_0}\left[\|I_1^{[i,N]}(\theta,t) - I_1^i(\theta,t)\|\right] \leq K_{\theta_0}\sqrt{\alpha(N)^2},$$

$$\mathbb{E}_{\theta_0}\left[\|I_2^{[i,N]}(\theta,t) - I_2^i(\theta,t)\|\right] \leq K_{\theta_0}\sqrt{\frac{\alpha(N)^2}{t}}.$$

Combining these bounds, and using the triangle inequality, we obtain the stated bound for $\mathcal{L}_t^{[i,N]}(\cdot)$. The result for $\mathcal{L}_t^{[N]}(\cdot)$ follows by arguing as at the end of Lemma 2. □

Lemma 5. *Assume that Conditions B.1–B.2, C.1, A.1 and F.1 hold. Then, for all $N \in \mathbb{N}$, we have, both almost surely and in $\mathbb{L}^1$, that*

$$\lim_{t\to\infty}\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| = 0 \quad \textit{and} \quad \lim_{t\to\infty}\|\nabla_\theta \tilde{\mathcal{L}}^{N}(\theta_t^{N})\| = 0.$$

Proof. We will prove the first statement of the lemma, with the second proved identically. We will use a modified version of the approach in [94], which itself is a continuous-time version of the approach first introduced in [8]. We will require the following additional notation. Define an arbitrary constant $\kappa > 0$, with $\lambda = \lambda(\kappa) > 0$ to be determined. Set $\sigma = 0$, and define the cycle of random stopping times $0 = \sigma_0 \leq \tau_1 \leq \sigma_1 \leq \tau_2 \leq \sigma_2 \leq \dots$ via

$$\tau_k = \inf\left\{t > \sigma_{k-1} : \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| \geq \kappa\right\}$$

$$\sigma_k = \sup\Big\{ t > \tau_k : \tfrac{1}{2}\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N})\| \leq \|\nabla \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\| \leq 2\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N})\| \ \ \forall s \in [\tau_k, t], \quad \int_{\tau_k}^t \gamma(s)\mathrm{d}s \leq \rho \Big\}.$$

The purpose of these stopping times is to control the periods of time for which $\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\|$ is close to zero, and those for which it is away from zero. In addition, let $\eta > 0$, and set $\sigma_{k,\eta} = \sigma_k + \eta$. We are now ready to prove this result. First consider the case in which there are a finite number of stopping times τ_k. In this case, there exists finite t_0 such that, for all $t \geq t_0$,

$\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| < \kappa$. Now consider the case in which there are an infinite number of stopping times τ_k. Then, using Lemmas 12–13 (see Appendix B), there exist $0 < \beta_1 < \beta$, and $k_0 \in \mathbb{N}$, such that for all $k \geq k_0$, almost surely,

$$\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_k}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k}) \geq \beta \quad \text{and} \quad \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_{k-1}}) \geq -\beta_1.$$

It follows straightforwardly that

$$\begin{aligned}\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_{n+1}}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_{k_0}}) &= \sum_{k=k_0}^{n} \left[\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_k}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k}) + \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_{k+1}}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_k})\right] \\ &\geq (n+1-k_0)(\beta-\beta_1).\end{aligned}$$

Since $\beta - \beta_1 > 0$, this implies that $\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_{n+1}}) \to \infty$ as $n \to \infty$. But this contradicts Lemma 7 (see Appendix B), which states that $\tilde{\mathcal{L}}^{i,N}(\theta)$ is bounded from above. Thus, there must exist a finite time t_0 such that, for all $t \geq t_0$, $\|\tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| < \kappa$. Since our original choice of κ was arbitrary, this completes the proof that, for all $N \in \mathbb{N}$, $\lim_{t\to\infty} \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\| = 0$ a.s. Finally, we observe that, by Lemma 7, $\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta)\|$ is bounded above for all $\theta \in \mathbb{R}^p$. Thus, we also have convergence in $\mathbb{L}^1$ via Lebesgue's dominated convergence theorem (e.g., [104, Chapter 5]). □

4.3. *Proof of Theorems 1.2 , 1.2*, 2.2, and 2.2**

Proof. The proof of this result closely follows the proof of Theorem 2.7 in [96], adapted appropriately to our particular case. We will focus on Theorem 2.2*, later outlining how our approach can be adapted to obtain the other results. We begin by rewriting the parameter update equation (14) in the following form

$$\begin{aligned}\mathrm{d}\theta_t^{i,N} &= \gamma_t \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})\mathrm{d}t + \gamma_t\big(\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N}) - \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})\big)\mathrm{d}t \\ &\quad + \gamma_t\big(\nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\big)\mathrm{d}t \\ &\quad + \gamma_t \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i. \end{aligned} \tag{38}$$

Using a first order Taylor expansion, and the fact that $\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_0) = 0$, we have that

$$\nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N}) = \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_0) + \nabla^2_\theta \tilde{\underline{\mathcal{L}}}(\tilde{\underline{\theta}}_t^{i,N})(\theta_t^{i,N} - \theta_0) = \nabla^2 \tilde{\underline{\mathcal{L}}}(\tilde{\underline{\theta}}_t^{i,N})(\theta_t^{i,N} - \theta_0), \tag{39}$$

where $\nabla^2 \underline{\mathcal{L}}(\cdot)$ denotes the Hessian, and $\tilde{\underline{\theta}}_t^{i,N}$ is a point in the segment connecting $\theta_t^{i,N}$ and θ_0. Substituting (39) into (38), we obtain the following equations for $Z_t^{i,N} = \theta_t^{i,N} - \theta_0$

$$\begin{aligned}\mathrm{d}Z_t^{i,N} &= \gamma_t \nabla^2_\theta \tilde{\underline{\mathcal{L}}}(\tilde{\underline{\theta}}_t^{i,N}) Z_t^{i,N}\mathrm{d}t + \gamma_t\big(\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N}) - \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})\big)\mathrm{d}t \\ &\quad + \gamma_t\big(\nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\big)\mathrm{d}t \\ &\quad + \gamma_t \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i. \end{aligned}$$

Applying Itô's formula to the function $\|\cdot\|^2$, and using the strong concavity of $\tilde{\underline{\mathcal{L}}}(\theta)$ (Condition E.2), it follows that

$$\begin{aligned}\mathrm{d}\|Z_t^{i,N}\|^2 + 2\eta\gamma_t \|Z_t^{i,N}\|^2 \mathrm{d}t &\leq \gamma_t\big\langle Z_t^{i,N}, \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N}) - \nabla_\theta \tilde{\underline{\mathcal{L}}}(\theta_t^{i,N})\big\rangle \mathrm{d}t \\ &\quad + \gamma_t\big\langle Z_t^{i,N}, \nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\big\rangle \mathrm{d}t \\ &\quad + \gamma_t\big\langle Z_t^{i,N}, \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\big\rangle \\ &\quad + \gamma_t^2 \big|\big|\nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\big|\big|_F^2 \mathrm{d}t, \end{aligned} \tag{40}$$

where $\|\cdot\|_F$ is the Frobenius norm. Now, let us define the function $\Phi_{t,t'} = \exp[-2\eta\int_t^{t'}\gamma_u \mathrm{d}u]$, with $\partial_t \Phi_{t,t'} = 2\eta\gamma_t \Phi_{t,t'}$. Using the product rule, and (40), we obtain

$$\begin{aligned}
\mathrm{d}\left[\Phi_{t,t'}\|Z_t^{i,N}\|^2\right] &= \Phi_{t,t'}\left[\mathrm{d}\|Z_t^{i,N}\|^2 + 2\eta\gamma_t\|Z_t^{i,N}\|^2\mathrm{d}t\right] \\
&\leq \gamma_t\Phi_{t,t'}\left\langle Z_t^{i,N}, \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N}) - \nabla_\theta\underline{\tilde{\mathcal{L}}}(\theta_t^{i,N})\right\rangle\mathrm{d}t \\
&+ \gamma_t\Phi_{t,t'}\left\langle Z_t^{i,N}, \nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})\right\rangle\mathrm{d}t \\
&+ \gamma_t\Phi_{t,t'}\left\langle Z_t^{i,N}, \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\right\rangle \\
&+ \gamma_t^2\Phi_{t,t'}\left|\left|\nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\right|\right|_F^2\mathrm{d}t.
\end{aligned}$$

Finally, rewriting this in integral form, setting $t' = t$, and taking expectations, we arrive at

$$\begin{aligned}
\mathbb{E}_{\theta_0}\left[\|Z_t^{i,N}\|^2\right] &\leq \mathbb{E}_{\theta_0}\left[\Phi_{1,t}\|Z_1^{i,N}\|^2\right] \qquad (41)\\
&+ \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s\Phi_{s,t}\left\langle Z_s^{i,N}, \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N}) - \nabla_\theta\underline{\tilde{\mathcal{L}}}(\theta_s^{i,N})\right\rangle\mathrm{d}s\right] \\
&+ \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s\Phi_{s,t}\left\langle Z_s^{i,N}, \nabla_\theta L(\theta_s^{i,N}, x_s^{i,N}, \mu_s^N, \mu_s^N) - \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\right\rangle\mathrm{d}s\right] \\
&+ \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s^2\Phi_{s,t}\left|\left|\nabla_\theta B(\theta_s^{i,N}, x_s^{i,N}, \mu_s^N)\right|\right|_F^2\mathrm{d}s\right] \\
&:= \mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(1)}\right] + \mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(2)}\right] + \mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(3)}\right] + \mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(4)}\right]. \qquad (42)
\end{aligned}$$

We will deal with each of these terms separately, beginning with $\Omega_{t,i,N}^{(1)}$. For this term, by Lemma 18 (the moment bounds for $\theta_s^{i,N}$), and Condition F.1 (the conditions on the learning rate), for sufficiently large t we have that

$$\mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(1)}\right] = \Phi_{1,t}\mathbb{E}_{\theta_0}\left[\|Z_1^{i,N}\|^2\right] \leq K_{\theta_0}^{(1)}\gamma_t. \qquad (43)$$

We now turn our attention to $\Omega_{t,i,N}^{(2)}$. For this term, using Lemmas 1, 2, 3, and 4, we have that, for all $\theta \in \mathbb{R}^p$,

$$\begin{aligned}
\|\nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta) - \nabla_\theta\underline{\tilde{\mathcal{L}}}(\theta)\| &= \lim_{t\to\infty}\|\nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta) - \nabla_\theta\underline{\tilde{\mathcal{L}}}(\theta)\| \\
&\leq \lim_{t\to\infty}\Big[\|\nabla_\theta\tilde{\underline{\mathcal{L}}}(\theta) - \tfrac{1}{t}\nabla_\theta\underline{\mathcal{L}}_t^i(\theta)\| + \|\tfrac{1}{t}\nabla_\theta\underline{\mathcal{L}}_t^i(\theta) - \tfrac{1}{t}\nabla_\theta\mathcal{L}_t^{[i,N]}(\theta)\| \\
&\quad + \|\tfrac{1}{t}\nabla_\theta\mathcal{L}_t^{[i,N]}(\theta) - \tfrac{1}{t}\nabla_\theta\mathcal{L}_t^{i,N}(\theta)\| + \|\tfrac{1}{t}\nabla_\theta\mathcal{L}_t^{i,N}(\theta) - \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta)\|\Big] \\
&\leq K_{\theta_0,1}\alpha(N) + \frac{K_{\theta_0,2}}{N^{\frac{1}{2}}}, \qquad (44)
\end{aligned}$$

where $\alpha : \mathbb{N} \to \mathbb{R}_+$ is the function defined in (19). Substituting this bound into (41), and using Condition F.1 (the conditions on the learning rate) to bound the final integral, we obtain

$$\begin{aligned}
\mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(2)}\right] &\leq \int_1^t \gamma_s\Phi_{s,t}\mathbb{E}_{\theta_0}\left[\|Z_s^{i,N}\|\sup_{\theta_s^{i,N}}\|\nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N}) - \nabla_\theta\underline{\tilde{\mathcal{L}}}(\theta_s^{i,N})\|\right]\mathrm{d}s \\
&\leq \left[K_{\theta_0,1}\alpha(N) + \frac{K_{\theta_0,2}}{N^{\frac{1}{2}}}\right]\int_1^t \gamma_s\Phi_{s,t}\mathrm{d}s \leq K_{\theta_0}^{(2,1)}\alpha(N) + \frac{K_{\theta_0}^{(2,2)}}{N^{\frac{1}{2}}}. \qquad (45)
\end{aligned}$$

We now turn our attention to $\Omega_{t,i,N}^{(3)}$. We will analyse this term by constructing an appropriate Poisson equation. Let us define $R^{i,N}(\theta, \hat{x}^N) = \langle\theta - \theta_0, \nabla_\theta\hat{L}^{i,N}(\theta, \hat{x}^N) - \nabla_\theta\tilde{\mathcal{L}}^{i,N}(\theta)\rangle$, where, as previously, $\hat{x}^N = (x^{1,N}, \ldots, x^{N,N})$. Due to Condition C.1, Lemma 7 (the boundedness

of the asymptotic log-likelihood and its derivatives), and Lemma 14 (the local Lipschitz and polynomial growth of $\hat{L}^{i,N}(\theta, \hat{x}^N)$ and its derivatives), this function satisfies all of the conditions of Lemmas 16 and 17 with $r = 1$. Thus, the Poisson equation

$$\mathcal{A}_x v^{i,N}(\theta, \hat{x}^N) = R^{i,N}(\theta, \hat{x}^N) \quad , \quad \int_{\mathbb{R}^d} v^{i,N}(\theta, \hat{x}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) = 0,$$

has a unique twice differentiable solution which satisfies

$$\sum_{k=0}^{2} \| \frac{\partial^k}{\partial \theta^k} v^{i,N}(\theta, \hat{x}^N)\| + \| \frac{\partial^2}{\partial \theta \partial \hat{x}^N} v^{i,N}(\theta, \hat{x}^N)\| \\ \leq K_{\theta_0}\Big[1 + \|\theta\|\Big]\Big[1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N} \|x^{j,N}\|^{q'}\Big].$$

Now, by Itô's formula, we have that

$$\begin{aligned} v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N) - v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) = &\int_s^t \mathcal{A}_\theta v^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\mathrm{d}u + \int_s^t \mathcal{A}_{\hat{x}^N} v^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\mathrm{d}u \\ &+ \int_s^t \gamma_u \partial_\theta v^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\nabla_\theta \hat{B}^{i,N}(\theta_u, \hat{x}_u^N)\mathrm{d}w_u^i \\ &+ \int_s^t \partial_x v^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\mathrm{d}\hat{w}_u^N \\ &+ \int_s^t \gamma_u \Big[\partial_\theta \partial_{\hat{x}} v^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\nabla_\theta \hat{B}^{i,N}(\theta_u^{i,N}, \hat{x}_u^N)\Big]\mathrm{d}u, \end{aligned}$$

where $\hat{w}_u^N$ was defined in (21). It follows, now writing $v_t^{i,N} := v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)$, that

$$\begin{aligned} R^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\mathrm{d}t &= \mathcal{A}_{\hat{x}^N} v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\mathrm{d}t \\ &= \mathrm{d}v_t^{i,N} - \mathcal{A}_\theta v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\mathrm{d}t \\ &- \gamma_t \partial_\theta v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\nabla_\theta \hat{B}^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\mathrm{d}w_t^i \\ &- \partial_{\hat{x}} v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\mathrm{d}\hat{w}_t^N \\ &- \gamma_t \Big[\partial_\theta \partial_{\hat{x}} v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\nabla_\theta \hat{B}^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\Big]\mathrm{d}t. \end{aligned}$$

Thus, we can rewrite $\Omega^{(3)}_{t,i,N}$ as

$$\begin{aligned} \Omega^{(3)}_{t,i,N} &= \int_1^t \gamma_s \Phi_{s,t} \underbrace{\Big\langle \theta_s^{i,N} - \theta_0, \nabla_\theta \hat{L}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\Big\rangle \mathrm{d}s}_{R^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s} \\ &= \int_1^t \gamma_s \Phi_{s,t}\mathrm{d}v_s^{i,N} - \int_1^t \gamma_s \Phi_{s,t}\mathcal{A}_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \\ &- \int_1^t \gamma_s^2 \Phi_{s,t}\partial_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}w_s^i - \int_1^t \gamma_s \Phi_{s,t}\partial_{\hat{x}} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}\hat{w}_s^N \\ &- \int_1^t \gamma_s^2 \Phi_{s,t}\partial_\theta \partial_x v^{i,N}(\theta_s, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s. \end{aligned} \tag{46}$$

We can rewrite the first term in this expression by applying Itô's formula to $f(s, v_s) = \gamma_s \Phi_{s,t} v_s$. This yields

$$\gamma_t \Phi_{t,t} v_t^{i,N} - \gamma_1 \Phi_{1,t} v_1^{i,N} = \int_1^t \gamma_s \Phi_{s,t} \mathrm{d}v_s^{i,N} + \int_1^t \dot{\gamma}_s \Phi_{s,t} v_s^{i,N} \mathrm{d}s + \int_1^t 2\eta\gamma_s^2 \Phi_{s,t} v_s^{i,N} \mathrm{d}s.$$

Substituting the resulting expression for $\int_1^t \gamma_s \Phi_{s,t} \mathrm{d}v_s^{i,N}$ into (46), and taking expectations, we obtain

$$\begin{aligned}
\mathbb{E}_{\theta_0}\left[\Omega_{t,i,N}^{(3)}\right] &= \mathbb{E}_{\theta_0}\left[\gamma_t \Phi_{t,t} v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\right] - \mathbb{E}_{\theta_0}\left[\gamma_1 \Phi_{1,t} v^{i,N}(\theta_1^{i,N}, \hat{x}_1^N)\right] \\
&- \mathbb{E}_{\theta_0}\left[\int_1^t \dot{\gamma}_s \Phi_{s,t} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s\right] - \mathbb{E}_{\theta_0}\left[\int_1^t 2\eta\gamma_s^2 \Phi_{s,t} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s\right] \\
&- \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s \Phi_{s,t} \mathcal{A}_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s\right] \\
&- \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s^2 \Phi_{s,t} \partial_\theta \partial_x v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s\right] \\
&\leq K_{\theta_0}\left[\gamma_t + \int_1^t \left(\dot{\gamma}_s + \gamma_s^2\right) \Phi_{s,t} \mathrm{d}s\right] \leq K_{\theta_0}^{(3)} \gamma_t,
\end{aligned} \tag{47}$$

where in the penultimate inequality we have used the polynomial growth of $v^{i,N}(\theta, \hat{x}^N)$ and $\partial_\theta \partial_x v^{i,N}(\theta, \hat{x}^N)$, Condition C.1 (which implies the polynomial growth of $\nabla_\theta \hat{B}^{i,N}(\theta, \hat{x}^N)$), Proposition 1 (the moment bounds for $\hat{x}_t^N$), Lemma 18 (the moment bounds for $\theta_s^{i,N}$), and in the final inequality we have used Condition F.1 (the conditions on the learning rate). It remains only to bound $\Omega_{t,i,N}^{(4)}$. For this term, using the same assumptions, we obtain

$$\mathbb{E}\left[\Omega_{t,i,N}^{(4)}\right] = \mathbb{E}_{\theta_0}\left[\int_1^t \gamma_s^2 \Phi_{s,t} \left|\left|\nabla_\theta B(\theta_s, x_s, \mu_s)\right|\right|_F^2 \mathrm{d}s\right] \leq K_{\theta_0} \int_1^t \gamma_s^2 \Phi_{s,t} \mathrm{d}s \leq K_{\theta_0}^{(4)} \gamma_t. \tag{48}$$

Combining inequalities (43), (45), (47), and (48), and setting $K_{\theta_0,1}^\dagger = \max\{K_{\theta_0}^{(1)}, K^{(3)}\}$, $K_{\theta_0,2}^\dagger = K_{\theta_0}^{(4)}$, $K_{\theta_0,3}^\dagger = K_{\theta_0}^{(2,1)}$, and $K_{\theta_0,4}^\dagger = K_{\theta_0}^{(2,2)}$, we thus have that

$$\mathbb{E}_{\theta_0}\left[\|\theta_t^{i,N} - \theta_0\|^2\right] \leq (K_{\theta_0,1}^\dagger + K_{\theta_0,2}^\dagger)\gamma_t + K_{\theta_0,3}^\dagger \alpha(N) + \frac{K_{\theta_0,4}^\dagger}{N^{\frac{1}{2}}}, \tag{49}$$

which completes the proof of the first statement of the theorem. Let us now turn our attention to the second statement. The proof of this bound goes through almost verbatim. Let us briefly highlight the main points of difference. To begin, we now have the following decomposition of the parameter update equation

$$\begin{aligned}
\mathrm{d}\theta_t^N &= \gamma_t \nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^N)\mathrm{d}t + \gamma_t \frac{1}{N}\sum_{i=1}^N \left(\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^N) - \nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^N)\right)\mathrm{d}t \\
&+ \gamma_t \frac{1}{N}\sum_{i=1}^N \left(\nabla_\theta L(\theta_t^N, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^N)\right)\mathrm{d}t \\
&+ \gamma_t \frac{1}{N}\sum_{i=1}^N \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i.
\end{aligned}$$

Using a first order Taylor expansion around θ_0, defining $Z_t^N = \theta_t^N - \theta_0$, applying Itô's formula to the function $\|Z_t^N\|^2$, and using the strong concavity of $\tilde{\mathcal{L}}(\theta)$, as in (39)–(40), we obtain

$$\begin{aligned}
\mathrm{d}\|Z_t^N\|^2 + 2\eta\gamma_t\|Z_t^N\|^2\mathrm{d}t \leq\ & \gamma_t \frac{1}{N}\sum_{i=1}^{N}\left\langle Z_t^N, \nabla_\theta \tilde{\mathcal{L}}^N(\theta_t^N) - \nabla_\theta \underline{\tilde{\mathcal{L}}}(\theta_t^N)\right\rangle \mathrm{d}t \\
&+ \gamma_t \frac{1}{N}\sum_{i=1}^{N}\left\langle Z_t^N, \nabla_\theta L(\theta_t^N, x_t^{i,N}, \mu_t^N, \mu_t^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_t^N)\right\rangle \mathrm{d}t \\
&+ \gamma_t \frac{1}{N}\sum_{i=1}^{N}\left\langle Z_t^N, \nabla_\theta B(\theta_t^N, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\right\rangle \\
&+ \gamma_t^2 \frac{1}{N^2}\sum_{i=1}^{N}\left|\left|\nabla_\theta B(\theta_t^N, x_t^{i,N}, \mu_t^N)\right|\right|_F^2 \mathrm{d}t.
\end{aligned}$$

Continuing to follow our previous arguments, we finally arrive at

$$\mathbb{E}_{\theta_0}\left[\|Z_t^N\|^2\right] \leq \frac{1}{N}\sum_{i=1}^{N}\left[\mathbb{E}_{\theta_0}\left[\tilde{\Omega}_{t,i,N}^{(1)}\right] + \mathbb{E}_{\theta_0}\left[\tilde{\Omega}_{t,i,N}^{(2)}\right] + \mathbb{E}_{\theta_0}\left[\tilde{\Omega}_{t,i,N}^{(3)}\right]\right] + \frac{1}{N^2}\sum_{i=1}^{N}\mathbb{E}_{\theta_0}\left[\tilde{\Omega}_{t,i,N}^{(4)}\right]$$

where, modulo changes in the arguments from $\theta_t^{i,N}$ to θ_t^N, $\tilde{\Omega}_{t,i,N}^{(1)}, \ldots, \tilde{\Omega}_{t,i,N}^{(4)}$ are identical to $\Omega_{t,i,N}^{(1)}, \ldots, \Omega_{t,i,N}^{(4)}$ as defined in (41)–(42). We thus have, using (43), (45), (47), and (48), that

$$\begin{aligned}
\mathbb{E}_{\theta_0}\left[\|Z_t^N\|^2\right] &\leq \frac{1}{N}\sum_{i=1}^{N}\left(K_{\theta_0,1}^{\dagger}\gamma_t + K_{\theta_0,3}^{\dagger}\alpha(N) + \frac{K_{\theta_0,4}^{\dagger}}{N^{\frac{1}{2}}}\right) + \frac{1}{N^2}\sum_{i=1}^{N}K_{\theta_0,2}^{\dagger}\gamma_t \\
&= (K_{\theta_0,1}^{\dagger} + \frac{K_{\theta_0,2}^{\dagger}}{N})\gamma_t + K_{\theta_0,3}^{\dagger}\alpha(N) + \frac{K_{\theta_0,4}^{\dagger}}{N^{\frac{1}{2}}}. \qquad (50)
\end{aligned}$$

It remains to outline how to adapt this proof to obtain the results in Theorems 1.2, 1.2*, and 2.2. Let us start with Theorem 2.2. In this case, we replace $\underline{\tilde{\mathcal{L}}}(\cdot)$ by $\tilde{\mathcal{L}}^{i,N}(\cdot)$ in (38), use a first order Taylor expansion for $\tilde{\mathcal{L}}^{i,N}(\cdot)$ instead of $\underline{\tilde{\mathcal{L}}}(\cdot)$ in (39), and the strong concavity of $\tilde{\mathcal{L}}^{i,N}(\cdot)$ (Condition E.1) rather than $\underline{\tilde{\mathcal{L}}}(\cdot)$ (Condition E.2) in (40). Following the subsequent arguments, we see that the second term $\Omega_{t,i,N}^{(2)}$ in (42) vanishes, with all other terms unchanged. The result is that, in (49) and (50), only the first two terms appear. This completes the proof of Theorem 2.2.

We now turn our attention to Theorem 1.2*. In this case, we replace $x_t^{i,N}$ by x_t^i, μ_t^N by $\mu_t^{[N]}$, $\theta_t^{i,N}$ by $\theta_t^{[i,N]}$, θ_t^N by $\theta_t^{[N]}$, $\mathcal{L}_t^{i,N}(\cdot)$ by $\mathcal{L}_t^{[i,N]}(\cdot)$, $\mathcal{L}_t^N(\cdot)$ by $\mathcal{L}_t^{[N]}(\cdot)$, $\tilde{\mathcal{L}}^{i,N}(\cdot)$ by $\tilde{\mathcal{L}}^{[i,N]}(\cdot)$, and $\tilde{\mathcal{L}}^N(\cdot)$ by $\tilde{\mathcal{L}}^{[N]}(\cdot)$. The arguments in our proof now go through essentially unchanged. The only material difference appears in the derivation of (44), where the third term vanishes. The result is that only the second term in (44) and hence (45) vanishes. Thus, in particular, the final term in (49) and (50) no longer appears. This completes the proof of Theorem 1.2*. Finally, Theorem 1.2 can be deduced from Theorem 1.2* in the same way that Theorem 2.2 was obtained from Theorem 2.2* (see above). □

5. Numerical examples

To illustrate the results of Section 3, we now provide two illustrative numerical examples. We will focus on Case II, in which we assume that we observe the trajectories of N interacting particles $(x_t^{i,N})_{t\geq 0}^{i=1,\ldots,N}$ from the IPS. In both examples, we implement the online estimators using a standard Euler–Maruyama scheme with $\Delta t = 0.1$.

5.1. Linear mean field dynamics

We begin by considering a one-dimensional linear mean field model, parametrised by $\theta = (\theta_1, \theta_2)^T \in \mathbb{R}^2$, which is defined according to

$$\mathrm{d}x_t^\theta = -\left[\theta_1 x_t^\theta + \theta_2 \int_{\mathbb{R}} (x_t^\theta - y)\mu_t^\theta(\mathrm{d}y)\right]\mathrm{d}t + \sigma \mathrm{d}w_t,$$

where $\sigma > 0$, $w = (w_t)_{t\geq 0}$ is a standard Brownian motion, and $x_0 \in \mathbb{R}$. This is clearly of the form of the McKean–Vlasov SDE (1) with $b(\theta, x) = -\theta_1 x$ and $\phi(\theta, x, y) = -\theta_2(x-y)$. In this model, the parameter θ_1 controls the strength of attraction of the non-linear process (or, in the IPS, of each individual particle) towards zero, while the strength of the parameter θ_2 controls the strength of the attraction of the non-linear process towards its mean (or, in the IPS, of each individual particle towards the empirical mean).

We will consider the case in which both parameters are unknown, and to be estimated from the data.[2] In this case, one can show that this model satisfies Conditions A.1, B.1–B.2, C.1, D.1–D.2, E.1 (but not Condition E.2), and F.1–F.2 (see Appendix C). Thus, the conditions of Theorem 2.1, Theorem 2.2 and Theorem 2.1* are satisfied, but those of Theorem 2.2* are not. In practice, this means that the log-likelihood of IPS admits a unique maximiser θ_0 and, for fixed $N \in \mathbb{N}$, the online parameter estimate is guaranteed to converge to θ_0 as $t \to \infty$. On the other hand, the asymptotic log-likelihood of the McKean–Vlasov SDE may not admit a unique maximiser, and thus the online parameter estimate may not convergence to θ_0 as $t \to \infty$ and $N \to \infty$.

In order to estimate $\theta_0 \in \mathbb{R}^2$, we will implement the online parameter update equations in (15). For this model, these update equations are given by

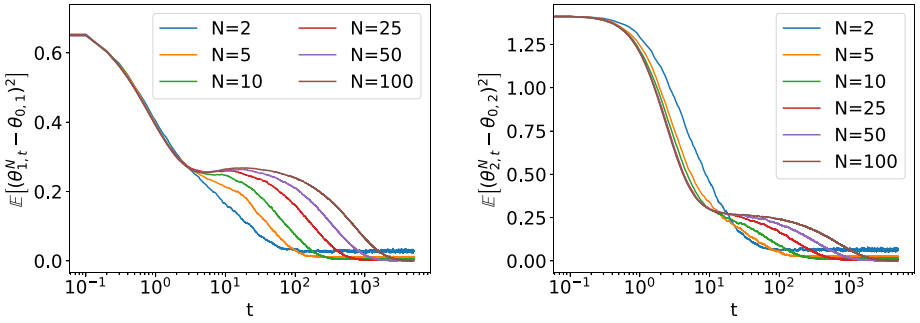
$$\mathrm{d}\theta_{1,t}^N = \frac{\gamma_{1,t}}{N\sigma^2}\sum_{i=1}^N \left[-x_t^{i,N}\mathrm{d}x_t^{i,N} - x_t^{i,N}(\theta_{1,t}^N x_t^{i,N} + \theta_{2,t}^N(x_t^{i,N} - \bar{x}_t^N))\mathrm{d}t\right],$$

$$\mathrm{d}\theta_{2,t}^N = \frac{\gamma_{2,t}}{N\sigma^2}\sum_{i=1}^N \left[-(x_t^{i,N} - \bar{x}_t^N)\mathrm{d}x_t^{i,N} - (x_t^{i,N} - \bar{x}_t^N)(\theta_{1,t}^N x_t^{i,N} + \theta_{2,t}^N(x_t^{i,N} - \bar{x}_t^N))\mathrm{d}t\right].$$

We will assume that the true parameter is given by $\theta_0 = (\theta_{0,1}, \theta_{0,2})^\top = (0.5, 0.1)^\top$. The initial parameter estimate $\theta_{\text{init}} = (\theta_{\text{init},1}, \theta_{\text{init},2})^\top$ is generated according to $\theta_{\text{init},1} \sim \mathcal{U}(-1, 2)$ and $\theta_{\text{init},2} \sim \mathcal{U}(-2, 2)$. Finally, we use constant learning rates, with $\gamma_{1,t} = 0.1$ and $\gamma_{2,t} = 0.2$. The performance of the stochastic gradient descent algorithm is illustrated in Fig. 1, in which we plot the mean-squared error (MSE) of the online parameter estimates for both of the unknown parameters, averaged over 500 random trials.

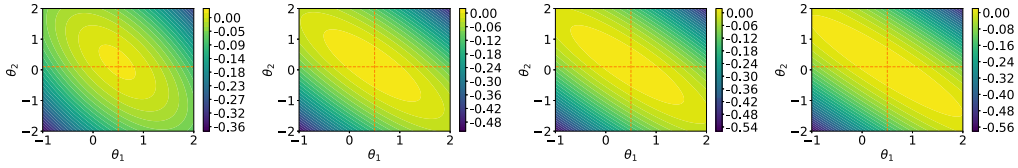
In this case, the evolution of the MSE indicates three distinct learning phases. In the initial phase, the performance of the online estimator improves as a function of the number of particles, with this improvement being more noticeable for the interaction parameter θ_2. In the middle phase, the online estimator performs significantly better for smaller values of N. In the final learning phase, the MSE of the online parameter estimate decreases as a function of the number of particles.

[2] Additional simulations for case in which one of the parameters is fixed are provided in the extended version of this paper [93].



(a) $\theta^N_{1,t}$. (b) $\theta^N_{2,t}$.

Fig. 1. MSE of the online parameter estimates for $T \in [0, 5000]$ and $N = \{2, 5, 10, 25, 50, 100\}$.



(a) $N = 2$. (b) $N = 5$. (c) $N = 10$. (d) $N = 100$.

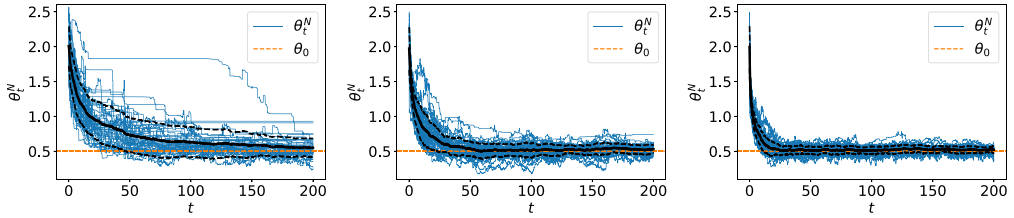
Fig. 2. Contour plots of the asymptotic log-likelihood $\tilde{\mathcal{L}}^N(\theta)$ for $N = \{2, 5, 10, 100\}$.

These observations are readily explained by considering the asymptotic log-likelihood of the IPS for different numbers of particles (see Fig. 2), and our theoretical results. Regarding the initial learning phase, we note that, far from the global maximum at $\theta_0 = (0.5, 0.1)$, the asymptotic log-likelihood decreases more steeply as N increases. Thus, in this phase, the online parameter estimate moves more quickly towards the global maximum for larger N. On the other hand, close to the global maximum at θ_0, the asymptotic log-likelihood exhibits an increasingly large plateau as N increases; that is, the maximum is increasingly 'flat'. In fact, in the mean field limit, the asymptotic log-likelihood does not even admit a unique maximum at θ_0, but instead is maximised by any $\theta = (\theta_1, \theta_2)^T$ on the line $\theta_1 + \theta_2 = \theta_{0,1} + \theta_{0,2}$ (see Appendix C). This region of the optimisation landscape is largely responsible for the middle learning phase, which explains the slower convergence of the estimator for larger numbers of particles. In the final learning phase, the MSE is governed by the results in Theorem 2.2, which show that the asymptotic $\mathbb{L}_2$ error of the online parameter estimate implemented here decreases with the number of particles.

5.2. *Stochastic opinion dynamics*

We now consider a stochastic opinion dynamics model, parametrised again by $\theta = (\theta_1, \theta_2)^\top \in \mathbb{R}^2$, defined according to

$$\mathrm{d}x^\theta_t = -\left[\int_{\mathbb{R}} \varphi_\theta(\|x^\theta_t - y\|)(x^\theta_t - y)\mu_t(\mathrm{d}y)\right]\mathrm{d}t + \sigma\,\mathrm{d}w_t,$$



(a) $N = 10$. (b) $N = 20$. (c) $N = 50$.

Fig. 3. Sequence on online parameter estimates (blue) for the range parameter θ_2, for 50 different random initialisations $\theta_2^0 \sim \mathcal{U}(1.5, 2.5)$, and $N = \{10, 20, 50\}$. We also plot the true parameter value (orange) and the mean online parameter estimate (black, solid) plus/minus one standard deviation (black, dashed).

where $\sigma > 0$, $w = (w_t)_{t\geq 0}$ is a standard Brownian motion, $x_0 \in \mathbb{R}$, and the interaction kernel $\varphi_\theta : \mathbb{R}_+ \to \mathbb{R}_+$ is given by the scaled indicator function $\varphi_\theta(r) = \theta_1 \mathbb{1}_{r\in[0,\theta_2]}$.[3]

In this model, we can interpret θ_1 as a scale parameter, which controls the strength of the attraction between particles, and θ_2 as a range parameter, which determines the range of the interaction. Models of this form arise in various applications, from biology to the social sciences, in which φ_θ determines how the dynamics of each particle (e.g., the opinions of each person) may influence the dynamics of other particles (e.g., the opinions of others). For a more detailed account of such models, we refer to [14,28,41,69,81] and references therein.

We provide illustrative results for the case in which the scale parameter θ_1 is fixed, and the range parameter θ_2 is to be estimated. While, in this case, we cannot verify all of the assumptions in Section 3, it is nonetheless interesting to consider whether our online estimators can recover the true parameter. We assume that $\theta_0 = (\theta_{0,1}, \theta_{0,2})^\top = (2, 0.5)^\top$. This corresponds to an interaction kernel with compact support on $[0, 0.5]$. The initial parameter estimates are generated uniformly at random on $[1.5, 2.5]$. Finally, we use constant learning rates with $\gamma_{2,t} = 0.002$. The performance of the recursive MLE is illustrated in Fig. 3 , in which we plot the sequence of online parameter estimates for θ_2, for several values of N, and for 50 different random initialisations. Encouragingly, almost all of the online parameter estimates converge to within a small neighbourhood of the true value of the parameter, suggesting that it is indeed possible to estimate the range of the interaction kernel in an online fashion. As before, the performance of the online estimator improves as the number of particles is increased. We should remark that, in this case, the performance is highly dependent on the initial conditions of the particles. This should not come as a surprise; indeed, if the distance between particles is greater than the support of the interaction kernel, then the interaction kernel (and its gradient) are identically zero, and thus so too are all of the terms in the parameter update equation. Thus, the value of the parameter estimate will remain unchanged. We see this phenomenon occurs for some of the trajectories in Fig. 3, particularly when there are fewer particles.

6. Conclusions

In this paper, we have studied the problem of online parameter estimation for a stochastic McKean–Vlasov equation and the associated system of weakly interacting particles. We

[3] In our numerics, we are required to use a differentiable approximation of this interaction kernel, and will use

$$\varphi_\theta(r) = \theta_1 \exp\left[-\frac{0.01}{1-(r-(\theta_2-1))^2}\right].$$

proposed an online estimator for the parameters of the stochastic McKean–Vlasov equation, based on observations of the trajectories of the associated IPS, or multiple independent replicates of the McKean–Vlasov SDE, and studied its convergence properties. In parallel work, we have also studied the asymptotic properties of the offline MLE (see [93] for some preliminary results in this direction).

Regarding directions for future research, it is of interest to extend our results to the case in which the diffusion coefficient is unknown, and must be estimated online (see [94] for online estimation of the diffusion coefficient in the linear case, and [52] for offline estimation of the diffusion coefficient in IPSs). This is a particularly interesting problem given that, for a broad class of McKean–Vlasov SDEs, the uniqueness (or non-uniqueness) of the invariant measure(s) is known to depend on the magnitude of the noise coefficient (e.g., [49]).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors are grateful to the two anonymous reviewers and the associate editor for their careful reading of the paper and their constructive feedback. The first author was funded by the Engineering and Physical Sciences Research Council (EPSRC) CDT in the Mathematics of Planet Earth (grant number EP/L016613/1) and the National Physical Laboratory. The second, third, and fourth authors were partially funded under a J.P. Morgan A.I. Research Award (2021). The third author and the fourth author were also supported by the EPSRC (grant numbers EP/W003317/1 and EP/P031587/1, respectively).

Appendix A. Existing results on the McKean–Vlasov SDE

In this appendix, we present some existing results on the McKean–Vlasov SDE (1), and the associated system of interacting particles (2). For completeness, we also include the proofs of these results. These results are well known, and have appeared in a number of existing papers (e.g., [11,26,72,73,102]), albeit under conditions which are different to our own. Under our particular set of assumptions, these results can also be found in [17].

Proposition 1 (*Moment Bounds*). *Assume that Conditions A.1 and B.1(i)–B.2(i) hold. Then, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$, and for all $q \geq 0$, there exists K_{θ_0} such that*

$$\sup_{t\geq 0} \mathbb{E}_{\theta_0} \|x_t\|^q \leq K_{\theta_0} \left(\int_{\mathbb{R}^d} \|x\|^q \mu_0(\mathrm{d}x) + 1 \right) \quad \textit{and}$$
$$\sup_{t\geq 0} \mathbb{E}_{\theta_0} \|x_t^{i,N}\|^q \leq K_{\theta_0} \left(\int_{\mathbb{R}^d} \|x\|^q \mu_0(\mathrm{d}x) + 1 \right). \tag{A.1}$$

Thus, in particular, for any $x \in \mathbb{R}^d$,

$$\sup_{t\geq 0} \mathbb{E}_{\theta_0,x} \|x_t\|^q \leq K_{\theta_0} \left(1 + \|x\|^q\right) \quad \textit{and} \quad \sup_{t\geq 0} \mathbb{E}_{\theta_0,x} \|x_t^{i,N}\|^q \leq K_{\theta_0} \left(1 + \|x\|^q\right). \tag{A.2}$$

Proof. We follow the approach in [17, Lemma 2.3.1] (see also [26, Corollary 2.3]). We will establish the first bound in (A.1), with the second bound obtained similarly. We begin by applying Itô's Lemma to $\|\cdot\|^{2q}$, and taking expectations, which yields

$$\mathbb{E}_{\theta_0}\left[\|x_t\|^{2q}\right] = \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] + \mathbb{E}_{\theta_0}\left[\int_0^t 2q\|x_s\|^{2q-2}\langle x_s, B(\theta_0, x_s, \mu_s)\rangle \mathrm{d}s\right] + [qd + 2q(q-1)]\mathbb{E}_{\theta_0}\left[\int_0^t \|x_s\|^{2q-2}\mathrm{d}s\right].$$

Due to Conditions B.1(i)–B.2(i), there exists $K_{\theta_0,1} > 0$ such that $\langle x_s, B(\theta_0, x_s, \mu_s)\rangle \leq -(\alpha_{\theta_0} - L_{\theta_0,2})\|x_s\|^2 + K_{\theta_0,1}\|x_s\| + L_{\theta_0,2}\|x_s\|\mathbb{E}_{\theta_0}[\|x_s\|]$. It follows, defining $K_1 = \frac{d}{2} + q - 1 > 0$, that

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\|x_t\|^{2q}\right] \leq \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] &- 2q\mathbb{E}_{\theta_0}\left[\int_0^t (\alpha_{\theta_0} - L_{\theta_0,2})\|x_s\|^{2q}\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t K_{\theta_0,1}\|x_s\|^{2q-1}\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t L_2\|x_s\|^{2q-1}\mathbb{E}_{\theta_0}[\|x_s\|]\,\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t K_1\|x_s\|^{2q-2}\mathrm{d}s\right].\end{aligned}$$

By Young's inequality, we then have

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\|x_t\|^{2q}\right] \leq \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] &- 2q\mathbb{E}_{\theta_0}\left[\int_0^t (\alpha_{\theta_0} - L_{\theta_0,2})\|x_s\|^{2q}\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t K_{\theta_0,1}\|x_s\|^{2q-1}\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t L_2\frac{2q-1}{2q}\|x_s\|^{2q}\mathrm{d}s\right] + 2q\mathbb{E}_{\theta_0}\left[\int_0^t L_2\frac{1}{2q}\|x_s\|^{2q}\mathrm{d}s\right] \\ &+ 2q\mathbb{E}_{\theta_0}\left[\int_0^t K_1\|x_s\|^{2q-2}\mathrm{d}s\right].\end{aligned}$$

It follows, writing $\lambda_{\theta_0} = \alpha_{\theta_0} - L_{\theta_0,2}$, that

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\|x_t\|^{2q}\right] &\leq \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] \\ &\quad + 2q\left[\int_0^t \mathbb{E}_{\theta_0}\left[-\lambda_{\theta_0}\|x_s\|^{2q} + K_{\theta_0,1}\|x_s\|^{2q-1} + K_1\|x_s\|^{2q-2}\right]\mathrm{d}s\right].\end{aligned}$$

We next use the fact that, for any $a > 0$, there exists b such that $K_{\theta_0,1}\|x\|^{2q-1} + K_1\|x\|^{2q-2} \leq a\|x\|^{2q} + b$. This implies, in particular, that there exists $K_{\theta_0,2} > 0$, $K_{\theta_0,3} > 0$ such that

$$\mathbb{E}_{\theta_0}\left[\|x_t\|^{2q}\right] \leq \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] + 2q\left[\int_0^t -K_{\theta_0,2}\mathbb{E}_{\theta_0}\left[\|x_s\|^{2q}\right] + K_{\theta_0,3}\mathrm{d}s\right].$$

The result now follows via an extension of Grönwall's inequality (see, e.g., [26, Proposition 2.1]). The second statement, namely (A.2), now follows immediately by setting $\mu_0(\cdot) = \delta_x(\cdot)$. □

Proposition 2 (*Asymptotic Moment Bounds*). *Assume that Conditions A.1 and B.1–B.2 hold. Then, for all $N \in \mathbb{N}$, for all $i = 1, \ldots, N$, and for all $q \in \mathbb{N}$, there exists a positive constant*

$K_{\theta_0} > 0$ such that

$$\int_{\mathbb{R}^d} \|x\|^q \mu_\infty(\mathrm{d}x) \leq K_{\theta_0} \quad , \quad \int_{(\mathbb{R}^d)^N} \|x^{i,N}\|^q \hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) \leq K_{\theta_0}.$$

Proof. By Proposition 3 (see below), the McKean–Vlasov SDE (1) admits a unique equilibrium measure μ_∞ which is independent of μ_0. By the ergodic theorem (e.g., [87, Chapter X]), we thus have, for all $k \in \mathbb{N}$,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \|x_s\|^q \mathrm{d}s = \int_{\mathbb{R}^d} \|x\|^q \mu_\infty(\mathrm{d}x) \ , \quad \text{a.s.} \tag{A.3}$$

Using Jensen's inequality and the moment bounds in Proposition 1, we obtain uniform integrability of the family $\{\frac{1}{t}\int_0^t \|x_s\|^q \mathrm{d}s\}_{t>0}$. It follows, taking expectations of (A.3), using uniform integrability in order to interchange the limit and the expectation, and once more making use of Proposition 1, that

$$\begin{aligned}\int_{\mathbb{R}^d} \|x\|^q \mu_\infty(\mathrm{d}x) &= \lim_{t\to\infty} \left[\frac{1}{t}\int_0^t \mathbb{E}_{\theta_0}\left[\|x_s\|^q\right] \mathrm{d}s\right] \\ &\leq \sup_{\mu_0 \in \mathcal{P}_k(\mathbb{R}^d)} K_{\theta_0} \left(\int_{\mathbb{R}^d} \|x\|^q \mu_0(\mathrm{d}x) + 1\right) \leq K_{\theta_0},\end{aligned}$$

where in the final inequality we have used Condition A.1 and, as elsewhere, allowed the constant K_{θ_0} to increase between inequalities. The proof of the bound for the IPS is identical, noting that all of the relevant results (Propositions 1 and 3) apply to both the McKean–Vlasov SDE and the IPS. □

Proposition 3 (*Unique Invariant Measure*)**.** *Assume that Conditions A.1 and B.1–B.2 hold. Then, for $\theta = \theta_0$, the McKean–Vlasov SDE (1) and IPS (2) admit unique invariant measures μ_∞ and $\hat{\mu}_\infty^N$ which are independent of the initial conditions μ_0 and $(\mu_0)^{\otimes N}$. Moreover, writing $\lambda_{\theta_0} = \alpha_{\theta_0} - 2L_{\theta_0,2}$, and $\hat{\mu}_t^{(k),N}$ for the law of a subset of $1 \leq k \leq N$ interacting particles, the following contraction rates hold*

$$\mathbb{W}_2(\mu_t, \mu_\infty) \leq e^{-\lambda_{\theta_0} t} \mathbb{W}_2(\mu_0, \mu_\infty) \quad , \quad \mathbb{W}_2(\hat{\mu}_t^{(k),N}, \hat{\mu}_\infty^{(k),N}) \leq e^{-\lambda_{\theta_0} t} \mathbb{W}_2((\mu_0)^{\otimes k}, \hat{\mu}_\infty^{(k),N}). \tag{A.4}$$

Proof. We follow the approach in [17, Theorem 2.3.3] (see also [103, Theorem 3.1]). We prove the first inequality in (A.4), with the second derived in a similar fashion. Let $(x_t^x)_{t\geq 0}$ and $(x_t^y)_{t\geq 0}$ be solutions of (1), starting from $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^d$, respectively. We will write μ_t for the law of x_t^x, and ν_t for the law of x_t^y. By Itô's formula, we have

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_{\theta_0}\left[\|x_t^x - x_t^y\|^2\right] = 2\mathbb{E}_{\theta_0}\left[\left\langle x_t^x - x_t^y, B(\theta_0, x_t^x, \mu_t) - B(\theta_0, x_t^y, \nu_t)\right\rangle\right].$$

Using Conditions B.1 and B.2, it follows that

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_{\theta_0}\left[\|x_t^x - x_t^y\|^2\right] \leq -2(\alpha_{\theta_0} - 2L_{\theta_0,2})\mathbb{E}_{\theta_0}\left[\|x_t^x - x_t^y\|^2\right],$$

and thus, writing $\lambda_{\theta_0} = \alpha_{\theta_0} - 2L_{\theta_0,2}$, that $\mathbb{E}_{\theta_0}[\|x_t^x - x_t^y\|^2] \leq e^{-2\lambda_{\theta_0} t}\|x - y\|^2$. Let π_0 be an arbitrary coupling of μ_0 and ν_0. By Condition A.1, we have that $\mu_0, \nu_0 \in \mathcal{P}_2(\mathbb{R}^d)$. We thus

have from the previous inequality that

$$\int_{\mathbb{R}^d} \mathbb{E}_{\theta_0}\left[\|x_t^x - x_t^y\|^2\right] \pi_0(\mathrm{d}x, \mathrm{d}y) \leq e^{-2\lambda_{\theta_0} t} \int_{\mathbb{R}^d} \|x - y\|^2 \pi_0(\mathrm{d}x, \mathrm{d}y).$$

It follows, taking the infimum over all coupling measures π_0, and taking square roots, that

$$\mathbb{W}_2(\mu_t, \nu_t) \leq e^{-\lambda_{\theta_0} t} \mathbb{W}_2(\mu_0, \nu_0). \tag{A.5}$$

We now establish the existence of an equilibrium measure for the McKean–Vlasov SDE. Let $(x_t)_{t\geq 0}$ denote a solution of (1) with initial law μ_0, and $(y_t)_{t\geq 0}$ the solution of (1) with initial law $\mathcal{L}(x_s)$, for some $s > 0$. Using (A.5), homogeneity, and Proposition 1, there exists $K > 0$ such that

$$\mathbb{W}_2(\mathcal{L}(x_t), \mathcal{L}(x_{t+s})) = \mathbb{W}_2(\mathcal{L}(x_t), \mathcal{L}(y_t)) \leq e^{-\lambda_{\theta_0} t} \mathbb{W}_2(\mu_0, \mathcal{L}(x_s)) \leq K_{\theta_0} e^{-\lambda_{\theta_0} t}.$$

Thus, $\mathcal{L}(x_t)$ converges to a unique limit μ_∞, which is independent of the initial measure. We will now demonstrate that this equilibrium measure is also the unique invariant measure of the McKean–Vlasov SDE. Let $0 < t_1 < t_2$. Consider the functional $P_{t_1} : \mathcal{P}(\mathbb{R}^d) \to \mathcal{P}(\mathbb{R}^d)$, defined such that $P_{t_1}(\mu)$ is the law of the solution of the McKean–Vlasov SDE (1) at time t_1 when the initial law is μ. Similarly, define $P_{t_2} : \mathcal{P}(\mathbb{R}^d) \to \mathcal{P}(\mathbb{R}^d)$. Following (A.5), each of these functionals admits a unique fixed point. We will denote these fixed points by $\bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$. Note that, by homogeneity of (1), that $P_{nt_i}(\bar{\mu}_i) = \bar{\mu}_i$ for all $n \in \mathbb{N}$. It follows, again using (A.5), that

$$\mathbb{W}_2(\bar{\mu}_1, \mu_\infty) = \mathbb{W}_2(P_{nt_1}(\bar{\mu}_1), \mu_\infty) \leq K_{\theta_0,1} e^{-\lambda_{\theta_0} n t_1},$$
$$\mathbb{W}_2(\bar{\mu}_2, \mu_\infty) = \mathbb{W}_2(P_{nt_2}(\bar{\mu}_2), \mu_\infty) \leq K_{\theta_0,2} e^{-\lambda_{\theta_0} n t_2}.$$

Taking the limit as $N \to \infty$, and using the existence of a unique limit, we have that $\bar{\mu}_1 = \bar{\mu}_2 = \mu_\infty$. Since t_1 was chosen arbitrarily, it follows that μ_∞ is the unique invariant measure in $[t_1, \infty)$. It remains to check that μ_∞ that this is also true in $(0, t_1)$. This is indeed the case. Let $(x_t)_{t\geq 0}$ be a solution of (1) starting from $\bar{\mu}_1$, and write $\mathcal{L}(x_t)$ for the law of this solution. By construction, we have that $\mathcal{L}(x_{t_1}) = \bar{\mu}_1 = \bar{\mu}_2 = \mu_\infty$. Let $(y_t)_{t\geq 0}$ be a solution of (1) starting from $\mu_{t_1} = \mu_\infty$, with law $\mathcal{L}(y_t) = \mathcal{L}(x_{t+t_1})$. Since μ_∞ is the unique invariant measure in $[t_1, \infty)$, and $\mathcal{L}(x_{t_1}) = \mu_\infty$, we must have $\mathcal{L}(x_{t+t_1}) = \mu_\infty$ for $t \in [0, \infty)$. But this implies that $\mathcal{L}(y_t) = \mathcal{L}(x_{t+t_1}) = \mu_\infty$ for all $t \in (0, \infty)$. This concludes the proof.

The bound in (A.4) now follows straightforwardly by setting $\nu_0 = \mu_\infty$ in (A.5), and using the fact that μ_∞ is an invariant measure. □

Proposition 4 (*Propagation of Chaos*). *Let $x^i = (x_t^i)_{t\geq 0}$ be N independent copies of the solution of (1) driven by independent Brownian motions $w^i = (w_t^i)_{t\geq 0}$. Assume that Conditions A.1 and B.1–B.2 hold. Then there exist $0 < K_{\theta_0} < \infty$, independent of time, such that*

$$\sup_{t\geq 0} \mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right] \leq \frac{K_{\theta_0}}{N}.$$

Proof. We follow the approach in [17, Lemma 2.4.1] (see also [72, Theorem 3.3], [26, Theorem 3.1]). Let $(x_t)_{t\geq 0}$ be a solution of (1), independent of $(x_t^i)_{t\geq 0}$ for all $i = 1, \ldots, N$. Using Itô's formula, and Condition B.1, we have

$$\mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right] \leq -2\alpha_{\theta_0} \int_0^t \mathbb{E}_{\theta_0}\left[\|x_s^i - x_s^{i,N}\|^2\right] \mathrm{d}s$$

$$+ 2\int_0^t \mathbb{E}_{\theta_0}\Big[\|x_s^{i,N} - x_s^i\| \underbrace{\|\frac{1}{N}\sum_{j=1}^n \phi(\theta_0, x_s^{i,N}, x_s^{j,N}) - \int_{\mathbb{R}^d}\phi(\theta_0, x_s^i, y)\mu_s(\mathrm{d}y)\|}_{\psi_s(N)}\Big].$$

Meanwhile, using Condition B.2, and the triangle inequality, we then have

$$\psi_s(N) \leq L_{\theta_0,2}\big[\|x_s^{i,N} - x_s^i\| + \frac{1}{N}\sum_{j=1}^N \|x_s^{j,N} - x_s^j\|\big] + \|\frac{1}{N}\sum_{j=1}^N \phi(\theta_0, x_s^i, x_s^j) - \int_{\mathbb{R}^d}\phi(\theta_0, x_s^i, y)\mu_s(\mathrm{d}y)\|.$$

Using this bound, the Holdër inequality, and that the particles are identically distributed, it follows that

$$\mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right] \leq -2(\alpha_{\theta_0} - 2L_{\theta_0,2})\int_0^t \mathbb{E}_{\theta_0}\left[\|x_s^i - x_s^{i,N}\|^2\right]\mathrm{d}s + 2\int_0^t \mathbb{E}_{\theta_0}[\|x_s^{i,N} - x_s^i\|^2]^{\frac{1}{2}}\mathbb{E}_{\theta_0}[\big|\big|\frac{1}{N}\sum_{j=1}^N \phi(\theta_0, x_s^i, x_s^j) - \int_{\mathbb{R}^d}\phi(\theta_0, x_s^i, y)\mu_s(\mathrm{d}y)\big|\big|^2]^{\frac{1}{2}}\mathrm{d}s.$$

Let $\xi_s^j = \phi(\theta_0, x_s^i, x_s^j) - \int_{\mathbb{R}^d}\phi(\theta_0, x_s^i, y)\mu_s(\mathrm{d}y)$. If $j \neq k$, one of j and k is not equal to i, and thus $\mathbb{E}_{\theta_0}\left[\langle \xi_s^j, \xi_s^k\rangle\right] = 0$, since the random variables $(x_s^j)_{j=1}^N$ are independent copies of the solution of (1) with law μ_s. It follows, using also Condition B.2, and the moment bounds in Proposition 1, that

$$\begin{aligned}
&\mathbb{E}_{\theta_0}[\big|\big|\frac{1}{N}\sum_{j=1}^N \phi(\theta_0, x_s^i, x_s^j) - \int_{\mathbb{R}^d}\phi(\theta_0, x_s^i, y)\mu_s(\mathrm{d}y)\big|\big|^2] \\
&\leq \frac{1}{N^2}\sum_{j=1}^N \mathbb{E}_{\theta_0}\left[\|\phi(\theta, x_s^i, x_s^j) - \int_{\mathbb{R}^d}\phi(\theta, x_s^i, y)\mu_s(\mathrm{d}y)\|^2\right] \\
&\leq \frac{L_{\theta_0,2}}{N^2}\sum_{j=1}^N \mathbb{E}_{\theta_0}\left[\int_{\mathbb{R}^d}\|x_s^j - y\|^2\mu_s(\mathrm{d}y)\right] \\
&\leq \frac{2L_{\theta_0,2}}{N^2}\sum_{j=1}^N \mathbb{E}_{\theta_0}\left[\|x_s^j\|^2\right] + \mathbb{E}_{\theta_0}\left[\|x_s\|^2\right] \leq \frac{K_{\theta_0,1}}{N}
\end{aligned}$$

for some $K_{\theta_0,1} > 0$. Thus, substituting this into our previous bound, and taking derivatives, we have

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right] \leq -2(\alpha_{\theta_0} - 2L_{\theta_0,2})\mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right] + \frac{K_{\theta_0,1}}{\sqrt{N}}\mathbb{E}_{\theta_0}\left[\|x_t^{i,N} - x_t^i\|^2\right]^{\frac{1}{2}}.$$

The result now follows from an application of Grönwall's inequality (see, e.g., [72, Theorem 3.1]). □

Appendix B. Additional lemmas for Theorems 1.1, 1.1*, 2.1 and 2.1*

Lemma 6. *Assume that Conditions A.1 and B.1–B.2 hold. Then, for all $N \in \mathbb{N}$, $i = 1, \dots, N$, for all $t \geq 0$, and for all $k \geq 1$, there exists $K_{\theta_0} > 0$ such that*

$$\mathbb{E}_{\theta_0}\left[\sup_{0\leq s\leq t} \|x_s\|^q\right] \leq K_{\theta_0} t^{\frac{1}{2}} \quad \text{and} \quad \mathbb{E}_{\theta_0}\left[\sup_{0\leq s\leq t} \|x_s^{i,N}\|^q\right] \leq K_{\theta_0} t^{\frac{1}{2}}.$$

Proof. We will prove the first claim (the proof of the second being essentially identical). By Itô's Lemma, we have

$$\begin{aligned}\|x_t\|^{2q} = \|x_0\|^{2q} &+ \int_0^t 2q\|x_s\|^{2q-2}\langle x_s, B(\theta_0, x_s, \mu_s)\rangle \mathrm{d}s \\ &+ \int_0^t q\|x_s\|^{2q-2}\mathrm{Tr}[I_d + (q-2)[x_s^i x_s^j]_{i,j=1}^d \|x_s\|^{-2}]\mathrm{d}s \\ &+ \int_0^t 2q\|x_s\|^{2q-2}\langle \mathrm{d}x_s, \mathrm{d}w_s\rangle.\end{aligned}$$

It follows, taking the supremum and taking expectations, that

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\sup_{0\leq s\leq t} \|x_t\|^{2q}\right] \leq \mathbb{E}_{\theta_0}\left[\|x_0\|^{2q}\right] &+ \underbrace{2q\int_0^t \mathbb{E}_{\theta_0}\left[\left|\|x_s\|^{2q-2}\langle x_s, B(\theta_0, x_s, \mu_s)\rangle\right|\right]\mathrm{d}s}_{\Pi_t^1} \\ &+ \underbrace{[qd + 2q(q-1)]\int_0^t \mathbb{E}_{\theta_0}\left[\|x_s\|^{2q-2}\right]\mathrm{d}s}_{\Pi_t^2} \\ &+ \underbrace{2q\mathbb{E}_{\theta_0}\left[\sup_{0\leq s\leq t}\int_0^t \|x_s\|^{2q-2}\langle x_s, \mathrm{d}w_s\rangle\right]}_{\Pi_t^3}. \qquad \text{(B.1)}\end{aligned}$$

We begin by bounding the first term. Due to Conditions B.1–B.2, there exist $K_{\theta_0,1} > 0$ such that

$$\begin{aligned}\langle x_s, B(\theta, x_s, \mu_s)\rangle &\leq -(\alpha_{\theta_0} - L_{\theta_0,2})\|x_s\|^2 + K_{\theta_0,1}\|x_s\| + L_{\theta_0,2}\|x_s\|\mathbb{E}_{\theta_0}[\|x_s\|] \\ &\leq K_{\theta_0,1}\left[\|x_s\|^2 + \|x_s\| + \|x_s\|\mathbb{E}_{\theta_0}[\|x_s\|]\right]\end{aligned}$$

where, as elsewhere, the constant is allowed to increase between inequalities. It follows straightforwardly, using the moment bounds given by Proposition 1, that

$$\Pi_t^1 \leq 2q\int_0^t K_{\theta_0,1}\left[\mathbb{E}_{\theta_0}\left[\|x_s\|^{2q}\right] + \mathbb{E}_{\theta_0}\left[\|x_s\|^{2q-1}\right] + \mathbb{E}_{\theta_0}\left[\|x_s\|^{2q-2}\right]\mathbb{E}_{\theta_0}[\|x_s\|]\right]\mathrm{d}s \leq K_{\theta_0,1}t, \qquad \text{(B.2)}$$

$$\Pi_t^2 \leq [qd + 2q(q-1)]\int_0^t \mathbb{E}_{\theta_0}\left[\|x_s\|^{2q-2}\right]\mathrm{d}s \leq K_{\theta_0,2}t, \qquad \text{(B.3)}$$

$$\Pi_t^3 \leq 2q\mathbb{E}_{\theta_0}\left[\int_0^t \|x_s\|^{4q-4}\|x_s\|^2\mathrm{d}s\right]^{\frac{1}{2}} \leq 2q\left[\int_0^t \mathbb{E}_{\theta_0}[\|x_s\|^{4q-2}]\mathrm{d}s\right]^{\frac{1}{2}} \leq K_{\theta_0,3}t^{\frac{1}{2}}, \qquad \text{(B.4)}$$

where we have used the Burkholder–Davis–Gundy inequality, and once more Proposition 1. Combining Eqs. (B.1), (B.2), (B.3), and (B.4), using the Hölder inequality, and suitably defining the constant K_{θ_0} in terms of $K_{\theta_0,1}$, $K_{\theta_0,2}$, and $K_{\theta_0,3}$, the conclusion follows. □

Lemma 7. *Assume that Conditions A.1, B.1–B.2, C.1 hold. Then, for $k = 0, 1, 2, 3$, there exists $K_{\theta_0} > 0$ such that, for all $\theta \in \mathbb{R}^p$, $\|\nabla_\theta^k \underline{\tilde{\mathcal{L}}}(\theta)\| \leq K_{\theta_0}$ and $\|\nabla_\theta^k \tilde{\mathcal{L}}^{i,N}(\theta)\| \leq K_{\theta_0}$.*

Proof. Using the definition of $\nabla_\theta^k \underline{\tilde{\mathcal{L}}}(\theta)$ (Lemma 1), the polynomial growth of $\nabla_\theta^k L(\theta, x, \mu, \mu)$ (from Lemma 14), and the finite moments of the invariant measure of the McKean–Vlasov SDE (Proposition 2), we have

$$\begin{aligned}\|\nabla_\theta^k \underline{\tilde{\mathcal{L}}}(\theta)\| &\leq \int_{\mathbb{R}^d} \|\nabla_\theta^k L(\theta, x, \mu_\infty, \mu_\infty)\| \mu_\infty(\mathrm{d}x) \\ &\leq K \int_{\mathbb{R}^d} \Big[1 + \|x\|^q + \Big[\int_{\mathbb{R}^d} \|y\|^q \mu_\infty(\mathrm{d}y)\Big]\Big] \mu_\infty(\mathrm{d}x) \leq K_{\theta_0}.\end{aligned}$$

The bound for $\nabla_\theta^k \tilde{\mathcal{L}}^{i,N}(\theta)$ follows identically, this time using the definition of $\nabla_\theta^k \tilde{\mathcal{L}}^{i,N}(\theta)$ (Lemma 2), the relevant bound on $\nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{x}^N)$ in Lemma 14, and the finite moments of the invariant measure of the IPS (Proposition 2). □

B.1. Additional lemmas for Lemma 3

Lemma 8. *Assume that Conditions A.1 and B.1–B.2 hold. For all Lipschitz functions φ, there exists $K_{\theta_0} > 0$ such that, for all $t \geq 0$, for all $N \in \mathbb{N}$,*

$$\mathbb{E}_{\theta_0}\left[\left\| \int_{\mathbb{R}^d} \varphi(y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^{i,N})\right\|^2\right] \leq \frac{K_{\theta_0}}{N}.$$

Proof. Let x_t^i, $i = 1, \ldots, N$ denote independent solutions of the McKean–Vlasov SDE (1). We then have, using the elementary inequality $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$, that

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\left|\left| \int_{\mathbb{R}^d} \varphi(y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^{i,N})\right|\right|^2\right] \leq& 2\mathbb{E}_{\theta_0}\left[\left|\left| \int_{\mathbb{R}^d} \varphi(y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^i)\right|\right|^2\right] \\ &+ 2\mathbb{E}_{\theta_0}\left[\left|\left| \frac{1}{N}\sum_{i=1}^N \left(\varphi(x_t^i) - \varphi(x_t^{i,N})\right)\right|\right|^2\right].\end{aligned}$$

For the first term, using independence of $(x_t^i)_{t\geq 0}^{i=1,\ldots,N}$, the Lipschitz property of φ, and Proposition 1 (bounded moments of the McKean–Vlasov SDE), we have that

$$\begin{aligned}\mathbb{E}_{\theta_0}\left[\left|\left| \int_{\mathbb{R}^d} \varphi(y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^i)\right|\right|^2\right] &\leq \frac{1}{N}\mathbb{E}_{\theta_0}\left[\left|\left|\varphi(x_t^1) - \mathbb{E}[\varphi(x_t^1)]\right|\right|^2\right] \\ &\leq \frac{1}{N}\mathbb{E}_{\theta_0}\left[\left|\left|\varphi(x_t^1) - \varphi(\mathbb{E}[x_t^1])\right|\right|^2\right] \leq \frac{K_{\theta_0}}{N},\end{aligned}$$

for some constant $K_{\theta_0} > 0$. For the second term, using the Cauchy–Schwarz inequality, the Lipschitz property of φ, and Proposition 4 (uniform-in-time propagation of chaos), we obtain

$$\mathbb{E}_{\theta_0}\Big[\Big|\Big|\frac{1}{N}\sum_{i=1}^N \big(\varphi(x_t^i) - \varphi(x_t^{i,N})\big)\Big|\Big|^2\Big] \leq \frac{K_{\theta_0}}{N}\sum_{i=1}^N \mathbb{E}_{\theta_0}\big[\big|\big|x_t^i - x_t^{i,N}\big|\big|^2\big] \leq \frac{K_{\theta_0}}{N}.$$

The result follows immediately. □

Lemma 9. *Assume that Conditions A.1 and B.1–B.2 hold. Suppose also that $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$. Let x_t^i denote a solution of the McKean–Vlasov SDE, driven by $w^i = (w_t^i)_{t\geq 0}$. Then, for all Lipschitz functions φ, there exists $K_{\theta_0} > 0$ such that, for all $t \geq 0$, for all $N \in \mathbb{N}$,*

$$\mathbb{E}_{\theta_0}\left[\left\| \int_{\mathbb{R}^d} \varphi(x_t^i, y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^{i,N}, x_t^{j,N})\right\|^2\right] \leq \frac{K_{\theta_0}}{N}.$$

Proof. The is an immediate corollary of Lemma 8. Using the Hölder inequality, and the assumption that the function φ is Lipschitz, we have

$$\begin{aligned}
&\left|\left| \int_{\mathbb{R}^d} \varphi(x_t^i, y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{j=1}^N \varphi(x_t^{i,N}, x_t^{j,N})\right|\right|^2 \\
&\quad \leq 2\left|\left| \int_{\mathbb{R}^d} \varphi(x_t^i, y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{j=1}^N \varphi(x_t^i, x_t^{j,N})\right|\right|^2 + \frac{2K_{\theta_0}}{N}\sum_{j=1}^N \left|\left|x_t^i - x_t^{i,N}\right|\right|^2.
\end{aligned}$$

We thus have, applying the triangle inequality, the bound from Lemma 8, and Proposition 4, that

$$\begin{aligned}
&\mathbb{E}_{\theta_0}\left[\left\| \int_{\mathbb{R}^d} \varphi(x_t^i, y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{i=1}^N \varphi(x_t^{i,N}, x_t^{j,N})\right\|^2\right] \\
&\leq 2\underbrace{\mathbb{E}_{\theta_0}\left[\left|\left| \int_{\mathbb{R}^d} \varphi(x_t^i, y)\mu_t(\mathrm{d}y) - \frac{1}{N}\sum_{j=1}^N \varphi(x_t^i, x_t^{j,N})\right|\right|^2\right]}_{\leq \frac{K_{\theta_0}}{N} \text{ by Lemma 8}} + \frac{2K_{\theta_0}}{N}\underbrace{\mathbb{E}_{\theta_0}\left[\sum_{j=1}^N \left|\left|x_t^i - x_t^{i,N}\right|\right|^2\right]}_{\leq \frac{K_{\theta_0}}{N} \text{ by Proposition 4}} \\
&\leq \frac{K_{\theta_0}}{N},
\end{aligned}$$

where, as elsewhere, the value of the constant K_{θ_0} has been allowed to increase between inequalities. □

B.2. Additional lemmas for Lemma 5

B.2.1. Main lemmas

The lemmas in this section are modified versions of Lemmas 3.1–3.5 in [94]. We provide the proofs of these results in full, appropriately adapted to the current setting.

Lemma 10. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Define, with $\hat{x}^N = (x^{1,N}, \ldots, x^{N,N})$,*

$$\Gamma_{k,\eta} = \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \left(\nabla_\theta \hat{L}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\right) \mathrm{d}s.$$

Then, almost surely, $\|\Gamma_{k,\eta}\| \to 0$ as $k \to \infty$.

Proof. Let $\hat{x}^N = (x^{1,N}, \dots, x^{N,N}) \in (\mathbb{R}^d)^N$. Consider the function $S^{i,N}(\theta, \hat{x}^N) = \nabla_\theta \hat{L}^{i,N}(\theta, \hat{x}^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta)$. We begin by noting that this function is 'centred' with respect to the invariant measure $\hat{\mu}_\infty(\cdot)$, using the definition of $\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\cdot)$ from Lemma 2. In addition, due to Lemma 7 (the boundedness of the asymptotic log-likelihood and its derivatives) and Lemma 14 (the local Lipschitz and polynomial growth of $\hat{L}^{i,N}(\theta, \hat{x}^N)$ and its derivatives), there exist positive constants $q, K > 0$ (which are independent of N), such that, for $k = 0, 1, 2$,

$$\|\partial_\theta^k S^{i,N}(\theta, \hat{x}^N) - \partial_\theta^k S^{i,N}(\theta, \hat{y}^N)\| \leq K\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}}\right] \cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|y^{j,N}\|^q\right].$$

Thus, the function $S^{i,N}(\theta, \hat{x}^N)$ satisfies the conditions of Lemmas 16 and 17 with $r = 0$. It follows that, for all $i = 1, \dots, N$, the Poisson equation

$$\mathcal{A}_{\hat{x}} v^{i,N}(\theta, \hat{x}^N) = S^{i,N}(\theta, \hat{x}^N) \quad , \quad \int_{(\mathbb{R}^d)^N} v^{i,N}(\theta, \hat{x}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) = 0$$

has a unique twice differentiable solution which satisfies

$$\sum_{k=0}^2 \|\frac{\partial^k v^{i,N}}{\partial \theta^k}(\theta, \hat{x}^N)\| + \|\frac{\partial^2 v^{i,N}}{\partial\theta\partial\hat{x}^N}(\theta, \hat{x}^N)\| \leq K_{\theta_0}(1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^{q'}), \quad \text{(B.5)}$$

where the constants $q', K_{\theta_0} > 0$ are also independent of N. Let $u^{i,N}(t, \theta, \hat{x}^N) = \gamma_t v^{i,N}(\theta, \hat{x}^N)$. Applying Ito's formula to each component of this vector-valued function, we obtain, for $l = 1, \dots, p$,

$$\begin{aligned} u_l^{i,N}(t_2, \theta_{t_2}^{i,N}, \hat{x}_{t_2}^N) - u_l^{i,N}(t_1, \theta_{t_1}^{i,N}, \hat{x}_{t_1}^N) &= \int_{t_1}^{t_2} \partial_s u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \\ &+ \int_{t_1}^{t_2} \mathcal{A}_{\hat{x}} u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \\ &+ \int_{t_1}^{t_2} \mathcal{A}_\theta u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \\ &+ \int_{t_1}^{t_2} \gamma_s \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\partial_\theta \partial_{\hat{x}} u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N)\right]\mathrm{d}s \\ &+ \int_{t_1}^{t_2} \partial_{\hat{x}} u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N) \cdot \mathrm{d}\hat{w}_s^N \\ &+ \int_{t_1}^{t_2} \gamma_s \partial_\theta u_l^{i,N}(s, \theta_s^{i,N}, \hat{x}_s^N) \cdot \nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}w_s^i, \end{aligned}$$

where $\mathcal{A}_{\hat{x}}$ and $\mathcal{A}_\theta$ are the infinitesimal generators of $\hat{x}^N$ and θ, and we recall from (21) that $\hat{w}_t^N = (w_t^1, \dots, w_t^N)^T$. Rearranging this identity, and also recalling that $v^{i,N}(\theta, \hat{x}^N)$ is the solution of the Poisson equation, we obtain

$$\begin{aligned} \Gamma_{k,\eta} &= \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathcal{A}_{\hat{x}} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \\ &= \gamma_{\sigma_{k,\eta}} v^{i,N}(\theta_{\sigma_{k,\eta}}^{i,N}, \hat{x}_{\sigma_{k,\eta}}^N) - \gamma_{\tau_k} v^{i,N}(\theta_{\tau_k}^{i,N}, \hat{x}_{\tau_k}^N) - \int_{\tau_k}^{\sigma_{k,\eta}} \dot{\gamma}_s v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s \end{aligned}$$

$$- \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathcal{A}_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s - \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s^2 \mathrm{Tr}\big[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{s}_x^N)\partial_\theta \partial_{\hat{x}} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\big]\mathrm{d}s$$
$$- \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \partial_{\hat{x}} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \cdot \mathrm{d}\hat{w}_s^N - \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s^2 \partial_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \cdot \nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}w_s^i.$$

We now prove the convergence of each term on the right hand side of this equation. We begin by defining $J_t^{(1)} = \gamma_t \| v^{i,N}(\theta_t^{i,N}, \hat{x}_t^N)\|$. We have, using the polynomial growth of $v^{i,N}(\theta, \hat{x}^N)$ in (B.5), and Proposition 1 (the bounded moments of the IPS), that

$$\mathbb{E}_{\theta_0}[|J_t^{(1)}|^2] \leq K_{\theta_0}\gamma_t^2\left(1 + \mathbb{E}_{\theta_0}[\|x_t^{i,N}\|^{q'}] + \frac{1}{N}\sum_{j=1}^N \mathbb{E}_{\theta_0}[\|x_t^{j,N}\|^{q'}]\right) \leq K_{\theta_0}\gamma_t^2.$$

Applying the Borel–Cantelli argument as in [96, Appendix B], it follows that $J_t^{(1)}$ converges to zero with probability one. We next consider the term

$$J_{0,t}^{(2)} = \int_0^t \partial_s \dot{\gamma}_s v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s + \int_0^t \gamma_s \mathcal{A}_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}s$$
$$+ \int_0^t \gamma_s^2 \mathrm{Tr}\big[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\partial_\theta \partial_x v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\big]\mathrm{d}s.$$

This term obeys the bound

$$\sup_{t>0} \mathbb{E}_{\theta_0}|J_{0,t}^{(2)}| \leq K_{\theta_0}\int_0^\infty (|\dot{\gamma}_s| + \gamma_s^2)(1 + \mathbb{E}_{\theta_0}[\|x_s^{i,N}\|^{q'}] + \frac{1}{N}\sum_{j=1}^N \mathbb{E}_{\theta_0}[\|x_s^{j,N}\|^{q'}])\mathrm{d}s$$
$$\leq K_{\theta_0}\int_0^\infty (|\dot{\gamma}_s| + \gamma_s^2)\mathrm{d}s < \infty.$$

Here, the first inequality follows from the growth properties of the $v^{i,N}(\theta, \hat{x}^N)$, the second inequality from Proposition 1 (the bounded moments of the IPS), and the final inequality from Condition F.1 (the properties of the learning rate). It follows that there exists a finite random variable $J_{0,\infty}^{(2)}$ such that, with probability one,

$$J_{0,t}^{(2)} \to J_{0,\infty}^{(2)}, \quad \text{as } t \to \infty. \tag{B.6}$$

The last term to consider is the stochastic integral

$$J_{0,t}^{(3)} = \int_0^t \gamma_s \partial_{\hat{x}} v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \cdot \mathrm{d}\hat{w}_s^N + \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s^2 \partial_\theta v^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \cdot \nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\mathrm{d}w_s^i.$$

In this case, using the BDG inequality, and the same bounds as above, we have

$$\mathbb{E}_{\theta_0}\left[|J_{0,t}^{(3)}|^2\right] \leq K_{\theta_0}\int_0^\infty (\gamma_s^2 + \gamma_s^4)\left[1 + \mathbb{E}_{\theta_0}[\|x_s^{i,N}\|^{q'}] + \frac{1}{N}\sum_{j=1}^N \mathbb{E}_{\theta_0}[\|x_s^{j,N}\|^{q'}]\right]\mathrm{d}s$$
$$\leq K_{\theta_0}\int_0^\infty \gamma_s^2 \mathrm{d}s < \infty.$$

Thus, by Doob's martingale convergence theorem, there exists a square integrable random variable $J_{0,\infty}^{(3)}$ such that, both almost surely and in $\mathbb{L}^2$,

$$J_{0,t}^{(3)} \to J_{0,\infty}^{(3)}, \quad \text{as } t \to \infty. \tag{B.7}$$

It remains only to observe, combining (B.6) and (B.7), we have

$$\|\Gamma_{k,\eta}\| \leq J^{(1)}_{\sigma_{k,\eta}} + J^{(1)}_{\tau_k} + J^{(2)}_{\tau_k,\sigma_{k,\eta}} + J^{(3)}_{\tau_k,\sigma_{k,\eta}} \overset{k\to\infty}{\to} 0. \quad \square$$

Lemma 11. *Assume that Conditions A.1, B.1–B.2, C.1 and F.1 hold. Let $\rho > 0$ be such that, for a given $\kappa > 0$, it is true that $3\rho + \frac{\rho}{4\kappa} = \frac{1}{2L}$, where L denotes the Lipschitz constant of $\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta)$. For k large enough, and for $\eta > 0$ small enough (potentially random, and depending on k), one has*

$$\int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s > \rho \quad \textit{and, a.s.,} \quad \frac{\rho}{2} \leq \int_{\tau_k}^{\sigma_k} \gamma_s \mathrm{d}s \leq \rho.$$

Proof. We proceed by contradiction. Let us assume that $\int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s \leq \rho$. Choose arbitrary $\varepsilon > 0$ such that $\varepsilon \leq \frac{\rho}{8}$. We begin with the observation that, via the Itô isometry, we have that

$$\begin{aligned}
&\sup_{t\geq 0} \mathbb{E}_{\theta_0}\Big|\Big| \int_0^t \gamma_s \frac{\kappa}{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|} \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s \Big|\Big|^2 \\
&\quad \leq \int_0^t K\gamma_s^2 \left(1 + \mathbb{E}_{\theta_0}[\|x^{i,N}_s\|^q] + \frac{1}{N}\sum_{j=1}^N \mathbb{E}_{\theta_0}[\|x^{j,N}_s\|^q]\right) \mathrm{d}s < \infty,
\end{aligned}$$

where we have used Condition C.1, which implies that $\|\nabla_\theta \hat{B}^{i,N}(\theta, \hat{x})\| \leq K(1 + \|x_i\|^q + \frac{1}{N}\sum_{j=1}^N \|x_j\|^q)$ for some K which is independent of N, Proposition 1 (the bounded moments of the IPS), and Condition F.1 (the properties of the learning rate). Thus, by the Doob's martingale convergence theorem, there exists a finite random variable M such that, both almost surely and in $\mathbb{L}^2$,

$$\int_0^t \gamma_s \frac{\kappa}{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|} \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s \to M.$$

It follows that, for the chosen $\varepsilon > 0$, there exists k such that

$$\int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \frac{\kappa}{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|} \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s < \varepsilon. \tag{B.8}$$

Let us now also assume that, for the given k, η is small enough such that for all $s \in [\tau_k, \sigma_{k,\eta}]$, we have $\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)\| \leq 3\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|$. We can then compute

$$\begin{aligned}
\|\theta^{i,N}_{\sigma_{k,\eta}} - \theta^{i,N}_{\tau_k}\| &= \Big|\Big| \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \nabla_\theta \hat{L}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}s + \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \langle \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s), \mathrm{d}w^i_s \rangle \Big|\Big| \\
&\leq 3\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s \\
&\quad + \Big|\Big| \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s [\nabla_\theta \hat{L}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)] \mathrm{d}s \Big|\Big| \\
&\quad + \frac{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|}{\kappa} \Big|\Big| \int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \frac{\kappa}{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|} \langle \nabla_\theta B(\theta^{i,N}_s, \hat{x}^N_s), \mathrm{d}w^i_s \rangle \Big|\Big| \\
&\leq 3\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|\rho + \varepsilon + \frac{\|\nabla \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|}{\kappa}\varepsilon \leq \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \left[3\rho + \frac{\rho}{4\kappa}\right],
\end{aligned}$$

where in the penultimate line we have used Lemma 10 and (B.8), and in the final line we have used the fact that our choice of ε satisfies $\varepsilon \leq \frac{\rho}{8}$. We thus obtain

$$\|\theta^{i,N}_{\sigma_{k,\eta}} - \theta^{i,N}_{\tau_k}\| \leq \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \left[3\rho + \frac{\rho}{4\kappa}\right] \leq \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \frac{1}{2L}.$$

Thus, using also the definition of the Lipschitz constant L, we have

$$\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_{k,\eta}}) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \leq L\|\theta^{i,N}_{\sigma_{k,\eta}} - \theta^{i,N}_{\tau_k}\| \leq \frac{1}{2}\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|,$$

which then yields

$$\frac{1}{2}\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\| \leq \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_{k,\eta}})\| \leq 2\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|.$$

But this implies that $\sigma_{k,\eta} \in [\tau_k, \sigma_k]$, which is a contradiction. Thus we do indeed have $\int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s > \rho$. We now turn our attention to the second part of the Lemma. By definition, we have that $\int_{\tau_k}^{\sigma_k} \gamma_s \mathrm{d}s \leq \rho$. Thus, it remains only to show that $\frac{\rho}{2} \leq \int_{\tau_k}^{\sigma_k} \gamma_s \mathrm{d}s$. From the first part of the Lemma, we have that $\int_{\tau_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s > \rho$. Moreover, for k sufficiently large and η sufficiently small, we must have $\int_{\sigma_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s \leq \frac{\rho}{2}$. We thus obtain

$$\int_{\tau_k}^{\sigma_k} \gamma_s \mathrm{d}s \geq \rho - \int_{\sigma_k}^{\sigma_{k,\eta}} \gamma_s \mathrm{d}s \geq \rho - \frac{\rho}{2} = \frac{\rho}{2}. \quad \square$$

Lemma 12. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Suppose that there are an infinite number of intervals $[\tau_k, \sigma_k)$. Then there exists a fixed constant $\beta = \beta(\kappa) > 0$ such that, for k large enough, almost surely,*

$$\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_k}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k}) \geq \beta.$$

Proof. By Itô's formula, we have that

$$\begin{aligned}
&\tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\sigma_k}) - \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k}) \\
&= \underbrace{\int_{\tau_k}^{\sigma_k} \gamma_s \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)\|^2 \mathrm{d}s}_{A^{i,N}_{1,k}} + \underbrace{\int_{\tau_k}^{\sigma_k} \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s), \nabla_\theta \hat{L}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)\rangle \mathrm{d}s}_{A^{i,N}_{2,k}} \\
&+ \underbrace{\int_{\tau_k}^{\sigma_k} \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s), \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s \rangle}_{A^{i,N}_{3,k}} \\
&+ \underbrace{\int_{\tau_k}^{\sigma_k} \frac{1}{2}\gamma_s^2 \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s)^T \nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s) \mathrm{d}s\right]}_{A^{i,N}_{4,k}}.
\end{aligned}$$

We will deal with each of these terms individually. First consider $A^{i,N}_{1,k}$. For this term, we have that

$$A^{i,N}_{1,k} = \int_{\tau_k}^{\sigma_k} \gamma_s \|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)\|^2 \mathrm{d}s \geq \frac{\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|^2}{4} \int_{\tau_k}^{\sigma_k} \gamma_s \mathrm{d}s \geq \frac{\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|^2}{8} \rho$$

where, in the first inequality, we have used the definition of the $\{\tau_k\}_{k\geq 0}$, namely, that $\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s)\| \geq \frac{1}{2}\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_{\tau_k})\|$ for all $s \in [\tau_k, \sigma_k]$, and in the second inequality we have used Lemma 11.

We now turn our attention to $A^{i,N}_{2,k}$. We will handle this term using a very similar to approach to that used in the proof of Lemma 10. Let us consider the function $T^{i,N}(\theta, \hat{x}^N) = \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta), \nabla_\theta \hat{L}^{i,N}(\theta, \hat{x}^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta) \rangle$. Due to Lemma 7 (the boundedness of the asymptotic log-likelihood and its derivatives) and Lemma 14 (the local Lipschitz and polynomial growth of $\hat{L}^{i,N}(\theta, \hat{x}^N)$), there exist $q, K > 0$, independent of N, such that, for $k = 0, 1, 2$,

$$\|\partial^k_\theta T^{i,N}(\theta, \hat{x}^N) - \partial^k_\theta T^{i,N}(\theta, \hat{y}^N)\| \leq K \left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N} \sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2 \right)^{\frac{1}{2}} \right] \cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^q \right].$$

Moreover, it is straightforward to show that this function satisfies $\int_{(\mathbb{R}^d)^N} T^{i,N}(\theta, \hat{x}^N) \hat{\mu}_\infty(\mathrm{d}\hat{x}^N) = 0$. Thus, by Lemmas 16 and 17 with $r = 0$, the Poisson equation

$$\mathcal{A}_{\hat{x}} v^{i,N}(\theta, \hat{x}^N) = T^{i,N}(\theta, \hat{x}^N) \quad , \quad \int_{(\mathbb{R}^d)^N} v^{i,N}(\theta, \hat{x}^N) \mu_\infty(\mathrm{d}\hat{x}^N) = 0$$

has a unique twice differentiable solution which satisfies a polynomial growth property of the form

$$\sum_{k=0}^2 \| \frac{\partial^k v^{i,N}}{\partial \theta^k}(\theta, \hat{x}^N)\| + \| \frac{\partial^2 v^{i,N}}{\partial \theta \partial \hat{x}^N}(\theta, \hat{x}^N)\| \leq K_{\theta_0}(1 + \|x^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'}),$$

for some constants $q', K_{\theta_0} > 0$ which are also independent of N. Using the same steps as in the proof of Lemma 10, we can then prove that, a.s.,

$$\left\| \int_{\tau_k}^{\sigma_k} \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s), \nabla_\theta \hat{L}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s) \rangle \mathrm{d}s \right\| \overset{k \to \infty}{\to} 0.$$

We next consider $A^{i,N}_{3,k}$. Using Itô's isometry, Lemma 7 (the bound on the asymptotic log-likelihood of the IPS), the polynomial growth of the function $\nabla_\theta \hat{B}^{i,N}(\theta, \hat{x})$, Proposition 1 (the moment bounds for solutions of the IPS) and Condition F.1 (the square summability of the learning rate), we have that

$$\begin{aligned} &\sup_{t \geq 0} \mathbb{E}_{\theta_0} \left[\left| \int_0^t \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s), \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s \rangle \right|^2 \right] \\ &\quad \leq K_{\theta_0} \mathbb{E}_{\theta_0} \int_0^\infty \gamma_s^2 \|\nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s)\|^2 \mathrm{d}s \\ &\quad \leq K_{\theta_0} \int_0^\infty \gamma_s^2 \left[1 + \mathbb{E}_{\theta_0} \left[\|x^{i,N}_s\|^q \right] + \frac{1}{N} \sum_{j=1}^N \mathbb{E}_{\theta_0} \left[\|x^{j,N}_s\|^q \right] \right] \mathrm{d}s < \infty. \end{aligned}$$

Thus, by Doob's martingale convergence theorem, there exists a finite random variable $A^{i,N}_{3,\infty}$ such that, both almost surely and in $\mathbb{L}^2$,

$$\int_0^t \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta^{i,N}_s), \nabla_\theta \hat{B}^{i,N}(\theta^{i,N}_s, \hat{x}^N_s) \mathrm{d}w^i_s \rangle \to A^{i,N}_{3,\infty}$$

as $t \to \infty$. It follows that $A_{3,k}^{i,N} \to 0$ a.s. as $k \to \infty$. Finally, we turn our attention to $A_{4,k}^{i,N}$. For this term, we observe that

$$\sup_{t\geq 0} \mathbb{E}_{\theta_0} \left\| \int_0^t \frac{1}{2}\gamma_s^2 \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)^T \nabla_\theta^2 \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\right] \mathrm{d}s \right\|$$
$$\leq K_{\theta_0} \int_0^\infty \gamma_s^2 (1 + \mathbb{E}_{\theta_0}\left[\|x_s^{i,N}\|^q\right] + \frac{1}{N}\sum_{j=1}^N \mathbb{E}_{\theta_0}\left[\|\hat{x}_s^{j,N}\|^q\right])\mathrm{d}s < \infty,$$

where, as above, we have used Lemma 7, the polynomial growth of $\nabla_\theta \hat{B}^{i,N}(\theta, \hat{x})$, Proposition 1, and Condition F.1. It follows that the random variable

$$\int_0^\infty \frac{1}{2}\gamma_s^2 \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)^T \nabla_\theta^2 \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\right] \mathrm{d}s$$

is finite a.s., which in turn implies that there exists a finite random variable $A_{4,\infty}^{i,N}$ such that

$$\int_0^t \frac{1}{2}\gamma_s^2 \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)^T \nabla_\theta^2 \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\right] \mathrm{d}s \to A_4^\infty$$

almost surely. It follows, in particular, that $A_{4,k}^{i,N} \to 0$ a.s. as $k \to \infty$. Summarising, we thus have that, for all $\varepsilon > 0$, there exists k such that

$$\tilde{\mathcal{L}}^{i,N}(\theta_{\sigma_k}^{i,N}) - \tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N}) \geq A_{1,k}^{i,N} - \|A_{2,k}^{i,N}\| - \|A_{3,k}^{i,N}\| - \|A_{4,k}^{i,N}\| = \frac{\|\nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N})\|^2}{8}\rho - 3\varepsilon.$$

The claim follows by setting $\varepsilon = \frac{\rho(\kappa)\kappa^2}{32}$ and $\beta = \frac{\rho(\kappa)\kappa^2}{32}$. □

Lemma 13. *Assume that Conditions A.1, B.1–B.2, C.1, and F.1 hold. Suppose that there are an infinite number of intervals $[\tau_k, \sigma_k)$. Then there exists a fixed constant $0 < \beta_1 < \beta$ such that, for k large enough,*

$$\tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N}) - \tilde{\mathcal{L}}^{i,N}(\theta_{\sigma_{k-1}}^{i,N}) \geq -\beta_1.$$

Proof. Using Itô's formula, we have that

$$\tilde{\mathcal{L}}^{i,N}(\theta_{\tau_k}^{i,N}) - \tilde{\mathcal{L}}^{i,N}(\theta_{\sigma_{k-1}}^{i,N}) \geq \underbrace{\int_{\sigma_{k-1}}^{\tau_k} \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N}), \nabla_\theta \hat{L}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) - \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\rangle \mathrm{d}s}_{B_{1,k}^{i,N}}$$
$$+ \underbrace{\int_{\sigma_{k-1}}^{\tau_k} \gamma_s \langle \nabla_\theta \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N}), \nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N) \mathrm{d}w_s^i\rangle}_{B_{2,k}^{i,N}}$$
$$+ \underbrace{\int_{\sigma_{k-1}}^{\tau_k} \frac{1}{2}\gamma_s^2 \mathrm{Tr}\left[\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)\nabla_\theta \hat{B}^{i,N}(\theta_s^{i,N}, \hat{x}_s^N)^T \nabla_\theta^2 \tilde{\mathcal{L}}^{i,N}(\theta_s^{i,N})\mathrm{d}s\right]}_{B_{3,k}^{i,N}}.$$

Arguing as in the proof of Lemma 12, the magnitude of each of the terms converges to zero a.s. as $k \to \infty$. This is sufficient for the conclusion. □

B.2.2. Technical lemmas: An additional bound

Lemma 14. *Assume that Condition C.1 holds. Then, for $k = 0, 1, 2, 3$, there exist constants $q, K < \infty$ such that the functions $\nabla_\theta^k L(\theta, x, \mu, \mu')$ satisfy*

$$\begin{aligned}\|\nabla_\theta^k L(\theta, x, \mu, \mu') - \nabla_\theta^k L(\theta, x', \nu, \nu')\| \leq K \big[\|x - x'\| + \mathbb{W}_2(\mu, \nu) + \mathbb{W}_2(\mu', \nu')\big] \\ \cdot \big[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^q) \\ + \mu'(\|\cdot\|^q) + \nu(\|\cdot\|^q) + \nu'(\|\cdot\|^q)\big].\end{aligned}$$

Thus, in particular, for $N \in \mathbb{N}$, $i = 1, \ldots, N$, there exist constant $q, K < \infty$, independent of N, such that the functions $\nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{x}^N)$ satisfy

$$\begin{aligned}\left|\nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{x}^N) - \nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{y}^N)\right| \leq K\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}}\right] \\ \cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^N \|y^{j,N}\|^{q'}\right],\end{aligned}$$

where $\hat{x}^N = (x^{1,N}, \ldots, x^{N,N}) \in (\mathbb{R}^d)^N$ and $\hat{y}^N = (x^{1,N}, \ldots, x^{N,N}) \in (\mathbb{R}^d)^N$.

Proof. We begin by recalling that, from Condition C.1 there exist constants $q, K < \infty$ such that, for $k = 0, 1, 2, 3$,

$$\|\nabla_\theta^k \phi(\theta, x, y) - \nabla_\theta^k \phi(\theta, x', y')\| \leq K[\|x - x'\| + \|y - y'\|][1 + \|x\|^q + \|x'\|^q + \|y\|^q + \|y'\|^q]. \tag{B.9}$$

It follows, substituting (B.9) with $y' = y$, and applying the Hölder inequality, that

$$\begin{aligned}&\left\|\int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x, y)\nu(\mathrm{d}y) - \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x', y)\nu(\mathrm{d}y)\right\| \\ &\quad\leq \int_{\mathbb{R}^d} \left\|\nabla_\theta^k \phi(\theta, x, y)\nu(\mathrm{d}y) - \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x', y)\right\| \nu(\mathrm{d}y) \\ &\quad\leq K\|x - x'\|\left[1 + \|x\|^q + \|x'\|^q + \nu\left(\|\cdot\|^{2q}\right)^{\frac{1}{2}}\right].\end{aligned} \tag{B.10}$$

Meanwhile, letting $\pi \in \Pi(\mu, \nu)$, substituting (B.9) with $x' = x$, and once again using the Hölder inequality, we have that

$$\begin{aligned}&\left\|\int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x, y)\mu(\mathrm{d}y) - \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x, y')\nu(\mathrm{d}y')\right\| \\ &\leq K\left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|y - y'\|\left[1 + \|x\|^q + \|x'\|^q + \|y\|^q + \|y'\|^q\right]\pi(\mathrm{d}y, \mathrm{d}y')\right] \\ &\leq K\left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|y - y'\|^2 \pi(\mathrm{d}y, \mathrm{d}z)\right]^{\frac{1}{2}} \left[1 + \|x\|^q + \|x'\|^q + \left[\int_{\mathbb{R}^d} \|y\|^{2q}\mu(\mathrm{d}y)\right]^{\frac{1}{2}}\right. \\ &\quad\left. + \left[\int_{\mathbb{R}^d} \|y'\|^{2q}\nu(\mathrm{d}y')\right]^{\frac{1}{2}}\right] \\ &\leq K\mathbb{W}_2(\mu, \nu)\left[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^{2q})^{\frac{1}{2}} + \nu(\|\cdot\|^{2q})^{\frac{1}{2}}\right],\end{aligned} \tag{B.11}$$

where the final line follows from the fact that π was chosen arbitrarily. Thus, using the triangle inequality, substituting both (B.10) and (B.11), and once more using the Hölder inequality, we have

$$\begin{aligned}&\left\| \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x, y)\mu(\mathrm{d}y) - \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x', y')\nu(\mathrm{d}y') \right\| \\ &\leq K \left[\|x - x'\| + \mathbb{W}_2(\mu, \nu) \right] \left[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^{2q})^{\frac{1}{2}} + \nu(\|\cdot\|^{2q})^{\frac{1}{2}} \right],\end{aligned}$$

where, as elsewhere, the constant K is allowed to increase from line to line. Thus, recalling the definition of the function B, and once more using Condition C.1, we obtain

$$\begin{aligned}&\|\nabla_\theta^k B(\theta, x, \mu) - \nabla_\theta^k B(\theta, x', \nu)\| \\ &\leq \left\| \nabla_\theta^k b(\theta, x) - \nabla_\theta^k b(\theta, x') \right\| + \left\| \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x, y)\mu(\mathrm{d}y) - \int_{\mathbb{R}^d} \nabla_\theta^k \phi(\theta, x', y')\nu(\mathrm{d}y') \right\| \\ &\leq K \left[\|x - x'\| + \mathbb{W}_2(\mu, \nu) \right] \left[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^{2q})^{\frac{1}{2}} + \nu(\|\cdot\|^{2q})^{\frac{1}{2}} \right].\end{aligned}$$

Using the definition of the function G, the triangle inequality, and substituting the previous bound, we then have

$$\begin{aligned}&\|\nabla_\theta^k G(\theta, x, \mu, \mu') - \nabla_\theta^k G(\theta, x', \nu, \nu')\| \\ &= \|\nabla_\theta^k \left[B(\theta, x, \mu) - B(\theta_0, x, \mu') \right] - \nabla_\theta^k \left[B(\theta, x', \nu) - B(\theta_0, x', \nu') \right] \| \\ &\leq \|\nabla_\theta^k B(\theta, x, \mu) - \nabla_\theta^k B(\theta, x', \nu)\| + \|\nabla_\theta^k B(\theta_0, x, \mu') - \nabla_\theta^k B(\theta, x', \nu')\| \\ &\leq K \left[\|x - x'\| + \mathbb{W}_2(\mu, \nu) + \mathbb{W}_2(\mu', \nu') \right] \\ &\quad \cdot \left[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^{2q})^{\frac{1}{2}} + \mu'(\|\cdot\|^{2q})^{\frac{1}{2}} + \nu(\|\cdot\|^{2q})^{\frac{1}{2}} + \nu'(\|\cdot\|^{2q})^{\frac{1}{2}} \right].\end{aligned}$$

Finally, recalling the definition of the function L, combining the previous two bounds, and using the triangle inequality and the Hölder inequality, we have, now allowing also the constant q to increase from the previous line,

$$\begin{aligned}&\|L(\theta, x, \mu, \mu') - L(\theta, x', \nu, \nu')\| = \frac{1}{2} \left\| G(\theta, x, \mu, \mu') - G(\theta, x', \nu, \nu') \right\|^2 \\ &\leq \frac{1}{2} \|G(\theta, x, \mu, \mu') - G(\theta, x', \nu, \nu')\| \left\| G(\theta, x, \mu, \mu') + G(\theta, x', \nu, \nu') \right\| \\ &\leq K \left[\|x - x'\| + \mathbb{W}_2(\mu, \nu) + \mathbb{W}_2(\mu', \nu') \right] \\ &\quad \cdot \left[1 + \|x\|^q + \|x'\|^q + \mu(\|\cdot\|^q)^{\frac{1}{2}} + \mu'(\|\cdot\|^q)^{\frac{1}{2}} + \nu(\|\cdot\|^q)^{\frac{1}{2}} + \nu'(\|\cdot\|^q)^{\frac{1}{2}} \right].\end{aligned}$$

The bounds on $\|\nabla_\theta^k L(\theta, x, \mu, \mu') - \nabla_\theta^k L(\theta, x', \nu, \nu')\|$, $k = 1, 2, 3$, follow via analogous arguments. We now turn our attention to the second statement in the lemma. In fact, this is a straightforward corollary of the previous result. Indeed, recall from the definition that $\hat{L}^{i,N}(\theta, \hat{x}^N) := L(\theta, x^{i,N}, \hat{\mu}^N, \hat{\mu}^N)$. Using this, and the previous result, it follows immediately that

$$\begin{aligned}&\left\| \nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{x}^N) - \nabla_\theta^k \hat{L}^{i,N}(\theta, \hat{y}^N) \right\| \\ &\leq K \left[\|x^{i,N} - y^{i,N}\| + \mathbb{W}_2\left(\frac{1}{N} \sum_{j=1}^N \delta_{x^{i,N}}, \frac{1}{N} \sum_{j=1}^N \delta_{y^{j,N}} \right) \right]\end{aligned}$$

$$\cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|y^{j,N}\|^q \right]$$

$$\leq K\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}} \right]$$

$$\cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|y^{j,N}\|^q \right]. \quad \square$$

B.2.3. Technical lemmas: On a related Poisson equation

Lemma 15. *Assume that Conditions A.1 and B.1–B.2 hold. Suppose that, for each $\theta \in \mathbb{R}^p$, $f(\cdot, x)$ is a locally Lipschitz function, which satisfies a polynomial growth condition in the sense that*

$$\|f(\theta, x) - f(\theta, y)\| \leq K\|x - y\|\left[1 + \|\theta\|^r\right][1 + \|x\|^q + \|y\|^q], \tag{B.12}$$

for some constants $r, q, K > 0$, where $x, y \in \mathbb{R}^d$. Then, for all $\theta \in \mathbb{R}^p$, $x, y \in \mathbb{R}^d$, $t \geq 0$, there exist positive constants $q', K_{\theta_0} > 0$, such that

$$\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \int_{\mathbb{R}^d} f(\theta, z)\mu_\infty(\mathrm{d}z)\right| \leq K_{\theta_0}[1 + \|\theta\|^r][1 + \|x\|^{q'}]e^{-\lambda_{\theta_0}t}.$$

$$\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \mathbb{E}_{\theta_0,y}[f(\theta, x_t)]\right| \leq K_{\theta_0}\|x - y\|[1 + \|\theta\|^r][1 + \|x\|^{q'} + \|y\|^{q'}]e^{-\lambda_{\theta_0}t}. \tag{B.13}$$

Alternatively, suppose that, for all $\theta \in \mathbb{R}^p$, and for some $i \in \{1, \ldots, N\}$, the function $f(\cdot, \hat{x}^N)$ is locally Lipschitz, and satisfies a polynomial growth condition in the sense that

$$\left|f(\theta, \hat{x}^N) - f(\theta, \hat{y}^N)\right| \leq K\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}}\right]\left[1 + \|\theta\|^r\right]$$

$$\cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N}\sum_{j=1}^N \|y^{j,N}\|^q\right] \tag{B.14}$$

for some constants r, q, K which are independent of N, where $\hat{x}^N = (x^{1,N}, \ldots, x^{N,N})^\top \in (\mathbb{R}^d)^N$. Then, for all $N \in \mathbb{N}$, $i = 1, \ldots, N$, and for all $\theta \in \mathbb{R}^p$, $\hat{x}^N, \hat{y}^N \in (\mathbb{R}^d)^N$, $t \geq 0$, there exists positive constants $q', K_{\theta_0} > 0$, independent of N, such that

$$\left|\mathbb{E}_{\theta_0,\hat{x}^N}[f(\theta, \hat{x}_t^N)] - \int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N)\right|$$

$$\leq K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^{q'}\right]e^{-\lambda_{\theta_0}t} \tag{B.15}$$

$$\left|\mathbb{E}_{\theta_0,\hat{x}^N}[f(\theta, \hat{x}_t^N)] - \mathbb{E}_{\theta_0,\hat{y}^N}[f(\theta, \hat{x}_t^N)]\right|$$

$$\leq K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}}\right]$$

$$\left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}(\|x^{j,N}\|^{q'} + \|y^{j,N}\|^{q'})\right]e^{-\lambda_{\theta_0}t}. \tag{B.16}$$

Proof.

We begin with the first statement of the lemma. Let $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, and $\pi \in \Pi(\mu, \nu)$. Using the Hölder inequality and the local Lipschitz and polynomial growth assumption in (B.12), we have that

$$\begin{aligned}
&\int_{\mathbb{R}^d\times\mathbb{R}^d} |f(\theta, y) - f(\theta, z)|\,\pi(\mathrm{d}y, \mathrm{d}z) \\
&\leq K\left[1 + \|\theta\|^r\right]\left[\int_{\mathbb{R}^d\times\mathbb{R}^d} \|y - z\|^2\pi(\mathrm{d}y, \mathrm{d}z)\right]^{\frac{1}{2}}\left[1 + \left[\int_{\mathbb{R}^d\times\mathbb{R}^d} \|y\|^{2q}\pi(\mathrm{d}y, \mathrm{d}z)\right]^{\frac{1}{2}}\right. \\
&\qquad\qquad \left. + \left[\int_{\mathbb{R}^d\times\mathbb{R}^d} \|y\|^{2q}\pi(\mathrm{d}y, \mathrm{d}z)\right]^{\frac{1}{2}}\right] \\
&= K\left[1 + \|\theta\|^r\right]\left[\int_{\mathbb{R}^d\times\mathbb{R}^d} \|y - z\|^2\pi(\mathrm{d}y, \mathrm{d}z)\right]^{\frac{1}{2}}\left[1 + \left[\int_{\mathbb{R}^d} \|y\|^{2q}\mu(\mathrm{d}y)\right]^{\frac{1}{2}}\right. \\
&\qquad\qquad \left. + \left[\int_{\mathbb{R}^d} \|z\|^{2q}\nu(\mathrm{d}z)\right]^{\frac{1}{2}}\right].
\end{aligned} \tag{B.17}$$

Now let $(x_t)_{t\geq 0}$ be a solution of the McKean–Vlasov SDE (1) starting from $x \in \mathbb{R}^d$. Let μ_t^x denote the law of x_t, and let μ_∞ denote the unique invariant measure. Moreover, let $\pi_t^{x,\infty}$ denote an arbitrary coupling of μ_t^x and μ_∞.

$$\begin{aligned}
\left|\mathbb{E}_{\theta_0,x}\left[f(\theta, x_t)\right] - \int_{\mathbb{R}^d} f(\theta, z)\mu_\infty(\mathrm{d}z)\right| &= \left|\int_{\mathbb{R}^d} f(\theta, y)\mu_t^x(\mathrm{d}y) - \int_{\mathbb{R}^d} f(\theta, z)\mu_\infty(\mathrm{d}z)\right| \\
&\leq \int_{\mathbb{R}^d\times\mathbb{R}^d} |f(\theta, y) - f(\theta, z)|\pi_t^{x,\infty}(\mathrm{d}y, \mathrm{d}z).
\end{aligned}$$

Substituting (B.17) with $\mu = \mu_t^x$ and $\nu = \mu_\infty$, and using the fact that the chosen coupling was arbitrary, we have that

$$\begin{aligned}
&\left|\mathbb{E}_{\theta_0,x}\left[f(\theta, x_t)\right] - \int_{\mathbb{R}^d} f(\theta, z)\mu_\infty(\mathrm{d}z)\right| \\
&\quad\leq K\mathbb{W}_2(\mu_t^x, \mu_\infty)\left[1 + \|\theta\|^r\right]\left[1 + \left[\int_{\mathbb{R}^d} \|y\|^{2q}\mu_t^x(\mathrm{d}y)\right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z\|^{2q}\mu_\infty(\mathrm{d}z)\right]^{\frac{1}{2}}\right].
\end{aligned} \tag{B.18}$$

We will bound each of these terms. Using Proposition 1 (the moment bounds for the McKean–Vlasov SDE), Proposition 2 (the bounded moments of the invariant measure of the McKean–Vlasov SDE), we have

$$\left[1 + \left[\int_{\mathbb{R}^d} \|y\|^{2q}\mu_t^x(\mathrm{d}y)\right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z\|^{2q}\mu_\infty(\mathrm{d}z)\right]^{\frac{1}{2}}\right] \leq K_{\theta_0}(1 + \|x\|^q). \tag{B.19}$$

Meanwhile, using also Proposition 3 (the exponential contractivity of the McKean–Vlasov SDE), we have that

$$\mathbb{W}_2(\mu_t^x, \mu_\infty) \leq \mathbb{W}_2(\mu_0^x, \mu_\infty)e^{-\lambda_{\theta_0}t} \leq \left[\int_{\mathbb{R}^d} \|x - y\|^2\mu_\infty(\mathrm{d}y)\right]^{\frac{1}{2}} e^{-\lambda_{\theta_0}t}$$

$$\leq 2\left[\|x\|^2 + \int_{\mathbb{R}^d} \|y\|^2 \mu_\infty(\mathrm{d}y)\right]^{\frac{1}{2}} e^{-\lambda_{\theta_0} t} \leq K_{\theta_0}(1 + \|x\|^q) e^{-\lambda_{\theta_0} t} \tag{B.20}$$

for some $q \geq 1$. Substituting (B.19) and (B.20) into (B.18), using the elementary inequality $(a+b)^2 \leq 2(a^2+b^2)$, and setting $q' = 2q$, we thus arrive at

$$\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \int_{\mathbb{R}^d} f(\theta, z)\mu_\infty(\mathrm{d}z)\right| \leq K_{\theta_0}\left[1 + \|\theta\|^r\right][1 + \|x\|^{q'}]e^{-\lambda_{\theta_0} t}.$$

We now turn our attention to the bound in (B.13). In fact, the proof of this bound is very similar. Let $(x_t)_{t\geq 0}$ and $(y_t)_{t\geq 0}$ be solution of the McKean–Vlasov SDE (1) starting from $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^d$, respectively. In addition, let μ_t^x and μ_t^y denote the laws of x_t and y_t, and let $\pi_t^{x,y}$ denote an arbitrary coupling of μ_t^x and μ_t^y. Then, starting from the LHS of the bound in (B.13), before using (B.17) with $\mu = \mu_t^x$ and $\nu = \mu_t^y$, we have

$$\begin{aligned}
&\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \mathbb{E}_{\theta_0,y}[f(\theta, y_t)]\right| = \left|\int_{\mathbb{R}^d} f(\theta, z_1)\mu_t^x(\mathrm{d}z_1) - \int_{\mathbb{R}^d} f(\theta, z_2)\mu_t^y(\mathrm{d}z_2)\right| \\
&\leq \int_{\mathbb{R}^d\times\mathbb{R}^d} |f(\theta, z_1) - f(\theta, z_2)|\pi_t^{x,y}(\mathrm{d}z_1, \mathrm{d}z_2) \\
&\leq K\left[1 + \|\theta\|^r\right]\left[\int_{\mathbb{R}^d\times\mathbb{R}^d} \|z_1 - z_2\|^2 \pi_t^{x,y}(\mathrm{d}z_1, \mathrm{d}z_2)\right]^{\frac{1}{2}} \\
&\quad \cdot \left[1 + \left[\int_{\mathbb{R}^d} \|z_1\|^{2q}\mu_t^x(\mathrm{d}z_1)\right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z_2\|^{2q}\mu_t^y(\mathrm{d}z_2)\right]^{\frac{1}{2}}\right].
\end{aligned}$$

Once again, since the coupling was chosen arbitrarily, we in fact have that

$$\begin{aligned}
&\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \mathbb{E}_{\theta_0,y}[f(\theta, y_t)]\right| \\
&\quad \leq K\mathbb{W}_2(\mu_t^x, \mu_t^y)\left[1 + \|\theta\|^r\right]\left[1 + \left[\int_{\mathbb{R}^d} \|z_1\|^{2q}\mu_t^x(\mathrm{d}z_1)\right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z_2\|^{2q}\mu_t^y(\mathrm{d}z_2)\right]^{\frac{1}{2}}\right].
\end{aligned} \tag{B.21}$$

We will bound each term, similar to before. On this occasion, once more appealing to Proposition 1 and the proof of Proposition 3 (see (A.5), and also Remark 3.6 in [11]) the required bounds are given, for some $K_{\theta_0} > 0$, by

$$\left[\int_{\mathbb{R}^d} \|z_1\|^{2q}\mu_t^x(\mathrm{d}z_1)\right]^{\frac{1}{2}} \leq \left[K_{\theta_0}(1 + \|x\|^{2q})\right]^{\frac{1}{2}} \leq K_{\theta_0}(1 + \|x\|^q), \tag{B.22}$$

$$\left[\int_{\mathbb{R}^d} \|z_2\|^{2q}\mu_t^y(\mathrm{d}z_2)\right]^{\frac{1}{2}} \leq \left[K_{\theta_0}(1 + \|y\|^{2q})\right]^{\frac{1}{2}} \leq K_{\theta_0}(1 + \|y\|^q), \tag{B.23}$$

$$\mathbb{W}_2(\mu_t^x, \mu_t^y) \leq \mathbb{W}_2(\mu_0^x, \mu_0^y)e^{-\lambda_{\theta_0} t} \leq \|x - y\|e^{-\lambda_{\theta_0} t}. \tag{B.24}$$

Substituting (B.22), (B.23), and (B.24) into (B.21), and this time $q' = q$, we have the bound in (B.13), namely,

$$\left|\mathbb{E}_{\theta_0,x}[f(\theta, x_t)] - \mathbb{E}_{\theta_0,y}[f(\theta, x_t)]\right| \leq K_{\theta_0}\|x - y\|[1 + \|\theta\|^r][1 + \|x\|^{q'} + \|y\|^{q'}]e^{-\lambda_{\theta_0} t}.$$

The proof of the second part of the lemma follows closely the previous proof, now using the statements in Propositions 1–3 that are relevant to the IPS. For completeness, and to emphasise that any and all constants are independent of N, we provide this proof in full. Suppose

now that $\mu^N, \nu^N \in \mathcal{P}((\mathbb{R}^d)^N)$, with marginals $\mu^{i,N}, \nu^{i,N} \in \mathcal{P}(\mathbb{R}^d)$ for $i = 1, \dots, N$. Let $\pi^N \in \Pi(\mu^N, \nu^N)$ denote an arbitrary coupling of μ^N and ν^N, and $\pi^{i,N}$ an arbitrary coupling of $\mu^{i,N}$ and $\pi^{i,N}$. Using the Hölder inequality and the local Lipschitz and polynomial growth assumption in (B.14), it is straightforward to obtain

$$
\begin{aligned}
&\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \left| f(\theta, \hat{y}^N) - f(\theta, \hat{z}^N) \right| \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \\
&\leq K \left[1 + \|\theta\|^r \right] \Bigg[\left[\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \|y^{i,N} - z^{i,N}\|^2 \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \\
&\qquad + \left[\frac{1}{N} \sum_{j=1}^N \int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \|y^{j,N} - z^{j,N}\|^2 \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \Bigg] \\
&\quad \cdot \Bigg[1 + \left[\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \|y^{i,N}\|^{2q} \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \\
&\qquad + \left[\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \|z^{i,N}\|^{2q} \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \\
&\qquad + \left[\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \left(\frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^q \right)^2 \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \\
&\qquad + \left[\int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} \left(\frac{1}{N} \sum_{j=1}^N \|z^{j,N}\|^q \right)^2 \pi^N(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N) \right]^{\frac{1}{2}} \Bigg] \\
&\leq K \left[1 + \|\theta\|^r \right] \Bigg[\left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|y^{i,N} - z^{i,N}\|^2 \pi^{i,N}(\mathrm{d}y^{i,N}, \mathrm{d}z^{i,N}) \right]^{\frac{1}{2}} \\
&\qquad + \left[\frac{1}{N} \sum_{j=1}^N \int_{\mathbb{R}^d \times \mathbb{R}^d} \|y^{j,N} - z^{j,N}\|^2 \pi^{j,N}(\mathrm{d}y^{j,N}, \mathrm{d}z^{j,N}) \right]^{\frac{1}{2}} \Bigg] \\
&\quad \cdot \Bigg[1 + \left[\int_{\mathbb{R}^d} \|y^{i,N}\|^{2q} \mu^{i,N}(\mathrm{d}y^{i,N}) \right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z^{i,N}\|^{2q} \nu^{i,N}(\mathrm{d}z^{i,N}) \right]^{\frac{1}{2}} \\
&\qquad + \frac{1}{N} \sum_{j=1}^N \left[\int_{\mathbb{R}^d} \|y^{j,N}\|^{2q} \mu^{j,N}(\mathrm{d}y^{j,N}) \right]^{\frac{1}{2}} + \frac{1}{N} \sum_{j=1}^N \left[\int_{\mathbb{R}^d} \|z^{j,N}\|^{2q} \nu^{j,N}(\mathrm{d}z^{j,N}) \right]^{\frac{1}{2}} \Bigg].
\end{aligned}
\tag{B.25}
$$

Let $\hat{x}_t^N = (x_t^{1,N}, \dots, x_t^{N,N})^\top$ be a solution of the IPS (2) starting from $\hat{x}^N = (x^{1,N}, \dots, x^{N,N})^\top \in (\mathbb{R}^d)^N$. Let $(\hat{\mu}_t^{\hat{x}^N,N})_{t \geq 0}$ denote the law of $(\hat{x}_t^N)_{t \geq 0}$, with marginals $(\hat{\mu}_t^{x^{i,N},i,N})_{t \geq 0}^{i=1:N}$, and let $\hat{\mu}_\infty^N$ denote the unique invariant measure, with marginals $(\hat{\mu}_\infty^{i,N})^{i=1:N}$. Moreover, let $\pi_t^{\hat{x}^N,\infty}$ denote an arbitrary coupling of $\hat{\mu}_t^{\hat{x}^N,N}$ and $\hat{\mu}_\infty^N$. Then

$$
\begin{aligned}
&\left| \mathbb{E}_{\theta_0, \hat{x}^N}[f(\theta, \hat{x}_t^N)] - \int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N) \right| \\
&\quad = \left| \int_{(\mathbb{R}^d)^N} f(\theta, \hat{y}^N) \hat{\mu}_t^{x,N}(\mathrm{d}\hat{y}^N) - \int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N) \right|
\end{aligned}
$$

$$\leq \int_{(\mathbb{R}^d)^N \times (\mathbb{R}^d)^N} |f(\theta, \hat{y}^N) - f(\theta, \hat{z}^N)| \pi_t^{\hat{x}^N, \infty}(\mathrm{d}\hat{y}^N, \mathrm{d}\hat{z}^N).$$

Using (B.25) with $\mu^N = \hat{\mu}_t^{\hat{x}^N, N}$ and $\nu^N = \hat{\mu}_\infty^N$, and using the fact that the chosen coupling was arbitrary, we have

$$\begin{aligned}
&\left| \mathbb{E}_{\theta_0, \hat{x}^N}[f(\theta, \hat{x}_t^N)] - \int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N) \right| \\
&\leq K \left[1 + \|\theta\|^r \right] \left[W_2(\hat{\mu}_t^{x^{i,N}, i, N}, \hat{\mu}_\infty^{i,N}) + \frac{1}{N} \sum_{j=1}^N W_2(\hat{\mu}_t^{x^{j,N}, j, N}, \hat{\mu}_\infty^{j,N}) \right] \\
&\quad \cdot \left[1 + \left[\int_{\mathbb{R}^d} \|y^{i,N}\|^{2q} \hat{\mu}_t^{x^{i,N}, i, N}(\mathrm{d}y^{i,N}) \right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z^{i,N}\|^{2q} \hat{\mu}_\infty^{i,N}(\mathrm{d}z^{i,N}) \right]^{\frac{1}{2}} \right. \\
&\quad \left. + \frac{1}{N} \sum_{j=1}^N \left[\left[\int_{\mathbb{R}^d} \|y^{j,N}\|^{2q} \hat{\mu}_t^{x^{j,N}, j, N}(\mathrm{d}y^{j,N}) \right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z^{j,N}\|^{2q} \hat{\mu}_\infty^{j,N}(\mathrm{d}z^{j,N}) \right]^{\frac{1}{2}} \right] \right].
\end{aligned} \tag{B.26}$$

Similar to before, we will now bound each of these terms. Using Proposition 1 (the moment bounds for the IPS), Proposition 2 (the bounded moments of the invariant measure of the IPS), and arguing precisely as in (B.19), we have for all $i = 1, \ldots, N$,

$$1 + \left[\int_{\mathbb{R}^d} \|y^{i,N}\|^{2q} \hat{\mu}_t^{x^{i,N}, i, N}(\mathrm{d}y^{i,N}) \right]^{\frac{1}{2}} + \left[\int_{\mathbb{R}^d} \|z^{i,N}\|^{2q} \hat{\mu}_\infty^{i,N}(\mathrm{d}z^{i,N}) \right]^{\frac{1}{2}} \leq K_{\theta_0}(1 + \|x^{i,N}\|^q) \tag{B.27}$$

where, importantly, the constant K_{θ_0} is independent of N. Meanwhile, using also Proposition 3 (the exponential contractivity of the IPS), and arguing as in (B.20), we have that, for any $i = 1, \ldots, N$,

$$\begin{aligned}
W_2(\hat{\mu}_t^{x^{i,N}, i, N}, \hat{\mu}_\infty^{i,N}) &\leq \mathbb{W}_2(\hat{\mu}_0^{x^{i,N}, i, N}, \hat{\mu}_\infty^{i,N}) e^{-\lambda_{\theta_0} t} \\
&\leq \left[\int_{\mathbb{R}^d} \|x^{i,N} - y^{i,N}\|^2 \hat{\mu}_\infty^{i,N}(\mathrm{d}y^{i,N}) \right]^{\frac{1}{2}} e^{-\lambda_{\theta_0} t} \\
&\leq 2 \left[\|x^{i,N}\|^2 + \int_{\mathbb{R}^d} \|y^{i,N}\|^2 \hat{\mu}_\infty^{i,N}(\mathrm{d}y^{i,N}) \right]^{\frac{1}{2}} e^{-\lambda_{\theta_0} t} \\
&\leq K_{\theta_0}(1 + \|x^{i,N}\|^q) e^{-\lambda_{\theta_0} t}
\end{aligned} \tag{B.28}$$

for some $q \geq 1$. We remark, once again, that the constants in (B.28) are independent of N. Now, substituting (B.27) and (B.28) into (B.26), using once again the fact that $(a+b)^2 \leq 2(a^2+b^2)$, and as before letting $q' = 2q$, we finally arrive at the bound in (B.15), viz

$$\begin{aligned}
&\left| \mathbb{E}_{\theta_0, \hat{x}^N}[f(\theta, \hat{x}_t^N)] - \int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N) \right| \\
&\leq K_{\theta_0} \left[1 + \|\theta\|^r \right] \left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} \right] e^{-\lambda_{\theta_0} t}.
\end{aligned}$$

It remains to establish the bound in (B.16). Similar to before, the proof of this bound is very similar to the proof of (B.15). Let $(\hat{x}_t^N)_{t \geq 0}$ and $(\hat{y}_t^N)_{t \geq 0}$ be two solutions of the IPS (2) starting

from $\hat{x}^N = (x^{1,N}, \dots, x^{N,N})^\top \in (\mathbb{R}^d)^N$ and $\hat{y}^N = (y^{1,N}, \dots, y^{N,N})^\top \in (\mathbb{R}^d)^N$, respectively. In addition, let $(\hat{\mu}_t^{\hat{x}^N,N})_{t\geq 0}$ and $(\hat{\mu}_t^{\hat{y}^N,N})_{t\geq 0}$ denote the laws of $(\hat{x}_t^N)_{t\geq 0}$ and $(\hat{y}_t^N)_{t\geq 0}$, with marginals $(\hat{\mu}_t^{x^{i,N},i,N})_{t\geq 0}^{i=1:N}$ and $(\hat{\mu}_t^{y^{i,N},i,N})_{t\geq 0}^{i=1:N}$. Finally, let $\hat{\pi}_t^{\hat{x}^N,\hat{y}^N}$ denote an arbitrary coupling of $\hat{\mu}_t^{\hat{x}^N,N}$ and $\hat{\mu}_t^{\hat{y}^N,N}$, and $\hat{\pi}_t^{x^{i,N},y^{i,N}}$ an arbitrary coupling of $\hat{\mu}_t^{x^{i,N},i,N}$ and $\hat{\mu}_t^{y^{i,N},i,N}$. Then, starting from the LHS of (B.16), and making use of the bound in (B.25) with $\mu^N = \hat{\mu}_t^{\hat{x}^N,N}$ and $\nu^N = \hat{\mu}_t^{\hat{y}^N,N}$ we have

$$
\begin{aligned}
&\left|\mathbb{E}_{\theta_0,\hat{x}^N}[f(\theta,\hat{x}_t^N)] - \mathbb{E}_{\theta_0,\hat{y}^N}[f(\theta,\hat{x}_t^N)]\right| \\
&= \left|\int_{(\mathbb{R}^d)^N} f(\theta,\hat{z}_1^N)\hat{\mu}_t^{\hat{x}^N,N}(\mathrm{d}\hat{z}_1^N) - \int_{(\mathbb{R}^d)^N} f(\theta,\hat{z}_2^N)\hat{\mu}_t^{\hat{y}^N,N}(\mathrm{d}\hat{z}_2)\right| \\
&\leq \int_{(\mathbb{R}^d)^N} \left|f(\theta,\hat{z}_1^N) - f(\theta,\hat{z}_2^N)\right|\hat{\pi}_t^{\hat{x}^N,\hat{y}^N}(\mathrm{d}\hat{z}_1^N,\mathrm{d}\hat{z}_2^N) \\
&\leq K\left[1+\|\theta\|^r\right]\left[\left[\int_{\mathbb{R}^d\times\mathbb{R}^d}\|z_1^{i,N}-z_2^{i,N}\|^2\pi_t^{x^{i,N},y^{i,N}}(\mathrm{d}z_1^{i,N},\mathrm{d}z_2^{i,N})\right]^{\frac{1}{2}}\right. \\
&\qquad\left.+\left[\frac{1}{N}\sum_{j=1}^N\int_{\mathbb{R}^d\times\mathbb{R}^d}\|z_1^{j,N}-z_2^{j,N}\|^2\pi_t^{x^{j,N},y^{j,N}}(\mathrm{d}z_1^{j,N},\mathrm{d}z_2^{j,N})\right]^{\frac{1}{2}}\right] \\
&\quad\cdot\left[1+\left[\int_{\mathbb{R}^d}\|z_1^{i,N}\|^{2q}\hat{\mu}_t^{x^{i,N},i,N}(\mathrm{d}z_1^{i,N})\right]^{\frac{1}{2}}+\left[\int_{\mathbb{R}^d}\|z_2^{i,N}\|^{2q}\hat{\mu}_t{}^{y^{i,N},i,N}(\mathrm{d}z_2{}^{i,N})\right]^{\frac{1}{2}}\right. \\
&\quad\left.+\frac{1}{N}\sum_{j=1}^N\left[\int_{\mathbb{R}^d}\|z_1^{j,N}\|^{2q}\hat{\mu}_t^{x^{j,N}j,N}(\mathrm{d}z_1^{j,N})\right]^{\frac{1}{2}}+\frac{1}{N}\sum_{j=1}^N\left[\int_{\mathbb{R}^d}\|z_2^{j,N}\|^{2q}\hat{\mu}_t^{y^{j,N},j,N}(\mathrm{d}z_2^{j,N})\right]^{\frac{1}{2}}\right].
\end{aligned}
$$

Once again, since the coupling was chosen arbitrarily, we in fact have that

$$
\begin{aligned}
&\left|\mathbb{E}_{\theta_0,\hat{x}^N}[f(\theta,\hat{x}_t^N)] - \mathbb{E}_{\theta_0,\hat{y}^N}[f(\theta,\hat{x}_t^N)]\right| \qquad \text{(B.29)}\\
&\leq K\left[1+\|\theta\|^r\right]\left[W_2(\hat{\mu}_t^{x^{i,N},i,N},\hat{\mu}_t^{y^{i,N},i,N}) + \frac{1}{N}\sum_{j=1}^N W_2(\hat{\mu}_t^{x^{j,N},j,N},\hat{\mu}_t^{x^{j,N},j,N})\right] \\
&\quad\cdot\left[1+\left[\int_{\mathbb{R}^d}\|z_1^{i,N}\|^{2q}\hat{\mu}_t^{x^{i,N},i,N}(\mathrm{d}z_1^{i,N})\right]^{\frac{1}{2}}+\left[\int_{\mathbb{R}^d}\|z_2^{i,N}\|^{2q}\hat{\mu}_\infty^{i,N}(\mathrm{d}z_2{}^{i,N})\right]^{\frac{1}{2}}\right. \\
&\quad\left.+\frac{1}{N}\sum_{j=1}^N\left[\left[\int_{\mathbb{R}^d}\|z_1^{j,N}\|^{2q}\hat{\mu}_t^{x^{j,N},j,N}(\mathrm{d}z_1^{j,N})\right]^{\frac{1}{2}}+\left[\int_{\mathbb{R}^d}\|z_2^{j,N}\|^{2q}\hat{\mu}_\infty^{j,N}(\mathrm{d}z_2^{j,N})\right]^{\frac{1}{2}}\right]\right]. \qquad \text{(B.30)}
\end{aligned}
$$

Finally, appealing once more to Proposition 1 (the moment bounds for the IPS) and Proposition 3 (exponential contractivity of the IPS), there exists K_{θ_0}, independent of N, such that, for all $i = 1, \dots, N$,

$$
\left[\int_{\mathbb{R}^d}\|z_1^{i,N}\|^{2q}\hat{\mu}_t^{x^{i,N},i,N}(\mathrm{d}z_1^{i,N})\right]^{\frac{1}{2}} \leq \left[K_{\theta_0}(1+\|x^{i,N}\|^{2q})\right]^{\frac{1}{2}} \leq K_{\theta_0}(1+\|x^{i,N}\|^q) \qquad \text{(B.31)}
$$

$$
\left[\int_{\mathbb{R}^d}\|z_2^{i,N}\|^{2q}\hat{\mu}_t^{y^{i,N},i,N}(\mathrm{d}z_2^{i,N})\right]^{\frac{1}{2}} \leq \left[K_{\theta_0}(1+\|y^{i,N}\|^{2q})\right]^{\frac{1}{2}} \leq K_{\theta_0}(1+\|y^{i,N}\|^q) \qquad \text{(B.32)}
$$

$$\mathbb{W}_2(\hat{\mu}_t^{x,i,N}, \hat{\mu}_t^{y,i,N}) \leq \mathbb{W}_2(\hat{\mu}_0^{x^{i,N},i,N}, \hat{\mu}_0^{y^{i,N},i,N}) e^{-\lambda_{\theta_0} t} \leq \|x^{i,N} - y^{i,N}\| e^{-\lambda_{\theta_0} t}. \tag{B.33}$$

Substituting (B.31), (B.32), and (B.33) into (B.30), and once again letting $q' = q$, we have the desired result, viz

$$\begin{aligned}
&\left| \mathbb{E}_{\theta_0, \hat{x}^N}[f(\theta, \hat{x}_t^N)] - \mathbb{E}_{\theta_0, \hat{y}^N}[f(\theta, \hat{x}_t^N)] \right| \\
&\quad \leq K_{\theta_0} \left[1 + \|\theta\|^r \right] \left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N} \sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2 \right)^{\frac{1}{2}} \right] \\
&\qquad \cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^{q'} \right] e^{-\lambda_{\theta_0} t}. \quad \square
\end{aligned}$$

Lemma 16. *Assume that Conditions A.1 and B.1–B.2 hold. Suppose that, for each $\theta \in \mathbb{R}^p$, and for some $i \in \{1, \dots, N\}$, the function $f(\theta, \hat{x}^N)$ satisfies*

$$\begin{aligned}
\left| f(\theta, \hat{x}^N) - f(\theta, \hat{y}^N) \right| &\leq K \left[1 + \|\theta\|^r \right] \left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N} \sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2 \right)^{\frac{1}{2}} \right] \\
&\quad \cdot \left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^q + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^q \right]
\end{aligned}$$

for some constants $q, r, K > 0$ which are independent of N, where $\hat{x}^N = (x^{1,N}, \dots, x^{N,N})^\top \in (\mathbb{R}^d)^N$. Moreover, suppose that for each $\theta \in \mathbb{R}^p$, $f(\theta, \hat{x}^N)$ is centred, in the sense that $\int_{(\mathbb{R}^d)^N} f(\theta, \hat{x}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) = 0$ for all $\theta \in \mathbb{R}^p$. Then, for all $N \in \mathbb{N}$, and for all $\theta \in \mathbb{R}^p$, the function

$$F(\theta, \hat{x}^N) = \int_0^\infty \mathbb{E}_{\hat{x}^N, \theta_0} \left[f(\theta, \hat{x}_t^N) \right] \mathrm{d}t \tag{B.34}$$

is a well defined, continuous function, of Sobolev class $\cap_{p \geq 1} W_{p,\mathrm{loc}}^2$, which satisfies the Poisson equation

$$\mathcal{A}_{\hat{x}^N, \theta_0} F(\theta, \hat{x}^N) = -f(\theta, \hat{x}^N). \tag{B.35}$$

Moreover, for all $\theta \in \mathbb{R}^p$, $F(\theta, \cdot)$ is centred, in the sense that $\int_{(\mathbb{R}^d)^N} F(\theta, \hat{x}^N) \hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) = 0$, and there exist constants $q', K_{\theta_0} > 0$, independent of N, such that

$$|F(\theta, \hat{x}^N)| \leq K_{\theta_0} \left[1 + \|\theta\|^r \right] \left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} \right], \tag{B.36}$$

$$\|\nabla_{\hat{x}^N} F(\theta, \hat{x})^N\| \leq K_{\theta_0} \left[1 + \|\theta\|^r \right] \left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} \right]. \tag{B.37}$$

Remark. In some sense, this proposition can be seen as a statement of [83, Theorem 1] adapted appropriately to the current setting. In our case, however, since we are interested in the solution of the Poisson equation associated with the generator of the IPS $\hat{x}^N = (x^{1,N}, \dots, x^{N,N}) \in (\mathbb{R}^d)^N$ for any $N \in \mathbb{N}$, additional care is needed to ensure that arguments in the proof of [83,

Theorem 1], in particular those used to establish that the solution is well defined, and that it satisfies the bounds in (B.36)–(B.37), are independent of N. Indeed, we are interested in the solution of this Poisson equation for arbitrarily large N, since we will later take the limit as $N \to \infty$. As an example, if we were to use [83, Theorem 1] directly, we would only have, in place of (B.36), the bound $|F(\theta, \hat{x})| \leq K(1 + \|\hat{x}^N\|^{q'})$, which, due to the $\|\hat{x}^N\|^{q'}$ term, is unbounded in the limit as $N \to \infty$.

Proof. We begin by showing that the function $F(\theta, \hat{x}^N)$ is well defined, and that it satisfies (B.36). Let $\hat{x}_t^N$ denote a solution of the IPS starting from $\hat{x}^N \in (\mathbb{R}^d)^N$. Let $\hat{\mu}_t^N$ denote the law of $\hat{x}_t^N$. Using the bounds in Lemma 15, and that f is centred, we have

$$\begin{aligned}\left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[f(\theta,\hat{x}_t^N)\right]\right| &= \left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[f(\theta,\hat{x}_t^N)\right] - \int_{(\mathbb{R}^d)^N} f(\theta,\hat{z}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N)\right| \\ &\leq K_{\theta_0}\left[1+\|\theta\|^r\right]\left[1+\|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^N \|x^{j,N}\|^{q'}\right]e^{-\lambda_{\theta_0}t}, \end{aligned} \tag{B.38}$$

where, from Lemma 15, the constants $p, q', K_{\theta_0}, \lambda_{\theta_0} > 0$ are all independent of N. Thus, the function F, as defined in (B.34), is absolutely integrable, and thus well defined. Moreover, upon integrating this expression, and using the triangle inequality, we immediately obtain the bound in (B.36).

We next verify that, for all $\theta \in \mathbb{R}^p$, $F(\theta, \cdot)$ is centred. Using (B.34), Fubini's Theorem, and the definition of the invariant measure, in this case we can verify directly that

$$\begin{aligned}\int_{(\mathbb{R}^d)^N} F(\theta,\hat{x}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) &= \int_{(\mathbb{R}^d)^N}\left[\int_0^\infty \mathbb{E}_{\hat{x}^N,\theta_0}\left[f(\theta,\hat{x}_t^N)\right]\mathrm{d}t\right]\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N) \\ &= \int_0^\infty\left[\int_{(\mathbb{R}^d)^N}\mathbb{E}_{\hat{x}^N,\theta_0}\left[f(\theta,\hat{x}_t^N)\right]\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N)\right]\mathrm{d}t \\ &= \int_0^t \underbrace{\left[\int_0^\infty f(\theta,\hat{x}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{x}^N)\right]}_{=0}\mathrm{d}t = 0,\end{aligned}$$

where in the final line we have used the assumption that, for all $\theta \in \mathbb{R}^p$, $f(\theta, \cdot)$ is centred.

We now turn our attention to the bound in (B.37). We will derive this bound directly, rather than appealing to the Sobolev embedding theorem as in [83]. Indeed, the latter argument does not apply in our case given the dependence of the dimensionality of the state space on N. By definition of the gradient, for any $\hat{y}^N \in (\mathbb{R}^d)^N$ such that $\hat{y}^N \to \hat{x}^N$, we have

$$F(\theta,\hat{y}^N) - F(\theta,\hat{x}^N) = \langle \nabla_{\hat{x}^N} F(\theta,\hat{x}^N), \hat{y}^N - \hat{x}^N\rangle + o(\|\hat{y}^N - \hat{x}^N\|). \tag{B.39}$$

Rearranging and then applying the triangle inequality, it follows that

$$|\langle \nabla_{\hat{x}^N} F(\theta,\hat{x}^N), \hat{y}^N - \hat{x}^N\rangle| \leq |F(\theta,\hat{y}^N) - F(\theta,\hat{x}^N)| + \left|o(\|\hat{y}^N - \hat{x}^N\|)\right|.$$

Thus, in particular, we have

$$\lim_{\hat{y}^N\to\hat{x}^N} \frac{|\langle \nabla_{\hat{x}^N} F(\theta,\hat{x}^N), \hat{y}^N - \hat{x}^N\rangle|}{\|\hat{y}^N - \hat{x}^N\|} \leq \lim_{\hat{y}^N\to\hat{x}^N} \frac{|F(\theta,\hat{y}^N) - F(\theta,\hat{x}^N)|}{\|\hat{y}^N - \hat{x}^N\|}. \tag{B.40}$$

Suppose we now consider $\hat{y}_\lambda^N = \hat{x}^N + \lambda\nabla_{\hat{x}^N} F(\theta,\hat{x}^N)$ or, equivalently, $\hat{y}_\lambda^N - \hat{x}^N = \lambda\nabla_{\hat{x}^N} F(\theta,\hat{x}^N)$, for some $\lambda \in \mathbb{R}$. Clearly, we have that $\hat{y}_\lambda^N \to \hat{x}^N$ as $\lambda \to 0$. We also have that

$|\langle \nabla_{\hat{x}^N} F(\theta, \hat{x}^N), \hat{y}_\lambda^N - \hat{x}^N \rangle| = \lambda \|\nabla_{\hat{x}^N} F(\theta, \hat{x}^N)\|^2$, from which it follows that

$$\lim_{\hat{y}_\lambda^N \to \hat{x}^N} \frac{|\langle \nabla_{\hat{x}^N} F(\theta, \hat{x}^N), \hat{y}_\lambda^N - \hat{x}^N \rangle|}{\|\hat{y}_\lambda^N - \hat{x}^N\|} = \lim_{\lambda \to 0} \frac{\lambda \|\nabla_{\hat{x}^N} F(\theta, \hat{x}^N)\|^2}{\lambda \|\nabla_{\hat{x}^N} F(\theta, \hat{x}^N)\|} = \|\nabla_{\hat{x}^N} F(\theta, \hat{x}^N)\|. \quad \text{(B.41)}$$

Substituting (B.41) into the LHS of (B.40), we thus have that

$$\|\nabla_{\hat{x}^N} F(\theta, \hat{x}^N)\| \leq \lim_{\hat{y}_\lambda^N \to \hat{x}^N} \frac{|F(\theta, \hat{y}_\lambda^N) - F(\theta, \hat{x}^N)|}{\|\hat{y}_\lambda^N - \hat{x}^N\|} \leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{|F(\theta, \hat{y}^N) - F(\theta, \hat{x}^N)|}{\|\hat{y}^N - \hat{x}^N\|}. \quad \text{(B.42)}$$

Using (B.42), as well as the bounds from Lemma 15, we then have

$$\begin{aligned}
\|\nabla_{x^N} F(\theta, \hat{x}^N)\| &\leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{|F(\theta, \hat{x}^N) - F(\theta, \hat{y}^N)|}{\|\hat{x}^N - \hat{y}^N\|} \\
&\leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{1}{\|\hat{x}^N - \hat{y}^N\|} \left| \int_0^\infty \mathbb{E}_{\theta_0, \hat{x}^N}\left[f(\theta, \hat{x}_t^N)\right] - \mathbb{E}_{\theta_0, \hat{y}^N}\left[f(\theta, \hat{x}_t^N)\right] \mathrm{d}t \right| \\
&\leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{1}{\|\hat{x}^N - \hat{y}^N\|} K_{\theta_0}\left[1 + \|\theta\|^r\right] \\
&\quad \cdot \left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N} \sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2 \right)^{\frac{1}{2}} \right] \\
&\quad \cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^{q'} \right].
\end{aligned} \quad \text{(B.43)}$$

We next observe that

$$\|x^{i,N} - y^{i,N}\| + \left[\frac{1}{N} \sum_{j=1}^N \|x^{j,N} - y^{j,N}\|^2 \right]^{\frac{1}{2}} \leq \left[1 + \frac{1}{N^{\frac{1}{2}}} \right] \|\hat{x}^N - \hat{y}^N\| \leq K_N \|\hat{x}^N - \hat{y}^N\|, \quad \text{(B.44)}$$

where we have introduced the constant $K_N := 1 + \frac{1}{N^{\frac{1}{2}}} < \infty$ for all $N \in \mathbb{N}$. Substituting (B.44) into (B.43), and absorbing K_N into the existing constant K_{θ_0},[4] we have

$$\begin{aligned}
\|\nabla_{x^N} F(\theta, \hat{x}^N)\| &\leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{1}{\|\hat{x}^N - \hat{y}^N\|} K_{\theta_0} \|\hat{x}^N - \hat{y}^N\| \left[1 + \|\theta\|^r\right] \\
&\quad \cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^{q} \right] \\
&\leq \limsup_{\hat{y}^N \to \hat{x}^N} K_{\theta_0} \left[1 + \|\theta\|^r\right] \Bigg[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} \\
&\qquad + \frac{1}{N} \sum_{j=1}^N \|x^{j,N}\|^{q'} + \frac{1}{N} \sum_{j=1}^N \|y^{j,N}\|^{q'} \Bigg]
\end{aligned}$$

[4] While the constant K_N does depend on the number of particles N, it remains finite as $N \to \infty$, and thus we will suppress this dependence.

$$= K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right], \tag{B.45}$$

where, as elsewhere, the constant K_{θ_0} has been allowed to increase from line to line. This completes the proof of this bound.

The remaining statements in Lemma 16 follow directly from Theorem 1 in [83] (see also Proposition 1 in [84]). In particular, having established suitable bounds on $|F(\theta, \hat{x}^N)|$, the arguments in (c), (d), and (f) in the proof of Theorem 1 in [83] can also be applied in our case. This shows that (B.34) defines a continuous solution of the Poisson equation (B.35), unique in the class of solutions belonging to $\cap_{p\geq 1} W^2_{p,\mathrm{loc}}$. □

Lemma 17. *Assume that Conditions A.1 and B.1–B.2 hold. Suppose that the function $f(\theta, \hat{x}^N)$ is twice differentiable with respect to θ, and for some $i \in \{1, \ldots, N\}$, satisfies, for $k = 1, 2$,*

$$\|\partial_\theta^k f(\theta, \hat{x}^N) - \partial_\theta^k f(\theta, \hat{y}^N)\| \leq K\left[\|x^{i,N} - y^{i,N}\| + \left(\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N} - y^{j,N}\|^2\right)^{\frac{1}{2}}\right]$$
$$\cdot\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^q + \|y^{i,N}\|^q + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^q + \frac{1}{N}\sum_{j=1}^{N}\|y^{j,N}\|^q\right]$$

for some constants $q, r, K > 0$ which are independent of N, where $\hat{x} = (x^{1,N}, \ldots, x^{N,N})^\top \in (\mathbb{R}^d)^N$. Moreover, suppose that, for each $\theta \in \mathbb{R}^p$, $f(\theta, \cdot)$ is centred in the same sense as Lemma 16. Then the solution (B.34) of the Poisson equation (B.35) is twice differentiable with respect to θ. Moreover, there exist constants $q', K_{\theta_0} > 0$, independent of N, such that

$$\sum_{k=0}^{2}\left|\frac{\partial^k}{\partial\theta^k}F(\theta, \hat{x}^N)\right| + \left|\frac{\partial^2}{\partial\hat{x}^N\partial\theta}F(\theta, \hat{x}^N)\right| \leq K_{\theta_0}\left[1+\|\theta\|^r\right]\left[1+\|x^{i,N}\|^{q'}+\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right].$$

Proof. The proof begins with the observation that, for $k = 1, 2$, the function $\partial_\theta^k f(\theta, \hat{x}^N)$ satisfies the bounds required by the second statement of Lemma 15. Thus, from Lemma 15, there exist constants $q', K_{\theta_0} > 0$, which are independent of N, such that

$$\left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{x}_t^N)\right] - \int_{(\mathbb{R}^d)^N}\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{z}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N)\right|$$
$$\leq K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right]e^{-\lambda_{\theta_0}t}. \tag{B.46}$$

Thus, using the fact that $f(\theta, \cdot)$ is centred, arguing similarly to in (B.38), and utilising the bound in (B.46), we have that

$$\left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{x}_t^N)\right]\right| = \left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{x}_t^N)\right] - \frac{\partial^k}{\partial\theta^k}\underbrace{\left[\int_{(\mathbb{R}^d)^N} f(\theta, \hat{z}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N)\right]}_{=0 \text{ since } f \text{ centred}}\right|$$
$$= \left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{x}_t^N)\right] - \int_{(\mathbb{R}^d)^N}\frac{\partial^k}{\partial\theta^k}f(\theta, \hat{z}^N)\hat{\mu}_\infty^N(\mathrm{d}\hat{z}^N)\right|$$

$$\leq K_{\theta_0}\left[1+\|\theta\|^r\right]\left[1+\|x^{i,N}\|^{q'}+\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right]e^{-\lambda_{\theta_0}t}.$$

Thus, due also to Proposition 1 (bounded moments of the IPS), we can differentiate under the integral sign to obtain,

$$\frac{\partial^k}{\partial\theta^k}F(\theta,\hat{x}^N)=\frac{\partial^k}{\partial\theta^k}\int_0^\infty \mathbb{E}_{\hat{x}^N,\theta_0}\left[f(\theta,\hat{x}_t^N)\right]\mathrm{d}t=\int_0^\infty \mathbb{E}_{\hat{x}^N,\theta_0}\left[\frac{\partial^k}{\partial\theta^k}f(\theta,\hat{x}_t^N)\right]\mathrm{d}t. \quad \text{(B.47)}$$

In particular, the required bound now follows straightforwardly as

$$\begin{aligned}\left|\frac{\partial^k F}{\partial\theta^k}(\theta,\hat{x}^N)\right| &\leq \int_0^\infty \left|\mathbb{E}_{\theta_0,\hat{x}^N}\left[\frac{\partial^k f}{\partial\theta^k}(\theta,\hat{x}^N)\right]\right|\mathrm{d}t\\ &\leq K_{\theta_0}\int_0^\infty\left[1+\|\theta\|^r\right]\left[1+\|x^{i,N}\|^{q'}+\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right]e^{-\lambda_{\theta_0}t}\mathrm{d}t\\ &\leq K_{\theta_0}\left[1+\|\theta\|^r\right]\left[1+\|x^{i,N}\|^{q'}+\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right].\end{aligned}$$

Thus, $F(\theta,\hat{x}^N)$ is indeed twice differentiable with respect to θ, and its derivatives satisfy the required bounds.

It remains to obtain the bound on the mixed derivative. In this case we follow similar steps to those used to obtain the bound (B.37) on $|\nabla_{\hat{x}^N}F(\theta,\hat{x}^N)|$ in Lemma 16, viz (B.39)–(B.45). We begin, arguing as in (B.39)–(B.42), with the bound

$$\begin{aligned}\left\|\frac{\partial^2}{\partial\hat{x}^N\partial\theta}F(\theta,\hat{x}^N)\right\| &= \left\|\frac{\partial}{\partial\hat{x}^N}\frac{\partial}{\partial\theta}F(\theta,\hat{x}^N)\right\|\\ &\leq \lim_{\hat{y}_\lambda^N\to\hat{x}^N}\frac{\left|\frac{\partial}{\partial\theta}F(\theta,\hat{y}_\lambda^N)-\frac{\partial}{\partial\theta}F(\theta,\hat{x}^N)\right|}{\|\hat{y}_\lambda^N-\hat{x}^N\|}\\ &\leq \limsup_{\hat{y}^N\to\hat{x}^N}\frac{\left|\frac{\partial}{\partial\theta}F(\theta,\hat{y}^N)-\frac{\partial}{\partial\theta}F(\theta,\hat{x}^N)\right|}{\|\hat{y}^N-\hat{x}^N\|},\end{aligned}$$

where $\hat{y}_\lambda^N$ and $\hat{y}^N$ are as defined in the proof of Lemma 16. It follows, substituting the expression for $\frac{\partial}{\partial\theta}F(\theta,\cdot)$ from (B.47) with $k=1$, that

$$\begin{aligned}&\left\|\frac{\partial^2}{\partial\hat{x}^N\partial\theta}F(\theta,\hat{x}^N)\right\|\\ &\quad\leq \limsup_{\hat{y}^N\to\hat{x}^N}\frac{1}{\|\hat{x}^N-\hat{y}^N\|}\int_0^\infty\left|\mathbb{E}_{\hat{x}^N,\theta_0}\left[\frac{\partial}{\partial\theta}f(\theta,\hat{x}_t^N)\right]-\mathbb{E}_{\hat{y}^N,\theta_0}\left[\frac{\partial}{\partial\theta}f(\theta,\hat{x}_t^N)\right]\right|\mathrm{d}t.\end{aligned} \quad \text{(B.48)}$$

Similar to above, we now appeal to Lemma 15, noting once again that $\partial_\theta f(\theta,\cdot)$ satisfies the bound required by this result. This yields

$$\begin{aligned}&\left|\mathbb{E}_{\theta_0,\hat{x}^N}[\frac{\partial}{\partial\theta}f(\theta,\hat{x}_t^N)]-\mathbb{E}_{\theta_0,\hat{y}^N}[\frac{\partial}{\partial\theta}f(\theta,\hat{x}_t^N)]\right| \quad \text{(B.49)}\\ &\leq K_{\theta_0}\left[1+\|\theta\|^r\right]\left[\|x^{i,N}-y^{i,N}\|+\left(\frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}-y^{j,N}\|^2\right)^{\frac{1}{2}}\right]\end{aligned}$$

$$\cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|y^{j,N}\|^{q'}\right] e^{-\lambda_{\theta_0} t}$$
$$\leq K_{\theta_0}\|\hat{x}^N - \hat{y}^N\|\left[1 + \|\theta\|^r\right]$$
$$\cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|y^{j,N}\|^{q'}\right] e^{-\lambda_{\theta_0}, t}, \tag{B.50}$$

where in the second line we have replaced the first term with the upper bound previously obtained in (B.44). Substituting (B.50) back into (B.48), and integrating, we finally arrive, as required, at

$$\begin{aligned}
\left\|\frac{\partial^2}{\partial \hat{x}^N \partial \theta} F(\theta, \hat{x}^N)\right\| &\leq \limsup_{\hat{y}^N \to \hat{x}^N} \frac{1}{\|\hat{x}^N - \hat{y}^N\|} K_{\theta_0}\|\hat{x}^N - \hat{y}^N\|[1 + \|\theta\|^r] \\
&\quad \cdot \left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|y^{j,N}\|^{q}\right] \\
&\leq \limsup_{\hat{y}^N \to \hat{x}^N} K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^{q'} + \|y^{i,N}\|^{q'}\right. \\
&\qquad \left. + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|y^{j,N}\|^{q'}\right] \\
&= K_{\theta_0}\left[1 + \|\theta\|^r\right]\left[1 + \|x^{i,N}\|^{q'} + \frac{1}{N}\sum_{j=1}^{N}\|x^{j,N}\|^{q'}\right]. \quad \square
\end{aligned}$$

*B.3. Additional lemmas for Theorem 2.2**

Lemma 18. *Assume that Conditions A.1, B.1–B.2, C.1, D.1–D.2, and F.1 hold. Then, for all $q \geq 1$, for all $i = 1, \ldots, N$, for all $N \in \mathbb{N}$, there exists K such that*

$$\sup_{t>0} \mathbb{E}_{\theta_0}\left[\|\theta_t\|^q\right] \leq K_{\theta_0} \quad \textit{and} \quad \sup_{t>0} \mathbb{E}_{\theta_0}\left[\|\theta_t^{i,N}\|^q\right] \leq K_{\theta_0}.$$

Proof. This Lemma follows straightforwardly as an extension of [96, Lemma A.1], making use of the appropriate bounds in Conditions D.1–D.2. $\square$

Appendix C. Verification of conditions for the linear mean field model

In this Appendix, we verify that the conditions of Theorems 2.1, 2.1* and 2.2 hold for the linear one-dimensional mean field model studied in Section 5.1. We also demonstrate that one of the conditions of Theorem 2.2* (Assumption E.2) is not satisfied.

Assumption A.1. We assume that $x_0 \in \mathbb{R}$ and thus this condition is trivially satisfied. We note that this condition would also be satisfied if, for example, $x_0 \sim \mathcal{N}(\mu, \sigma^2)$ for some $\mu, \sigma \in \mathbb{R}$.

Assumption B.1. For this model, we have $b(\theta_0, \cdot) : \mathbb{R} \to \mathbb{R}$, with $b(\theta_0, x) = -\theta_{1,0}x$. This function is Lipschitz continuous with constant θ_1, and satisfies $\langle x - x', b(\theta_0, x) - b(\theta_0, x')\rangle = \langle x - x', -\theta_{1,0}(x - x')\rangle = -\theta_{1,0}\|x - x'\|^2$. This verifies Condition B.1, provided $\theta_{1,0} > 0$.

Assumption B.2. For this model, we have $\phi(\theta_0, \cdot, \cdot) : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, with $\phi(\theta, x, y) = -\theta_{2,0}(x - y)$. This function is twice differentiable with respect to both of its arguments, and is globally Lipschitz with constant $|\theta_{2,0}|$. This verifies Condition B.2, provided $|\theta_{2,0}| \leq \frac{1}{2}\theta_{1,0}$.

Assumption C.1. The functions $b : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and $\phi : \mathbb{R} \times \mathbb{R} \times \mathbb{R}$ are infinitely differentiable with respect to all of their arguments. Moreover, we have that $\|\nabla^i_\theta b(\theta, x)\| = \|\nabla^i_\theta \phi(\theta, x, y)\| = 0$ for $i = 1, 2, 3$. Finally, $\|b(\theta, x) - b(\theta, x')\| \leq \|\theta_1\|\|x - x'\|$, $\|\phi(\theta, x, y) - \phi(\theta, x', y')\| \leq \|\theta_2\|(\|x - y\| + \|x' - y'\|)$. This verifies Condition C.1, provided either $\theta \in \Theta \subset \mathbb{R}^d$ for some compact subset Θ, or we allow for linear growth in $b(\theta, x)$ and $\phi(\theta, x, y)$ (see the remark after Condition C.1).

Assumption D.1. For this model, we have that $L(\theta, x, \mu, \mu) = -\frac{1}{2}[(\theta_1 - \theta_{1,0})x + (\theta_2 - \theta_{2,0})(x - \mathbb{E}_\mu[x])]^2$. For simplicity, let us focus on the 'pure interaction' case, in which $\theta_1 = \theta_{1,0} = 0$. In this case, we have $L(\theta, x, \mu, \mu) = -\frac{1}{2}(\theta_2 - \theta_{2,0})^2(x - \mathbb{E}_\mu[x])^2$, and thus $\nabla_\theta L(\theta, x, \mu, \mu) = -(\theta_2 - \theta_{2,0})(x - \mathbb{E}_\mu[x])^2$. It is then straightforward to compute

$$\langle \nabla_\theta L(\theta, x, \mu, \mu), \theta \rangle = -\theta_2(\theta_2 - \theta_{2,0})(x - \mathbb{E}_\mu[x])^2 = -\big(1 - \frac{\theta_{2,0}}{\theta_2}\big)(x - \mathbb{E}_\mu[x])^2\theta_2^2.$$

It follows that, for all $\|\theta_2\| \geq \|\theta_{2,0}\|$, we have $\langle \nabla_\theta L(\theta, x, \mu, \mu), \theta \rangle \leq -2(x - \mathbb{E}_\mu[x])^2\|\theta_2\|^2 = -\kappa(x, \mu)\|\theta_2\|^2$.

Assumption D.2. For this model, we recall that $\nabla_\theta B(\theta, x, \mu) = [-x, -(x - \mathbb{E}_\mu[x])]$. We thus have

$$\tau(\theta, x, \mu) = \Big\langle \nabla_\theta B(\theta, x, \mu)\nabla_\theta B^T(\theta, x, \mu)\frac{\theta}{\|\theta\|}, \frac{\theta}{\|\theta\|}\Big\rangle^{\frac{1}{2}} = \big[x^2 + (x - \mathbb{E}_\mu(x))^2\big]^{\frac{1}{2}}.$$

Thus implies, in particular, that $|\tau(\theta, x, \mu) - \tau(\theta', x, \mu)| = 0$, which verifies Condition D.2.

Assumption E.1. The unique invariant measure $\hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N)$ of the interacting particle system associated with the linear mean field model is multivariate normal, with mean 0 and covariance Σ^N_∞ given by the solution of the Lyapunov equation (e.g., [85, Chapter 3]) $A\Sigma^N_\infty + \Sigma^N_\infty A^T = I$, where $A = (\theta_{1,0} + \theta_{1,0})1_N - \frac{\theta_{2,0}}{N}J_N$, with I_N the $\mathbb{R}^{N\times N}$ identity matrix, and J_N the $\mathbb{R}^{N\times N}$ matrix of ones. We also have that

$$\hat{L}^{i,N}(\theta, \hat{x}^N) = -\frac{1}{2}\Big[(\theta_1 - \theta_{1,0})x^{i,N} + (\theta_2 - \theta_{2,0})(x^{i,N} - \bar{x}_N)\Big]^2,$$

where $\bar{x}_N = \frac{1}{N}\sum_{j=1}^N x^{j,N}$. We can thus obtain the asymptotic log-likelihood as

$$\begin{aligned}\tilde{\mathcal{L}}^{i,N}(\theta) = &-\frac{1}{2}(\theta_1 - \theta_{1,0})^2 \underbrace{\int_{\mathbb{R}^N}(x^{i,N})^2\hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N)}_{c_1^{i,N}} \\ &- (\theta_1 - \theta_{1,0})(\theta_2 - \theta_{2,0})\underbrace{\int_{\mathbb{R}^N} x^{i,N}\left(x^{i,N} - \bar{x}_N\right)\hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N)}_{c_2^{i,N}} \\ &- \frac{1}{2}(\theta_2 - \theta_{2,0})^2\underbrace{\int_{\mathbb{R}^N}\left(x^{i,N} - \bar{x}_N\right)^2\hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N)}_{c_3^{i,N}}.\end{aligned}$$

The Hessian can now be computed as

$$\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) = \begin{pmatrix} -C_1^{i,N} & -C_2^{i,N} \\ -C_2^{i,N} & -C_3^{i,N} \end{pmatrix}.$$

In order to establish strong concavity, we are required to show that there exists $\eta > 0$ such that, for all $\theta \in \mathbb{R}^2$, $\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) \preceq -\eta I$. That is, the matrix $\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I$ is negative semi-definite. We will show, equivalently, that $\text{trace}(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) < 0$ and $\det(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) > 0$. First note that

$$\text{trace}(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) = -C_1^{i,N} - C_3^{i,N} + 2\eta \leq -C_1^{i,N} + 2\eta.$$

Since Σ^N_∞ is a covariance matrix, we must have $C_1^{i,N} = [\Sigma^N_\infty]_{ii} > 0$. Thus, for any $\eta \in (0, \frac{1}{2}C_1^{i,N})$, we do indeed have $\text{trace}(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) < 0$. We now turn our attention to

$$\det(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) = \eta^2 - \left(C_1^{i,N} + C_2^{i,N}\right)\eta + \left(C_1^{i,N} C_3^{i,N} - (C_2^{i,N})^2\right).$$

The roots of this quadratic are given by

$$\eta_\pm = \frac{\left(C_1^{i,N} + C_2^{i,N}\right) \pm \sqrt{\left(C_1^{i,N} + C_2^{i,N}\right)^2 - 4\left(C_1^{i,N} C_3^{i,N} - (C_2^{i,N})^2\right)}}{2}.$$

These roots are real, since $(C_1^{i,N} + C_2^{i,N})^2 - 4(C_1^{i,N} C_3^{i,N} - (C_2^{i,N})^2) = (C_1^{i,N} - C_2^{i,N})^2 + 2(C_2^{i,N})^2 > 0$. It follows that $\det(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) > 0$ for $\eta \in (-\infty, \eta_-) \cup (\eta_+, \infty)$. We claim that η_- is positive, and thus $\det(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) > 0$ for all $\eta \in (0, \eta_-)$. To see this, note that

$$\begin{aligned} C_1^{i,N} C_3^{i,N} - (C_2^{i,N})^2 = &\int_{\mathbb{R}^N} (x^{i,N})^2 \hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N) \int_{\mathbb{R}^N} \left(x^{i,N} - \bar{x}_N\right)^2 \hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N) \\ &- \left(\int_{\mathbb{R}^N} x^{i,N} \left(x^{i,N} - \bar{x}_N\right) \hat{\mu}^N_\infty(\mathrm{d}\hat{x}^N)\right)^2 > 0 \end{aligned}$$

by the Cauchy–Schwarz inequality. It follows that $(C_1^{i,N} + C_2^{i,N})^2 - 4(C_1^{i,N} C_3^{i,N} - (C_2^{i,N})^2) < (C_1^{i,N} + C_2^{i,N})^2$, and thus η_- is indeed positive. This completes the verification of Condition E.1, since for any $0 < \eta < \min(\frac{1}{2}C_1^{i,N}, \eta_-)$, we have both $\text{trace}(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) < 0$ and $\det(\nabla^2_\theta \tilde{\mathcal{L}}^{i,N}(\theta) + \eta I) > 0$.

Assumption E.2. The unique invariant measure $\mu_\infty(\mathrm{d}x)$ of the linear mean-field model is normal, with mean $\mu_\infty = 0$ and variance $\sigma^2_\infty = \frac{1}{2(\theta_{1,0}+\theta_{2,0})}$ (e.g., [57]). It follows that $L(\theta, x, \mu_\infty, \mu_\infty) = -\frac{1}{2}[(\theta_1 - \theta_{1,0})x + (\theta_2 - \theta_{2,0})x]^2$. We can thus compute the asymptotic contrast function as

$$\underline{\mathcal{L}}(\theta) = -\frac{1}{2}[(\theta_1 + \theta_2) - (\theta_{1,0} + \theta_{2,0})]^2 \int_{\mathbb{R}} x^2 \mu_\infty(\mathrm{d}x) = -\frac{[(\theta_1 + \theta_2) - (\theta_{1,0} + \theta_{2,0})]^2}{4(\theta_{1,0} + \theta_{2,0})}.$$

This function is not strongly concave, and does not admit a unique maximiser. Indeed, it is maximised by any $\theta = (\theta_1, \theta_2)^T$ such that $\theta_1 + \theta_2 = \theta_{1,0} + \theta_{2,0}$. This demonstrates that Condition E.2 is *not* satisfied by the linear mean field model.

Assumption F.1-F.2. Conditions F.1–F.2 are satisfied by $\gamma_t = \min\{\gamma^0, \gamma^0 t^{-\delta}\}$, $\gamma^0 \in [0, \infty)$, $\delta \in (\frac{1}{2}, 1]$.

References

[1] C. Amorino, A. Heidari, V. Pilipauskaitė, M. Podolskij, Parameter estimation of discretely observed interacting particle systems, 2022, arXiv Preprint arXiv:2208.11965.

[2] J. Baladron, D. Fasoli, O. Faugeras, J. Touboul, Mean-field description and propagation of chaos in networks of Hodgkin–Huxley and FitzHugh–Nagumo neurons, J. Math. Neurosci. 2 (1) (2012) 10, http://dx.doi.org/10.1186/2190-8567-2-10.

[3] F. Ball, D. Sirl, Stochastic SIR in structured populations, in: T. Britton, E. Pardoux (Eds.), Stochastic Epidemic Models with Inference, Part II, Springer-Verlag, Cham, Switzerland, 2020, pp. 123–240.

[4] K. Bashiri, On the long-time behaviour of McKean-Vlasov paths, Electron. Commun. Probab. 25 (2020) http://dx.doi.org/10.1214/20-ECP330.

[5] M. Bauer, T. Meyer-Brandis, F. Proske, Strong solutions of mean-field stochastic differential equations with irregular drift, Electron. J. Probab. 23 (2018) 1–35, http://dx.doi.org/10.1214/18-EJP259.

[6] S. Benachour, B. Roynette, D. Talay, P. Vallois, Nonlinear self-stabilizing processes - I existence, invariant probability, propagation of chaos, Stochastic Process. Appl. 75 (2) (1998) 173–201, http://dx.doi.org/10.1016/S0304-4149(98)00018-0.

[7] D. Benedetto, E. Caglioti, M. Pulvirenti, A kinetic equation for granular media, Math. Model. Numer. Anal. 31 (5) (1997) 615–641, http://dx.doi.org/10.1051/m2an/1997310506151.

[8] D.P. Bertsekas, J.N. Tsitsiklis, Gradient convergence in gradient methods with errors, SIAM J. Optim. 10 (3) (2000) 627–642, http://dx.doi.org/10.1137/S1052623497331063.

[9] T. Bhudisaksang, Á. Cartea, Online drift estimation for jump-diffusion processes, Bernoulli 27 (4) (2021) 2494–2518, http://dx.doi.org/10.3150/20-BEJ1319.

[10] J.P.N. Bishwal, Estimation in interacting diffusions: continuous and discrete sampling, Appl. Math. 2 (9) (2011) 1154–1158, http://dx.doi.org/10.4236/am.2011.29160.

[11] F. Bolley, I. Gentil, A. Guillin, Uniform convergence to equilibrium for granular media, Arch. Ration. Mech. Anal. 208 (2) (2013) 429–445, http://dx.doi.org/10.1007/s00205-012-0599-z.

[12] V.S. Borkar, Stochastic Approximation: A Dynamical Systems Viewpoint, Cambridge University Press, Cambridge, UK, 2008, http://dx.doi.org/10.1007/978-93-86279-38-5.

[13] V.S. Borkar, A. Bagchi, Parameter estimation in continuous-time stochastic processes, Stochastics 8 (3) (1982) 193–212, http://dx.doi.org/10.1080/17442508208833238.

[14] C. Brugna, G. Toscani, Kinetic models of opinion formation in the presence of personal conviction., Phys. Rev. E 92 (5) (2015) 52818, http://dx.doi.org/10.1103/PhysRevE.92.052818.

[15] R. Buckdahn, J. Li, J. Ma, A mean-field stochastic control problem with partial observations, Ann. Appl. Probab. 27 (5) (2017) 3201–3245, http://dx.doi.org/10.1214/17-AAP1280.

[16] M. Burger, V. Capasso, D. Morale, On an aggregation model with long and short range interactions, Nonlinear Anal. RWA 8 (3) (2007) 939–958, http://dx.doi.org/10.1016/j.nonrwa.2006.04.002.

[17] A.G. Bustos, Approximations des distributions d'equilibre de certains systems stochastiques avec interactions McKean-Vlasov (Ph.D. thesis), Universite de Nice Sophia Antipolis, 2008.

[18] O.A. Butkovsky, On ergodic properties of nonlinear Markov chains and stochastic McKean–Vlasov equations, Theory Probab. Appl. 58 (4) (2014) 661–674, http://dx.doi.org/10.1137/S0040585X97986825.

[19] C. Canuto, F. Fagnani, P. Tilli, An Eulerian approach to the analysis of Krause's consensus models, SIAM J. Control Optim. 50 (1) (2012) 243–265, http://dx.doi.org/10.1137/100793177.

[20] P. Cardaliaguet, F. Delarue, J.-M. Lasry, P.-L. Lions, The master equation and the convergence problem in mean field games, vol. 2, Princeton University Press, 2019, http://dx.doi.org/10.2307/j.ctvckq7qf.

[21] P. Cardaliaguet, C.-A. Lehalle, Mean field game of controls and an application to trade crowding, Math. Financ. Econ. 12 (3) (2018) 335–363, http://dx.doi.org/10.1007/s11579-017-0206-z.

[22] R. Carmona, F. Delarue, Probabilistic Theory of Mean Field Games with Applications I, Springer-Verlag, Cham, Switzerland, 2018, http://dx.doi.org/10.1007/978-3-319-58920-6.

[23] J.A. Carrillo, R.S. Gvalani, G.A. Pavliotis, A. Schlichting, Long-time behaviour and phase transitions for the McKean-Vlasov equation on the torus, Arch. Ration. Mech. Anal. 235 (1) (2020) 635–690, http://dx.doi.org/10.1007/s00205-019-01430-4, URL https://doi.org/10.1007/s00205-019-01430-4.

[24] J.A. Carrillo, R. McCann, C. Villani, Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates., Rev. Mat. Iberoam. 19 (3) (2003) 971–1018, http://dx.doi.org/10.4171/RMI/376.

[25] J.A. Carrillo, R.J. McCann, C.D. Villani, Contractions in the 2-wasserstein length space and thermalization of granular media, Arch. Ration. Mech. Anal. 179 (2) (2006) 217–263, http://dx.doi.org/10.1007/s00205-005-0386-1.

[26] P. Cattiaux, A. Guillin, F. Malrieu, Probabilistic approach for granular media equations in the non-uniformly convex case, Probab. Theory Related Fields 140 (1) (2008) 19–40, http://dx.doi.org/10.1007/s00440-007-0056-3.

[27] P.E. Chaudru de Raynal, Strong well posedness of McKean–Vlasov stochastic differential equations with Hölder drift, Stochastic Process. Appl. 130 (1) (2020) 79–107, http://dx.doi.org/10.1016/j.spa.2019.01.006.

[28] B. Chazelle, Q. Jiu, Q. Li, C. Wang, Well-posedness of the limiting equation of a noisy consensus model in opinion dynamics, J. Differential Equations 263 (1) (2017) 365–397, http://dx.doi.org/10.1016/j.jde.2017.02.036.

[29] X. Chen, Maximum likelihood estimation of potential energy in interacting particle systems from single-trajectory data, Electron. Commun. Probab. 26 (2021) 1–13, http://dx.doi.org/10.1214/21-ECP416.

[30] H.F. Chen, C. Wu, Y. Zhu, Continuous-time stochastic approximation procedure with randomly varying truncations, Acta Math. Sci. 7 (1) (1987) 43–55, http://dx.doi.org/10.1016/S0252-9602(18)30487-9.

[31] F. Comte, V. Genon-Catalot, Nonparametric adaptive estimation for interacting particle systems, 2022, HAL Preprint.

[32] D. Crisan, J. Xiong, Approximate mckean-vlasov representations for a class of SPDEs, Stochastics 82 (1) (2010) 53–68, http://dx.doi.org/10.1080/17442500902723575.

[33] P. Del Moral, J. Tugaut, Uniform propagation of chaos and creation of chaos for a class of nonlinear diffusions, Stoch. Anal. Appl. 37 (6) (2019) 909–935, http://dx.doi.org/10.1080/07362994.2019.1622426.

[34] M.G. Delgadino, R.S. Gvalani, G.A. Pavliotis, S.A. Smith, Phase transitions, logarithmic Sobolev inequalities, and uniform-in-time propagation of chaos for weakly interacting diffusions, Commun. Math. Phys. (2023) http://dx.doi.org/10.1007/s00220-023-04659-z.

[35] L. Della Maestra, M. Hoffmann, Nonparametric estimation for interacting particle systems: Mckean-vlasov models, Probab. Theory Related Fields 182 (1) (2022) 551–613, http://dx.doi.org/10.1007/s00440-021-01044-6.

[36] J.L. Doob, Stochastic Processes, Wiley, New York, 1953.

[37] A. Durmus, A. Eberle, A. Guillin, R. Zimmer, An elementary approach to uniform in time propagation of chaos, Proc. Amer. Math. Soc. 148 (2020) 5387–5398, http://dx.doi.org/10.1090/proc/14612.

[38] A. Eberle, A. Guillin, R. Zimmer, Quantitative harris-type theorems for diffusions and mckean-vlasov processes, Trans. Amer. Math. Soc. 371 (2019) 7135–7173, http://dx.doi.org/10.1090/tran/7576.

[39] N. Fournier, A. Guillin, On the rate of convergence in wasserstein distance of the empirical measure, Probab. Theory Related Fields 162 (3) (2015) 707–738, http://dx.doi.org/10.1007/s00440-014-0583-7.

[40] T. Funaki, A certain class of diffusion processes associated with nonlinear parabolic equations, Zeitschrift FÜR Wahrscheinlichkeitstheorie Und Verwandte Gebiete 67 (3) (1984) 331–348, http://dx.doi.org/10.1007/BF00535008.

[41] J. Garnier, G. Papanicolaou, T.-W. Yang, Consensus convergence with stochastic effects, Vietnam J. Math. 45 (1) (2017) 51–75, http://dx.doi.org/10.1007/s10013-016-0190-2.

[42] V. Genon-Catalot, C. Larédo, Parametric inference for small variance and long time horizon McKean-Vlasov diffusion models, Electron. J. Stat. 15 (2) (2021a) 5811–5854, http://dx.doi.org/10.1214/21-EJS1922.

[43] V. Genon-Catalot, C. Larédo, Probabilistic properties and parametric inference of small variance nonlinear self-stabilizing stochastic differential equations, Stochastic Process. Appl. 142 (2021b) 513–548, http://dx.doi.org/10.1016/j.spa.2021.09.002.

[44] V. Genon-Catalot, C. Larédo, Inference for ergodic McKean–Vlasov stochastic differential equations with polynomial interactions, 2022, HAL Preprint.

[45] K. Giesecke, G. Schwenkler, J.A. Sirignano, Inference for large financial systems, Math. Finance 30 (1) (2020) 3–46, http://dx.doi.org/10.1111/mafi.12222.

[46] B.D. Goddard, B. Gooding, H. Short, G.A. Pavliotis, Noisy bounded confidence models for opinion dynamics: the effect of boundary conditions on phase transitions, IMA J. Appl. Math. 87 (1) (2022) 80–110, http://dx.doi.org/10.1093/imamat/hxab044.

[47] S.N. Gomes, A.M. Stuart, M.-T. Wolfram, Parameter estimation for macroscopic pedestrian dynamics models from microscopic data, SIAM J. Appl. Math. 79 (4) (2019) 1475–1500, http://dx.doi.org/10.1137/18M1215980.

[48] W. Hammersley, D. Šiška, L. Szpruch, McKean-Vlasov SDEs under measure dependent Lyapunov conditions, Ann. Inst. H. Poincaré Probab. Statist. 57 (2) (2021) 1032–1057, http://dx.doi.org/10.1214/20-AIHP1106.
[49] S. Herrmann, J. Tugaut, Non-uniqueness of stationary measures for self-stabilizing processes, Stochastic Process. Appl. 120 (7) (2010) 1215–1246, http://dx.doi.org/10.1016/j.spa.2010.03.009.
[50] C.C. Heyde, A quasi-likelihood approach to estimating parameters in diffusion-type processes, J. Appl. Probab. 31 (1994) 283–290, http://dx.doi.org/10.2307/3214962.
[51] K. Hu, Z. Ren, D. Siska, L. Szpruch, Mean-field langevin dynamics and energy landscape of neural networks, 2020, arXiv Preprint arXiv:1905.07769.
[52] X. Huang, F.-Y. Wang, Distribution dependent SDEs with singular coefficients, Stochastic Process. Appl. 129 (11) (2019) 4747–4770, http://dx.doi.org/10.1016/j.spa.2018.12.012.
[53] J.E. Hutton, P.I. Nelson, Quasi-likelihood estimation for semimartingales, Stochastic Process. Appl. 22 (2) (1986) 245–257, http://dx.doi.org/10.1016/0304-4149(86)90004-9.
[54] N. Ikeda, S. Watanabe, A comparison theorem for solutions of stochastic differential equations and its applications, Osaka J. Math. 14 (3) (1977) 619–633, http://dx.doi.org/10.18910/7664.
[55] B. Jourdain, S. Meleard, W.A. Woyczynski, Nonlinear SDEs driven by levy processes and related PDEs, Alea 4 (2008) 1–29.
[56] M. Kac, Foundations of kinetic theory, in: Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 3: Contributions To Astronomy and Physics, Berkeley, CA, 1956, pp. 171–197.
[57] R.A. Kasonga, Maximum likelihood theory for large interacting systems, SIAM J. Appl. Math. 50 (3) (1990) 865–875, http://dx.doi.org/10.1137/0150050.
[58] R. Khasminskii, Stochastic Stability of Differential Equations, second ed., Springer-Verlag, Berlin, Heidelberg, 2012, http://dx.doi.org/10.1007/978-3-642-23280-0.
[59] Y.A. Kutoyants, Parameter Estimation for Stochastic Processes, Heldermann, 1984, p. 206.
[60] Y.A. Kutoyants, Statistical Inference for Ergodic Diffusion Processes, Springer-Verlag, London, 2004, http://dx.doi.org/10.1007/978-1-4471-3866-2.
[61] D. Lacker, On a strong form of propagation of chaos for McKean–Vlasov equations, Electron. Commun. Probab. 23 (2018) 1–11, http://dx.doi.org/10.1214/18-ECP150.
[62] Q. Lang, F. Lu, Identifiability of interaction kernels in mean-field equations of interacting particles, 2021, arXiv Preprint arXiv:2106.05565.
[63] D. Levanony, A. Shwartz, O. Zeitouni, Recursive identification in continuous-time stochastic processes, Stochastic Process. Appl. 49 (2) (1994) 245–275, http://dx.doi.org/10.1016/0304-4149(94)90137-6.
[64] J. Li, H. Min, Weak solutions of mean-field stochastic differential equations and application to zero-sum stochastic differential games, SIAM J. Control Optim. 54 (3) (2016) 1826–1858, http://dx.doi.org/10.1137/15M1015583.
[65] R.S. Liptser, A.N. Shiryaev, Statistics of Random Processes, second ed., Springer-Verlag, Berlin, Heidelberg, 2001, http://dx.doi.org/10.1007/978-3-662-13043-8.
[66] Q. Liu, Stein variational gradient descent: a general purpose bayesian inference algorithm, in: Proceedings of the 30th Conference on Neural Information Processing Systems, NIPS 2016, Barcelona, Spain, 2016.
[67] M. Liu, H. Qiao, Parameter estimation of path-dependent mckean-vlasov stochastic differential equations, 2020, arXiv Preprint arXiv:2004.09580.
[68] W. Liu, L. Wu, C. Zhang, Long-time behaviors of mean-field interacting particle systems related to mckean-vlasov equations, Comm. Math. Phys. 387 (1) (2021) 179–214, http://dx.doi.org/10.1007/s00220-021-04198-5.
[69] F. Lu, M. Maggioni, S. Tang, Learning interaction kernels in heterogeneous systems of agents from multiple trajectories, J. Mach. Learn. Res. 22 (32) (2021) 1–67.
[70] F. Lu, M. Zhong, S. Tang, M. Maggioni, Nonparametric inference of interaction laws in systems of agents from trajectory data, Proc. Natl. Acad. Sci. 116 (29) (2019) 14424–14433, http://dx.doi.org/10.1073/pnas.1822012116.
[71] L.D. Maestra, M. Hoffmann, The LAN property for McKean-Vlasov models in a mean-field regime, Stochastic Process. Appl. 155 (2023) 109–146, http://dx.doi.org/10.1016/j.spa.2022.10.002.
[72] F. Malrieu, Logarithmic Sobolev inequalities for some nonlinear PDE's, Stochastic Process. Appl. 95 (1) (2001) 109–132, http://dx.doi.org/10.1016/S0304-4149(01)00095-3.

[73] F. Malrieu, Convergence to equilibrium for granular media equations and their Euler schemes, Ann. Appl. Probab. 13 (2) (2003) 540–560, http://dx.doi.org/10.1214/aoap/1050689593.

[74] H.P. McKean, A class of Markov processes associated with nonlinear parabolic equations, Proc. Natl. Acad. Sci. USA 56 (6) (1966) 1907–1911, http://dx.doi.org/10.1073/pnas.56.6.1907.

[75] H.P. McKean, Propagation of chaos for a class of non-linear parabolic equations, in: Lecture Series in Differential Equations, vol. 7, Catholic University, 1967, pp. 41–57.

[76] S. Mei, A. Montanari, P.-M. Nguyen, A mean field view of the landscape of two-layer neural networks, Proc. Natl. Acad. Sci. 115 (33) (2018) E7665 LP – E7671, http://dx.doi.org/10.1073/pnas.1806579115.

[77] S. Méléard, Asymptotic behaviour of some interacting particle systems; McKean–Vlasov and Boltzmann models, in: C. Graham, T.G. Kurtz, S. Méléard, P.E. Protter, M. Pulvirenti, D. Talay, D. Talay, L. Tubaro (Eds.), Probabilistic Models Nonlinear Partial Differ. Equ. 1627 (1996) 42–95, http://dx.doi.org/10.1007/BFb0093177.

[78] D.A. Messenger, D.M. Bortz, Learning mean-field equations from particle data using WSINDy, Physica D 439 (2022) 133406, http://dx.doi.org/10.1016/j.physd.2022.133406.

[79] Y.S. Mishura, A.Y. Veretennikov, Existence and uniqueness theorems for solutions of McKean–Vlasov stochastic equations, Theory Probab. Math. Statist. 103 (2020) 59–101, http://dx.doi.org/10.1090/tpms/1135.

[80] A. Mogilner, L. Edelstein-Keshet, A non-local model for a swarm, J. Math. Biol. 38 (6) (1999) 534–570, http://dx.doi.org/10.1007/s002850050158.

[81] S. Motsch, E. Tadmor, Heterophilious dynamics enhances consensus, SIAM Rev. 56 (4) (2014) 577–621, http://dx.doi.org/10.1137/120901866.

[82] K. Oelschlager, A martingale approach to the law of large numbers for weakly interacting stochastic processes, Ann. Probab. 12 (2) (1984) 458–479, http://dx.doi.org/10.1214/aop/1176993301.

[83] E. Pardoux, A.Y. Veretennikov, On the Poisson equation and diffusion approximation 1, Ann. Probab. 29 (3) (2001) 1061–1085, http://dx.doi.org/10.1214/aop/1015345596.

[84] E. Pardoux, A.Y. Veretennikov, On the Poisson equation and diffusion approximation 2, Ann. Probab. 31 (3) (2003) 1166–1192, http://dx.doi.org/10.1214/aop/1015345596.

[85] G.A. Pavliotis, Stochastic Processes and Applications, Springer-Verlag, New York, 2014, http://dx.doi.org/10.1007/978-1-4939-1323-7.

[86] G.A. Pavliotis, A. Zanoni, Eigenfunction martingale estimators for interacting particle systems and their mean field limit, SIAM J. Appl. Dyn. Syst. 21 (4) (2022) 2338–2370, http://dx.doi.org/10.1137/21M1464348.

[87] D. Revuz, M. Yor, Continuous Martingales and Brownian Motion, third ed., Springer-Verlag, Berlin, Heidelberg, 1999, http://dx.doi.org/10.1007/978-3-662-06400-9.

[88] H. Robbins, S. Monro, A stochastic approximation method, Ann. Math. Stat. 22 (3) (1951) 400–407.

[89] G.M. Rotskoff, E. Vanden-Eijnden, Trainability and accuracy of neural networks: an interacting particle system approach, Commun. Pure Appl. Math. 75 (9) (2022) 1889–1935, http://dx.doi.org/10.1002/cpa.22074.

[90] L. Sharrock, Two-timescale stochastic approximation for bilevel optimisation problems in continuous-time models, in: ICML 2022 Workshop on Continuous Time Methods for Machine Learning, 2022, http://dx.doi.org/10.48550/arXiv.2206.06995.

[91] L. Sharrock, N. Kantas, Joint online parameter estimation for the partially observed stochastic advection diffusion equation, SIAM /ASA J. Uncert. Quant. 10 (1) (2022) 55–95, http://dx.doi.org/10.1137/20M1375073.

[92] L. Sharrock, N. Kantas, Two-timescale stochastic gradient descent in continuous time with applications to joint online parameter estimation and optimal sensor placement, Bernoulli 29 (2) (2023) 1137–1165, http://dx.doi.org/10.3150/22-BEJ1493.

[93] L. Sharrock, N. Kantas, P. Parpas, G.A. Pavliotis, Parameter estimation for the McKean-Vlasov stochastic differential equation, 2021, arXiv Preprint arXiv:2106.13751.

[94] J. Sirignano, K. Spiliopoulos, Stochastic gradient descent in continuous time, SIAM J. Financial Math. 8 (1) (2017) 933–961, http://dx.doi.org/10.1137/17M1126825.

[95] J. Sirignano, K. Spiliopoulos, Mean field analysis of neural networks: a law of large numbers, SIAM J. Appl. Math. 80 (2) (2020a) 725–752, http://dx.doi.org/10.1137/18M1192184.

[96] J. Sirignano, K. Spiliopoulos, Stochastic gradient descent in continuous time: a central limit theorem, Stoch. Syst. 10 (2) (2020b) 124–151, http://dx.doi.org/10.1287/stsy.2019.0050.

[97] S.C. Surace, J. Pfister, Online maximum-likelihood estimation of the parameters of partially observed diffusion processes, IEEE Trans. Automat. Control 64 (7) (2019) 2814–2829, http://dx.doi.org/10.1109/TAC.2018.2880404.

[98] A.-S. Sznitman, Topics in propagation of chaos, ecole d'eté de probabilités de saint-flour XIX - 1989, in: D.L. Burkholder, E. Pardoux, A.-S. Sznitman, P.-L. Hennequin (Eds.), in: Lecture Notes in Mathematics, vol. 1464, Springer-Verlag, Berlin, Heidelberg, 1991, pp. 165–251.

[99] Y. Tamura, On asymptotic behaviors of the solution of a nonlinear diffusion equation, J. Fac. Sci. Uni. Tokyo 31 (1984) 195–221.

[100] J. Tugaut, Convergence to the equilibria for self-stabilizing processes in double-well landscape, Ann. Probab. 41 (3A) (2013) 1427–1460, http://dx.doi.org/10.1214/12-AOP749.

[101] J. Tugaut, Phase transitions of McKean-Vlasov processes in double-wells landscape, Stochastics 86 (2) (2014) 257–284, http://dx.doi.org/10.1080/17442508.2013.775287.

[102] A.Y. Veretennikov, On ergodic measures for McKean–Vlasov stochastic equations, in: H. Niederreiter, D. Talay (Eds.), Monte Carlo and Quasi-Monte Carlo Methods 2004, Springer-Verlag, Berlin, Heidelberg, 2006, pp. 471–486, http://dx.doi.org/10.1007/3-540-31186-6_29.

[103] F.-Y. Wang, Distribution dependent SDEs for Landau type equations, Stochastic Process. Appl. 128 (2) (2018) 595–621, http://dx.doi.org/10.1016/j.spa.2017.05.006.

[104] A.J. Weir, Lebesgue Integration and Measure, Cambridge University Press, Cambridge, UK, 1973, http://dx.doi.org/10.1017/9781139171458.

[105] J. Wen, X. Wang, S. Mao, X. Xiao, Maximum likelihood estimation of McKean–Vlasov stochastic differential equation and its application, Appl. Math. Comput. 274 (2016) 237–246, http://dx.doi.org/10.1016/j.amc.2015.11.019.

[106] T. Yamada, On a comparison theorem for solutions of stochastic differential equations and its applications, J. Math. Kyoto Univ. 13 (3) (1973) 497–512, http://dx.doi.org/10.1215/kjm/1250523321.

[107] R. Yao, X. Chen, Y. Yang, Mean-field nonparametric estimation of interacting particle systems, in: Proceedings of Thirty Fifth Conference on Learning Theory, COLT, London, UK, 2022, https://proceedings.mlr.press/v178/yao22a.html.