

LARGE DEVIATION PROPERTIES OF WEAKLY INTERACTING PROCESSES VIA WEAK CONVERGENCE METHODS

BY AMARJIT BUDHIRAJA¹, PAUL DUPUIS² AND MARKUS FISCHER³

*University of North Carolina at Chapel Hill, Brown University
and Brown University*

We study large deviation properties of systems of weakly interacting particles modeled by Itô stochastic differential equations (SDEs). It is known under certain conditions that the corresponding sequence of empirical measures converges, as the number of particles tends to infinity, to the weak solution of an associated McKean–Vlasov equation. We derive a large deviation principle via the weak convergence approach. The proof, which avoids discretization arguments, is based on a representation theorem, weak convergence and ideas from stochastic optimal control. The method works under rather mild assumptions and also for models described by SDEs not of diffusion type. To illustrate this, we treat the case of SDEs with delay.

1. Introduction. Collections of weakly interacting random processes have long been of interest in statistical physics and more recently have appeared in problems of engineering and operations research. A simple but important example of such a collection is a group of “particles,” each of which evolves according to the solution of an Itô-type stochastic differential equation (SDE). All particles have the same functional form for the drift and diffusion coefficients. The coefficients of particle i are, as usual, allowed to depend on the current state of particle i , but also depend on the current empirical distribution of all particle locations. When the number of particles is large the contribution of any given particle to the empirical distribution is small, and in this sense the interaction between any two particles is considered “weak.”

Received December 2009; revised October 2010.

¹Supported in part by the Army Research Office (W911NF-0-1-0080, W911NF-10-1-0158), NSF Grant DMS-10-04418 and the US-Israel Binational Science Foundation 2008466.

²Supported in part by NSF Grant DMS-07-06003, the Army Research Office W911NF-09-1-0155 and the Air Force Office of Scientific Research FA9550-09-1-0378.

³Supported by the German Research Foundation (DFG research fellowship), NSF Grant DMS-07-06003 and the Air Force Office of Scientific Research (FA9550-07-1-0544, FA9550-09-1-0378).

MSC2010 subject classifications. Primary 60F10, 60K35; secondary 60B10, 60H10, 34K50, 93E20.

Key words and phrases. Large deviations, interacting random processes, McKean–Vlasov equation, stochastic differential equation, delay, weak convergence, martingale problem, optimal stochastic control.

For various reasons, including model simplification and approximation, one may consider a functional law of large numbers (LLN) limit as the number of particles tends to infinity. The limit behavior of a single particle (under assumptions which guarantee that all particles are in some sense exchangeable) can be described by a two component Markov process. One component corresponds to the state of a typical particle, while the second corresponds to the limit of the empirical measures. Again using that all particles are exchangeable, under appropriate conditions one can show that the second component coincides with the distribution of the particle component. The limit process, which typically has an infinite-dimensional state, is sometimes referred to as a “nonlinear diffusion.” Because the particle’s own distribution appears in the state dynamics, the partial differential equations that characterize expected values and densities associated with this process are nonlinear, and hence the terminology.

In this paper we consider the large deviation properties of the particle system as the number of particles tends to infinity. Thus the deviations we study are those of the empirical measure of the prelimit process from the distribution of the nonlinear diffusion. Of particular interest, and a subject for further study, are deviations when the initial distribution of the single particle in the nonlinear diffusion is invariant under the joint particle/measure dynamics, and related questions of stability for both the limit and prelimit processes.

One of the basic references for large deviation results for weakly interacting diffusions is [10]. This paper considers a system of uniformly nondegenerate diffusions with interaction in the drift term and establishes a large deviation principle for the empirical measure using discretization arguments and careful exponential probability estimates (see Section 7.1). Properties related to a large deviation principle such as fluctuation theorems have been studied in [2, 3, 21, 26, 33]. A proof of the large deviation principle for systems with constant diffusion coefficient that is based on a comparison result for a related infinite-dimensional Hamilton–Jacobi–Bellman equation appears in [17], Section 13.3.

Later works have developed the theory for a variety of alternative models, including multilevel large deviations [11, 13], jump diffusions [24, 25], discrete-time systems [9, 12] and interacting diffusions with random interaction coefficients [1] or singular interaction [18]. In the current work we develop an approach which is very different from the one taken in any of these papers. Our proofs do not involve any time or space discretization of the system, and no exponential probability estimates are invoked. The main ingredients in the proof are weak convergence methods for functional occupation measures and certain variational representation formulas. Our proofs cover models with degenerate noise and allow for interaction in both drift and diffusion terms. In fact, the techniques are applicable to a wide range of model settings, and an example of stochastic delay equations is considered in Section 7 to illustrate the possibilities.

The starting point of our analysis is a variational representation for moments of nonnegative functionals of a Brownian motion [5]. Using this representation, the

proof of the large deviation principle reduces to the study of asymptotic properties of certain controlled versions of the original process. The key step in the proof is to characterize the weak limits of the control and controlled process as the large deviation parameter tends to its limit and under the same scaling that applies to the original process. More precisely, one needs to characterize the limit of the empirical measure of a large collection of controlled and weakly interacting processes. In the absence of control this characterization problem reduces to an LLN analysis of the original particle system, which has been studied extensively [19, 20, 27]. Our main tools for the study of the controlled analog are functional occupation measure methods. Indeed, these methods have been found to be quite useful for the study of averaging problems, but where the average is with respect to a time variable [23]. In the problem studied here the measure-valued processes of interest are obtained using averaging over particles rather than the time variable.

The approach presented here can be applied to interacting systems driven by general continuous time processes with jumps provided the systems are scaled in the right way. Indeed, the driving noise process could be a Brownian motion plus an independent Poisson random measure. A key step to make the approach work is a variational representation of Poisson functionals, which has recently been established in [8].

Finally, we remark that variational representations for Brownian motions and Poisson random measures [6–8] have proved to be useful for the study of small-noise large-deviation problems, and many recent papers have applied these results to a variety of infinite-dimensional small-noise systems. A small selection is [14, 29–31] (see [8] for a more complete list). We expect the current work to be similarly a starting point for the study, using variational representations, of a rather different collection of large deviation problems, namely asymptotics of a large number of interacting particles.

An outline of the paper is as follows. In Section 2 we introduce the interacting SDE particle model, the related controlled and LLN limit versions and discuss the relevant topologies and sense of uniqueness of solutions. Section 3 discusses the relation between Laplace and large-deviation principles, states assumptions and the main result of the paper and then outlines how this result will be proved using a representation theorem. In Section 4 we describe the martingale problems that will be used in the proof. The proof itself is divided into lower and upper bounds in Sections 5 and 6, respectively. The constructions in the proof are set up to handle a more general case than just the model introduced in Section 2, and in Section 7 we use this generality to state and prove a large deviation theorem for systems with delay. This section also reviews the prior work of [10]. The Appendix contains the proof of a technical point that was deferred for reasons of exposition.

2. The model. For each $N \in \mathbb{N}$, the N -particle prelimit model is described in terms of a system of N weakly coupled d -dimensional stochastic differential equations (SDEs). The system is considered over the fixed finite interval $[0, T]$.

Set $\mathcal{X} \doteq \mathbf{C}([0, T], \mathbb{R}^d)$, and equip \mathcal{X} with the maximum norm, which is denoted by $\|\cdot\|$. Similarly, set $\mathcal{W} \doteq \mathbf{C}([0, T], \mathbb{R}^{d_1})$ and equip \mathcal{W} with the maximum norm. Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space, and suppose that on this space there is a filtration (\mathcal{F}_t) satisfying the usual conditions [i.e., (\mathcal{F}_t) is right-continuous and \mathcal{F}_0 contains all \mathbf{P} -negligible sets], as well as a collection $\{W^i, i \in \mathbb{N}\}$ of independent standard d_1 -dimensional (\mathcal{F}_t) -Wiener processes.

Let b and σ be Borel measurable functions defined on $\mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$ taking values in \mathbb{R}^d and the space of real $d \times d_1$ -matrices, respectively. If $(\mathcal{S}, d_{\mathcal{S}})$ is a metric space, then $\mathcal{P}(\mathcal{S})$ denotes the space of probability measures on the Borel σ -field $\mathcal{B}(\mathcal{S})$. The space $\mathcal{P}(\mathcal{S})$ is equipped with the topology of weak convergence, which can be metricized, using, for example, the bounded Lipschitz metric, making it a Polish space.

The evolution of the state of the particles in the N -particle model is given by the solution to the system of SDEs

$$(2.1) \quad \begin{aligned} dX^{i,N}(t) &= b(X^{i,N}(t), \mu^N(t)) dt + \sigma(X^{i,N}(t), \mu^N(t)) dW^i(t), \\ X^{i,N}(0) &= x^{i,N}, \end{aligned}$$

where $x^{i,N} \in \mathbb{R}^d$, $i \in \{1, \dots, N\}$, and

$$\mu^N(t, \omega) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{X^{i,N}(t, \omega)}, \quad \omega \in \Omega,$$

is the empirical measure of $(X^{1,N}(t), \dots, X^{N,N}(t))$ for $t \in [0, T]$. By construction, $\mu^N(t)$ is a $\mathcal{P}(\mathbb{R}^d)$ -valued random variable. Denote by μ^N the empirical measure of $(X^{1,N}, \dots, X^{N,N})$ over the time interval $[0, T]$, that is, μ^N is the $\mathcal{P}(\mathcal{X})$ -valued random variable defined by

$$\mu_{\omega}^N \doteq \frac{1}{N} \sum_{i=1}^N \delta_{X^{i,N}(\cdot, \omega)}, \quad \omega \in \Omega.$$

Clearly, the distribution of $\mu^N(t)$ is identical to the marginal distribution of μ^N at time t , that is, $\mu^N(t) = \mu^N \circ \pi_t^{-1}$ where $\pi_t: \mathcal{X} \rightarrow \mathbb{R}^d$ is the projection map corresponding to the value at time t .

Our aim is to establish a Laplace principle for the family $\{\mu^N, N \in \mathbb{N}\}$ of $\mathcal{P}(\mathcal{X})$ -valued random variables. When $\frac{1}{N} \sum_{i=1}^N \delta_{x^{i,N}}$ converges weakly to ν_0 for some $\nu_0 \in \mathcal{P}(\mathbb{R}^d)$, the asymptotic behavior of μ^N as N tends to infinity can be characterized in terms of solutions to the nonlinear diffusion

$$(2.2) \quad \begin{aligned} dX(t) &= b(X(t), \text{Law}(X(t))) dt + \sigma(X(t), \text{Law}(X(t))) dW(t), \\ X(0) &\sim \nu_0, \end{aligned}$$

where W is a standard d_1 -dimensional Wiener process. Thus we are interested in the study of deviations of μ^N , N large, from its typical behavior, namely the probability law of the process solving (2.2).

In the formulation and proof of the Laplace principle, we will need to consider a controlled version of (2.1). For $N \in \mathbb{N}$, let \mathcal{U}_N be the space of all (\mathcal{F}_t) -progressively measurable functions $u : [0, T] \times \Omega \rightarrow \mathbb{R}^{N \times d_1}$ such that

$$\mathbf{E} \left[\int_0^T |u(t)|^2 dt \right] < \infty,$$

where \mathbf{E} denotes expectation with respect to \mathbf{P} , and $|\cdot|$ denotes the Euclidean norm of appropriate dimension. For $u \in \mathcal{U}_N$, we sometimes write $u = (u_1, \dots, u_N)$, where u_i is the i th block of d_1 components of u .

Given $u \in \mathcal{U}_N$, $u = (u_1, \dots, u_N)$, we consider the controlled system of SDEs

$$(2.3) \quad \begin{aligned} d\bar{X}^{i,N}(t) &= b(\bar{X}^{i,N}(t), \bar{\mu}^N(t)) dt + \sigma(\bar{X}^{i,N}(t), \bar{\mu}^N(t)) u_i(t) dt \\ &\quad + \sigma(\bar{X}^{i,N}(t), \bar{\mu}^N(t)) dW^i(t), \quad \bar{X}^{i,N}(0) = x^{i,N}, \end{aligned}$$

where $\bar{\mu}^N(t)$ and $\bar{\mu}^N$ are the empirical measures of $\bar{X}^{i,N}(t)$ and $\bar{X}^{i,N}$, respectively,

$$\bar{\mu}^N(t, \omega) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}^{i,N}(t, \omega)}, \quad \bar{\mu}_\omega^N \doteq \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}^{i,N}(\cdot, \omega)}, \quad \omega \in \Omega.$$

The ‘‘barred’’ symbols in the display above and in (2.3) refer to objects depending on a control, here u . We adopt this as a convention and indicate control-dependent objects by overbars. The existence and uniqueness of strong solutions to (2.3) will be a consequence of assumption (A3) made in Section 3; see comments below assumption (A5) there.

It will be convenient to have a path space which is Polish for the components u_i , $i \in \{1, \dots, N\}$, of a control process $u \in \mathcal{U}_N$. We choose the space of deterministic relaxed controls on $\mathbb{R}^{d_1} \times [0, T]$ with finite first moments. Let us first recall some facts about deterministic relaxed controls (see, e.g., [23], Section 3.2, for the case of a compact space of control actions). Denote by \mathcal{R} the space of all deterministic relaxed controls on $\mathbb{R}^{d_1} \times [0, T]$, that is, \mathcal{R} is the set of all positive measures r on $\mathcal{B}(\mathbb{R}^{d_1} \times [0, T])$ such that $r(\mathbb{R}^{d_1} \times [0, t]) = t$ for all $t \in [0, T]$. If $r \in \mathcal{R}$ and $B \in \mathcal{B}(\mathbb{R}^{d_1})$, then the mapping $[0, T] \ni t \mapsto r(B \times [0, t])$ is absolutely continuous, hence differentiable almost everywhere. Since $\mathcal{B}(\mathbb{R}^{d_1})$ is countably generated, the time derivative of r exists almost everywhere and is a measurable mapping $r_t : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^{d_1})$ such that $r(dy \times dt) = r_t(dy) dt$.

Denote by \mathcal{R}_1 the space of deterministic relaxed controls with finite first moments, that is,

$$\mathcal{R}_1 \doteq \left\{ r \in \mathcal{R} : \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r(dy \times dt) < \infty \right\}.$$

By definition, $\mathcal{R}_1 \subset \mathcal{R}$. The topology of weak convergence of measures turns \mathcal{R} into a Polish space (not compact in our case). We equip \mathcal{R}_1 with the topology of weak convergence of measures plus convergence of first moments. This topology

turns \mathcal{R}_1 into a Polish space (cf. [28], Section 6.3). It is related to the Monge–Kantorovich distances. For $T = 1$ (else one has to renormalize), the topology coincides with that induced by the Monge–Kantorovich distance with exponent one, also called the Kantorovich–Rubinstein distance or Wasserstein distance of order one. The topology is convenient because the controls appear in an unbounded (but affine) fashion in the dynamics. Thus ordinary weak convergence will not imply convergence of corresponding integrals, but convergence in \mathcal{R}_1 will.

Any \mathbb{R}^{d_1} -valued process v defined on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}})$ induces an \mathcal{R} -valued random variable ρ according to

$$(2.4) \quad \rho_\omega(B \times I) \doteq \int_I \delta_{v(t, \omega)}(B) dt, \quad B \in \mathcal{B}(\mathbb{R}^{d_1}), I \subset [0, T], \omega \in \tilde{\Omega}.$$

If v is such that $\int_0^T |v(t, \omega)| dt < \infty$ for all $\omega \in \tilde{\Omega}$, then the induced random variable ρ takes values in \mathcal{R}_1 . If v is progressively measurable with respect to a filtration $(\tilde{\mathcal{F}}_t)$ in $\tilde{\mathcal{F}}$, then ρ is adapted in the sense that the mapping $t \mapsto \rho(B \times [0, t])$ is $(\tilde{\mathcal{F}}_t)$ -adapted for all $B \in \mathcal{B}(\mathbb{R}^{d_1})$ [23], Section 3.3.

Given an adapted (in the above sense) \mathcal{R}_1 -valued random variable ρ and a Borel measurable mapping $v : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$, we will consider the controlled SDE

$$(2.5) \quad \begin{aligned} d\bar{X}(t) &= b(\bar{X}(t), v(t)) dt + \left(\int_{\mathbb{R}^{d_1}} \sigma(\bar{X}(t), v(t)) y \rho_t(dy) \right) dt \\ &+ \sigma(\bar{X}(t), v(t)) dW(t), \quad \bar{X}(0) \sim \nu(0), \end{aligned}$$

where W is a d_1 -dimensional $(\tilde{\mathcal{F}}_t)$ -adapted standard Wiener process. Equation (2.5) is a parameterized version of (2.7) below, the controlled analog of the limit SDE (2.2). We will only have to deal with weak solutions of (2.5) or, equivalently, with certain probability measures on $\mathcal{B}(\mathcal{Z})$, where

$$\mathcal{Z} \doteq \mathcal{X} \times \mathcal{R}_1 \times \mathcal{W}.$$

For a typical element in \mathcal{Z} let us write (φ, r, w) with the understanding that $\varphi \in \mathcal{X}$, $r \in \mathcal{R}_1$, $w \in \mathcal{W}$.

Notice that we include \mathcal{W} as a component of our canonical space \mathcal{Z} . This will allow identification of the joint distribution of the control and driving Wiener process. Indeed, if the triple (\bar{X}, ρ, W) defined on some filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}}, (\tilde{\mathcal{F}}_t))$ solves (2.5) for some measurable $v : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$, then the distribution of (\bar{X}, ρ, W) under $\tilde{\mathbf{P}}$ is an element of $\mathcal{P}(\mathcal{Z})$.

When (2.5) is used the mapping $v : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$ appearing in the coefficients will be determined by a probability measure on $\mathcal{B}(\mathcal{Z})$. To be more precise, let $\Theta \in \mathcal{P}(\mathcal{Z})$. Then Θ induces a mapping $\nu_\Theta : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$ which is defined by

$$(2.6) \quad \nu_\Theta(t)(B) \doteq \Theta(\{(\varphi, r, w) \in \mathcal{Z} : \varphi(t) \in B\}), \quad B \in \mathcal{B}(\mathbb{R}^d), t \in [0, T].$$

By construction, $\nu_\Theta(t)$ is the distribution under Θ of the first component of the coordinate process on $\mathcal{Z} = \mathcal{X} \times \mathcal{R}_1 \times \mathcal{W}$ at time t . Therefore, if Θ corresponds to a weak solution of (2.5) with $\nu = \nu_\Theta$, then Θ also corresponds to a weak solution of the controlled limit SDE

$$(2.7) \quad \begin{aligned} d\bar{X}(t) &= b(\bar{X}(t), \text{Law}(\bar{X}(t))) dt + \left(\int_{\mathbb{R}^{d_1}} \sigma(\bar{X}(t), \text{Law}(\bar{X}(t))) y \rho_t(dy) \right) dt \\ &+ \sigma(\bar{X}(t), \text{Law}(\bar{X}(t))) dW(t), \quad \bar{X}(0) \sim \nu_\Theta(0). \end{aligned}$$

Here W is a d_1 -dimensional standard Wiener process defined on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}})$ carrying a filtration $(\tilde{\mathcal{F}}_t)$, and ρ is an $(\tilde{\mathcal{F}}_t)$ -adapted \mathcal{R}_1 -valued random variable such that (\bar{X}, ρ, W) has distribution Θ under $\tilde{\mathbf{P}}$. The process triple (\bar{X}, ρ, W) can be given explicitly as the coordinate process on the probability space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}), \Theta)$ endowed with the canonical filtration (\mathcal{G}_t) in $\mathcal{B}(\mathcal{Z})$. More precisely, the processes \bar{X}, ρ, W are defined on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ by

$$\begin{aligned} \bar{X}(t, (\varphi, r, w)) &\doteq \varphi(t), & \rho(t, (\varphi, r, w)) &\doteq r|_{\mathcal{B}(\mathbb{R}^{d_1} \times [0, t])}, \\ W(t, (\varphi, r, w)) &\doteq w(t). \end{aligned}$$

Here we abuse notation and use $\rho(t, \cdot)$ to denote the restriction of a measure defined on $\mathcal{B}(\mathbb{R}^{d_1} \times [0, T])$ to $\mathcal{B}(\mathbb{R}^{d_1} \times [0, t])$. The canonical filtration is given by

$$\mathcal{G}_t \doteq \sigma((\bar{X}(s), \rho(s), W(s)) : 0 \leq s \leq t), \quad t \in [0, T].$$

Notice that $\rho(s)$ takes values in the space of deterministic relaxed controls on $\mathbb{R}^{d_1} \times [0, s]$ with finite first moments.

One of the assumptions we make below [assumption (A4) in Section 3] is the weak uniqueness of solutions to (2.7). If $((\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}}), (\tilde{\mathcal{F}}_t), (\bar{X}, \rho, W))$ is a weak solution of (2.7), then $\tilde{\mathbf{P}} \circ (\bar{X}, \rho, W)^{-1} \in \mathcal{P}(\mathcal{Z})$. The property of weak uniqueness can therefore be formulated in terms of probability measures on $\mathcal{B}(\mathcal{Z})$.

DEFINITION 1. *Weak uniqueness* is said to hold for (2.7) if whenever $\Theta, \tilde{\Theta} \in \mathcal{P}(\mathcal{Z})$ are such that $\Theta, \tilde{\Theta}$ both correspond to weak solutions of (2.7), $\nu_\Theta(0) = \nu_{\tilde{\Theta}}(0)$ and $\Theta|_{\mathcal{B}(\mathcal{R}_1 \times \mathcal{W})} = \tilde{\Theta}|_{\mathcal{B}(\mathcal{R}_1 \times \mathcal{W})}$, then $\Theta = \tilde{\Theta}$.

Thus, weak uniqueness for (2.7) means that, given any initial distribution for the state process, the joint distribution of control and driving Wiener process uniquely determines the distribution of the solution triple.

3. Laplace principle. A function $I : \mathcal{P}(\mathcal{X}) \rightarrow [0, \infty]$ is called a *rate function* if for each $M < \infty$ the set $\{\theta \in \mathcal{P}(\mathcal{X}) : I(\theta) \leq M\}$ is compact (some authors call such functions *good rate functions*). We say that a Laplace principle holds for the family $\{\mu^N, N \in \mathbb{N}\}$ with rate function I if for any bounded and continuous function $F : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$,

$$(3.1) \quad \lim_{N \rightarrow \infty} -\frac{1}{N} \log \mathbf{E}[\exp(-N \cdot F(\mu^N))] = \inf_{\theta \in \mathcal{P}(\mathcal{X})} \{F(\theta) + I(\theta)\}.$$

It is well known that in our setting the Laplace principle holds if and only if $\{\mu^N, N \in \mathbb{N}\}$ satisfies a large deviation principle with rate function I [16], Section 1.2.

Let us make the following assumptions about the functions b, σ and the family $\{x^{i,N}\} \subset \mathbb{R}^d$ of initial conditions:

(A1) For some $\nu_0 \in \mathcal{P}(\mathbb{R}^d)$, $\frac{1}{N} \sum_{i=1}^N \delta_{x^{i,N}} \rightarrow \nu_0$ as N tends to infinity.

(A2) The coefficients b, σ are continuous.

(A3) For all $N \in \mathbb{N}$, existence and uniqueness of solutions holds in the strong sense for the system of N equations given by (2.1).

(A4) Weak uniqueness of solutions holds for (2.7).

(A5) If $u^N \in \mathcal{U}_N, N \in \mathbb{N}$, are such that

$$\sup_{N \in \mathbb{N}} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] < \infty,$$

then $\{\bar{\mu}^N, N \in \mathbb{N}\}$ is tight as a family of $\mathcal{P}(\mathcal{X})$ -valued random variables, where $\bar{\mu}^N$ is the empirical measure of the solution to the system of (2.3) under u^N .

Assumption (A1) is a sort of law of large numbers for the deterministic initial conditions. The assumption is necessary for the convergence of the empirical measures μ^N associated with the state process. The continuity assumption (A2) implies that the coefficients b, σ are uniformly continuous and uniformly bounded on sets $B \times P$, where $B \subset \mathbb{R}^d$ is bounded and $P \subset \mathcal{P}(\mathbb{R}^d)$ is compact.

Assumption (A3) about strong existence and uniqueness of solutions for the prelimit model will be needed to justify a variational representation for the cumulant generating functionals appearing in (3.1); see (3.3) below. Assumption (A3) and an application of Girsanov's theorem show that (2.3) has a unique strong solution whenever $\int_0^T |u(t)|^2 dt \leq M$ \mathbf{P} -almost surely for some $M \in (0, \infty)$. In fact, there is a Borel measurable mapping $h^N = (h_1^N, \dots, h_N^N)$ with $h_i^N : \Omega \rightarrow \mathcal{X}, i \in \{1, \dots, N\}$, such that, for \mathbf{P} -almost all $\omega \in \Omega$, the unique strong solution of (2.1) is given as

$$X^{i,N}(\cdot, \omega) = h_i^N(W(\cdot, \omega)),$$

and under the above integrability condition on u , the unique strong solution of (2.3) equals \mathbf{P} -almost surely

$$\bar{X}^{i,N}(\cdot, \omega) = h_i^N \left(W(\cdot, \omega) + \int_0^\cdot u(s, \omega) ds \right).$$

By a localization argument one can now show that (2.3) in fact has a unique strong solution for all $u \in \mathcal{U}_N$, which is once more given by the above relation.

Weak uniqueness as stipulated in (A4) for the controlled nonlinear diffusions given by (2.7) is meant in the sense of Definition 1. It is typical that such weak uniqueness holds if it holds for the uncontrolled system (2.2).

Grant assumption (A1). Then assumptions (A2)–(A5) are all satisfied if b, σ are uniformly Lipschitz [with respect to the bounded Lipschitz metric on $\mathcal{P}(\mathbb{R}^d)$] or locally Lipschitz satisfying a suitable coercivity condition. A simple example of such a condition on b, σ would be that for some constant $C > 0$, all $x \in \mathbb{R}^d$ and all $v \in \mathcal{P}(\mathbb{R}^d)$,

$$2\langle b(x, v), x \rangle + \text{tr}(\sigma \sigma^\top)(x, v) \leq C(1 + |x|^2).$$

The reason for assumption (A5) being stated as it is, is that there are many different sets of conditions on the problem data (i.e., b and σ) and the initial conditions which imply tightness of the empirical measures of the $\bar{X}^{i,N}$. For instance, (A5) is automatically satisfied if the coefficients are bounded. It also holds if b, σ are Lipschitz continuous. More general conditions can be formulated in terms of the action of the infinitesimal generator associated with (2.7), given in (4.2) below, on some “Lyapunov function” $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$; also see Section 7.1.

For a probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$, recalling that $\mathcal{Z} = \mathcal{X} \times \mathcal{R}_1 \times \mathcal{W}$, let $\Theta_{\mathcal{X}}, \Theta_{\mathcal{R}}$ denote the first and second marginal, respectively. Let \mathcal{P}_∞ be the set of all probability measures $\Theta \in \mathcal{P}(\mathcal{Z})$ such that:

(i)

$$\int_{\mathcal{R}_1} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) < \infty;$$

(ii) Θ corresponds to a weak solution of (2.7);

(iii) $\nu_\Theta(0) = \nu_0$, where $\nu_0 \in \mathcal{P}(\mathbb{R}^d)$ is the initial distribution from assumption (A1).

The main result of this paper is the following.

THEOREM 3.1. *Suppose that assumptions (A1)–(A5) hold. Then the family of empirical measures $\{\mu^N, N \in \mathbb{N}\}$ satisfies the Laplace principle with rate function*

$$I(\theta) = \inf_{\Theta \in \mathcal{P}_\infty : \Theta_{\mathcal{X}} = \theta} \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr).$$

REMARK 3.2. The above expression for the rate function I is convenient for proving the Laplace principle. An alternative and perhaps more familiar form of the rate function is the following. By definition of \mathcal{P}_∞ , and since the control appears linearly in the limit dynamics, we can write

$$I(\theta) = \inf_{\Theta \in \mathcal{P}_\infty : \Theta_{\mathcal{X}} = \theta} \mathbf{E}_\Theta \left[\frac{1}{2} \int_0^T |u(t)|^2 dt \right],$$

where $\inf \emptyset \doteq \infty$ by convention, $u(t) = \int_{\mathbb{R}^{d_1}} y \rho_t(dy)$, (\bar{X}, W, ρ) is the canonical process on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$, and Θ -almost surely \bar{X} satisfies

$$(3.2) \quad d\bar{X}(t) = b(\bar{X}(t), \theta(t)) dt + \sigma(\bar{X}(t), \theta(t)) u(t) dt + \sigma(\bar{X}(t), \theta(t)) dW(t).$$

The proof of Theorem 3.1 is based on a representation for functionals of Brownian motion, a martingale characterization of weak solutions of (2.7) and weak convergence arguments.

By assumption (A3), for each $N \in \mathbb{N}$, the N -particle system of (2.1) possesses a unique strong solution for the given initial condition. By Theorem 3.6 in [6], for any $F \in \mathbf{C}_b(\mathcal{X})$ the prelimit expressions in (3.1) can be rewritten as

$$(3.3) \quad \begin{aligned} & -\frac{1}{N} \log \mathbf{E}[\exp(-N \cdot F(\mu^N))] \\ & = \inf_{u^N \in \mathcal{U}_N} \left\{ \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)] \right\}, \end{aligned}$$

where $\bar{\mu}^N$ is the empirical measure of the solution to the system of (2.3) under $u^N = (u_1^N, \dots, u_N^N) \in \mathcal{U}_N$. The representation in [6] applies to an infinite-dimensional Brownian motion, and thus strictly speaking the infimum would be over a collection of controls indexed by $i \in \mathbb{N}$. However, since those controls with $i > N$ have no effect on $\bar{\mu}^N$ we can and will assume they are zero.

Based on (3.3), the Laplace principle will be established in two steps. First, in Section 5, we establish the variational lower bound by showing that for any sequence $(u^N)_{N \in \mathbb{N}}$ with $u^N \in \mathcal{U}_N$,

$$(3.4) \quad \begin{aligned} & \liminf_{N \rightarrow \infty} \left\{ \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)] \right\} \\ & \geq \inf_{\Theta \in \mathcal{P}_\infty} \left\{ \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) + F(\Theta_{\mathcal{X}}) \right\}. \end{aligned}$$

Second, in Section 6, we verify the variational upper bound by showing that for any measure $\Theta \in \mathcal{P}_\infty$ there is a sequence $(u^N)_{N \in \mathbb{N}}$ with $u^N \in \mathcal{U}_N$ such that

$$(3.5) \quad \begin{aligned} & \limsup_{N \rightarrow \infty} \left\{ \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)] \right\} \\ & \leq \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) + F(\Theta_{\mathcal{X}}). \end{aligned}$$

To see that those two steps establish Theorem 3.1, first observe that

$$\begin{aligned} & \inf_{\theta \in \mathcal{P}(\mathcal{X})} \left\{ F(\theta) + \inf_{\Theta \in \mathcal{P}_\infty : \Theta_{\mathcal{X}} = \theta} \left\{ \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) \right\} \right\} \\ & = \inf_{\Theta \in \mathcal{P}_\infty} \left\{ \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) + F(\Theta_{\mathcal{X}}) \right\}. \end{aligned}$$

Hence, in view of (3.3), we have to show that for all $F \in \mathbf{C}_b(\mathcal{X})$,

$$\inf_{u \in \mathcal{U}_N} J_N^F(u) \xrightarrow{N \rightarrow \infty} \inf_{\Theta \in \mathcal{P}_\infty} J_\infty^F(\Theta),$$

where

$$J_N^F(u) \doteq \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)],$$

$$J_\infty^F(\Theta) \doteq \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) + F(\Theta_{\mathcal{X}}).$$

Let $\varepsilon > 0$. For the lower bound, choose $u^N \in \mathcal{U}_N$, $N \in \mathbb{N}$, such that $J_N^F(u^N) \leq \inf_{u \in \mathcal{U}_N} J_N^F(u) + \varepsilon$. Then (3.4) implies that

$$\liminf_{N \rightarrow \infty} \inf_{u \in \mathcal{U}_N} J_N^F(u) \geq \inf_{\Theta \in \mathcal{P}_\infty} J_\infty^F(\Theta) - \varepsilon.$$

For the upper bound, choose a probability measure $\Theta \in \mathcal{P}_\infty$ such that $J_\infty^F(\Theta) \leq \inf_{\Theta \in \mathcal{P}_\infty} J_\infty^F(\Theta) + \varepsilon$. Since $\inf_{u \in \mathcal{U}_N} J_N^F(u) \leq J_N^F(\tilde{u})$ for any $\tilde{u} \in \mathcal{U}_N$, (3.5) implies that

$$\limsup_{N \rightarrow \infty} \inf_{u \in \mathcal{U}_N} J_N^F(u) \leq \inf_{\Theta \in \mathcal{P}_\infty} J_\infty^F(\Theta) + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, the assertion follows.

There is a technical observation to be made about the probability spaces and filtrations underlying the stochastic control problems, namely that there is a certain flexibility in the choice of the stochastic bases. This flexibility will be needed in establishing the variational upper bound. To be more precise we note that the representation theorem in [6] holds for any stochastic basis rich enough to carry a sequence of independent standard $(\tilde{\mathcal{F}}_t)$ -Wiener processes. The filtration $(\tilde{\mathcal{F}}_t)$, which is assumed to satisfy the usual conditions, need not be the filtration induced by the Wiener processes, but may be strictly larger. As a consequence of assumption (A3), the left-hand side of (3.3) does not depend on the choice of the stochastic basis. The stochastic optimal control problem on the right-hand side of (3.3) can therefore be regarded in the weak sense, that is, the infimum is taken over all suitable stochastic bases (see Definition 4.2 in [34], page 64). The definition of the sets \mathcal{U}_N and assumption (A5) are to be understood accordingly.

As a consequence of the weak formulation of the control problems, in the proof of the variational lower bound, the control processes u^N , the driving Wiener processes W^1, \dots, W^N and thus the empirical measures $\bar{\mu}^N$ could live on stochastic bases which vary with N . While we do not make this variation explicit, it is easy to see that the arguments of Section 5, being weak convergence arguments, do not rely on having a common filtered probability space. The variational upper bound, on the other hand, will be established in Section 6 by taking an arbitrary $\Theta \in \mathcal{P}_\infty$ and then constructing a sequence of control processes and independent Wiener processes so that (3.5) holds. The prelimit processes will be coordinate processes on a common stochastic basis which, however, will depend on the limit probability measure Θ .

4. Auxiliary constructions. This section collects useful results for characterizing those probability measures in $\mathcal{P}(\mathcal{Z})$ which correspond to a weak solution of (2.7). Let $\Theta \in \mathcal{P}(\mathcal{Z})$. Recall from (2.6) the definition of the mapping $\nu_\Theta : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$ induced by Θ . The mapping ν_Θ is continuous. To check this, take any $t_0 \in [0, T]$ and any sequence $(t_n) \subset [0, T]$ such that $t_n \rightarrow t_0$. Then for all $f \in \mathbf{C}_b(\mathbb{R}^d)$, the fact that elements of \mathcal{X} are continuous and the bounded convergence theorem imply

$$\begin{aligned} \int_{\mathbb{R}^d} f(x) \nu_\Theta(t_n)(dx) &= \int_{\mathcal{X} \times \mathcal{R} \times \mathcal{W}} f(\varphi(t_n)) \Theta(d\varphi \times dr \times dw) \\ &\xrightarrow{n \rightarrow \infty} \int_{\mathcal{X} \times \mathcal{R} \times \mathcal{W}} f(\varphi(t_0)) \Theta(d\varphi \times dr \times dw) \\ &= \int_{\mathbb{R}^d} f(x) \nu_\Theta(t_0)(dx). \end{aligned}$$

Therefore $\nu_\Theta(t_n) \rightarrow \nu_\Theta(t)$ in $\mathcal{P}(\mathbb{R}^d)$. The continuity of ν_Θ implies that the set $\{\nu_\Theta(t) : t \in [0, T]\}$ is compact in $\mathcal{P}(\mathbb{R}^d)$.

The question of whether a probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ corresponds to a weak solution of (2.7), or, equivalently, of (2.5) with $\nu = \nu_\Theta$, can be conveniently phrased in terms of an associated local martingale problem. We summarize here the main facts that we will use (see [32], [23], Section 4.4, and [22], Section 5.4, e.g.).

Given $f \in \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$, define a real-valued process $(M_f^\Theta(t))_{t \in [0, T]}$ on the probability space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}), \Theta)$ by

$$(4.1) \quad \begin{aligned} M_f^\Theta(t, (\varphi, r, w)) &\doteq f(\varphi(t), w(t)) - f(\varphi(0), 0) \\ &\quad - \int_0^t \int_{\mathbb{R}^{d_1}} \mathcal{A}_s^\Theta(f)(\varphi(s), y, w(s)) r_s(dy) ds, \end{aligned}$$

where for $s \in [0, T]$, $x \in \mathbb{R}^d$, $y, z \in \mathbb{R}^{d_1}$,

$$(4.2) \quad \begin{aligned} \mathcal{A}_s^\Theta(f)(x, y, z) &\doteq \langle b(x, \nu_\Theta(s)) + \sigma(x, \nu_\Theta(s))y, \nabla_x f(x, z) \rangle \\ &\quad + \frac{1}{2} \sum_{j,k=1}^d (\sigma \sigma^\top)_{jk}(x, \nu_\Theta(s)) \frac{\partial^2 f}{\partial x_j \partial x_k}(x, z) \\ &\quad + \frac{1}{2} \sum_{l=1}^{d_1} \frac{\partial^2 f}{\partial z_l \partial z_l}(x, z) \\ &\quad + \sum_{k=1}^d \sum_{l=1}^{d_1} \sigma_{kl}(x, \nu_\Theta(s)) \frac{\partial^2 f}{\partial x_k \partial z_l}(x, z). \end{aligned}$$

The expression involving $\mathcal{A}_s^\Theta(f)$ in (4.1) is integrated against time and the time derivative measures r_s of any relaxed control r . The measures r_s are actually not needed in that we may use $r(dy \times ds)$ in place of $r_s(dy) ds$.

The key relation, which we formulate as a lemma, is a one-to-one correspondence between weak solutions of (2.7) and a local martingale problem.

LEMMA 4.1. *Let $\Theta \in \mathcal{P}(\mathcal{Z})$ be such that $\Theta(\{(\varphi, r, w) \in \mathcal{Z} : w(0) = 0\}) = 1$. Then Θ corresponds to a weak solution of (2.7) if and only if M_f^Θ is a local martingale under Θ with respect to the canonical filtration (\mathcal{G}_t) for all $f \in \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$.*

Moreover, in order to show that Θ corresponds to a weak solution of (2.7), it is enough to check the local martingale property for those M_f^Θ where the test function f is a monomial of first or second order, that is, for the test functions

$$\begin{aligned} (x, z) &\mapsto x_k, & k \in \{1, \dots, d\}, & & (x, z) &\mapsto x_j x_k, & j, k \in \{1, \dots, d\}, \\ (x, z) &\mapsto z_l, & l \in \{1, \dots, d_1\}, & & (x, z) &\mapsto z_j z_l, & j, l \in \{1, \dots, d_1\}, \\ (x, z) &\mapsto x_k z_l, & k \in \{1, \dots, d\}, l \in \{1, \dots, d_1\}. \end{aligned}$$

PROOF. See, for example, the proof of Proposition 5.4.6 in [22], page 315. Note that since the canonical process on the sample space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ includes a component which corresponds to the driving Wiener process, there is no need to extend the probability space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}), \Theta)$ even if the diffusion coefficient σ is degenerate. \square

REMARK 4.2. There is a technical point here concerning the canonical filtration (\mathcal{G}_t) in $\mathcal{B}(\mathcal{Z})$. That filtration is not necessarily Θ -complete or right-continuous, while in the literature solutions to SDEs are usually defined with respect to filtrations satisfying the usual conditions (i.e., containing all sets contained in a set of measure zero and being right-continuous). However, any stochastically continuous and uniformly bounded real-valued process defined on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}})$ which is a martingale under $\tilde{\mathbf{P}}$ with respect to some filtration $(\tilde{\mathcal{F}}_t)$, is also a martingale under $\tilde{\mathbf{P}}$ with respect to $(\tilde{\mathcal{F}}_{t+}^{\tilde{\mathbf{P}}})$, where $(\tilde{\mathcal{F}}_t^{\tilde{\mathbf{P}}})$ denotes the $\tilde{\mathbf{P}}$ -augmentation of $(\tilde{\mathcal{F}}_t)$ (see the solution to Exercise 5.4.13 in [22], page 392). The filtration $(\tilde{\mathcal{F}}_{t+}^{\tilde{\mathbf{P}}})$ satisfies the usual conditions. Since the localizing sequence of stopping times for a local martingale can always be chosen in such a way that the corresponding stopped processes are bounded martingales, it follows that if M_f^Θ is a local martingale under Θ with respect to (\mathcal{G}_t) , then it is also a local martingale under Θ with respect to $(\tilde{\mathcal{G}}_{t+}^\Theta)$. The local martingale property of the processes M_f^Θ under Θ with respect to the canonical filtration (\mathcal{G}_t) thus implies that the canonical process on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ solves (2.7) under Θ with respect to the filtration $(\tilde{\mathcal{G}}_{t+}^\Theta)$, which satisfies the usual conditions.

REMARK 4.3. The reason why we use a local martingale problem rather than the corresponding martingale problem is that it gives more flexibility in characterizing the convergence of Itô processes which are not necessarily of diffusion

type. In Section 7.2, we extend the Laplace principle of Theorem 3.1 to interacting systems described by SDEs with delay. In that case, the coefficients b, σ are progressive functionals; thus, they may depend on the entire trajectory of the solution process up to the current time. An appropriate choice of the stopping times in the local martingale problem gives control over the state process up to the current time and not only at the current time. In particular, the proof of Lemma 5.2 below, where the local martingale problem is used to identify certain limit distributions, continues to work also for the more general model of Section 7.2.

5. Variational lower bound. In the proof of the lower bound (3.4) we can assume that

$$(5.1) \quad \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] \leq 2\|F\|,$$

since otherwise the desired inequality is automatic. Let $(u^N)_{N \in \mathbb{N}}$ be a sequence of control processes such that (5.1) holds. This implies in particular that for \mathbf{P} -almost all $\omega \in \Omega$, all $N \in \mathbb{N}$, $i \in \{1, \dots, N\}$, $\int_0^T |u_i^N(t, \omega)| dt < \infty$. Modifying the sequence (u^N) on a set of \mathbf{P} -measure zero has no impact on the validity of (3.4). Thus, we may assume that $u_i^N(\cdot, \omega)$ has a finite first moment for all $\omega \in \Omega$.

For each $N \in \mathbb{N}$, define a $\mathcal{P}(\mathcal{Z})$ -valued random variable by

$$(5.2) \quad Q_\omega^N(B \times R \times D) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}^{i,N}(\cdot, \omega)}(B) \cdot \delta_{\rho_\omega^{i,N}}(R) \cdot \delta_{W^{i,N}(\cdot, \omega)}(D),$$

$B \times R \times D \in \mathcal{B}(\mathcal{Z})$, $\omega \in \Omega$, where $\bar{X}^{i,N}$ is the solution of (2.3) under $u^N = (u_1^N, \dots, u_N^N)$, and $\rho_\omega^{i,N}$ is the relaxed control induced by $u_i^N(\cdot, \omega)$ according to (2.4). Notice that $\rho_\omega^{i,N} \in \mathcal{R}_1$. The functional occupation measures Q^N , $N \in \mathbb{N}$, just defined are related to the Laplace principle by the fact that

$$(5.3) \quad \begin{aligned} & \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)] \\ &= \int_\Omega \left[\int_{\mathcal{R}_1} \left(\frac{1}{2} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \right) Q_{\omega, \mathcal{R}}^N(dr) \right. \\ & \quad \left. + F(Q_{\omega, \mathcal{X}}^N) \right] \mathbf{P}(d\omega), \end{aligned}$$

where $Q_{\omega, \mathcal{X}}^N$, $Q_{\omega, \mathcal{R}}^N$ denote the first and second marginal of $Q_\omega^N \in \mathcal{P}(\mathcal{Z})$, respectively, and we recall that $\mathcal{Z} = \mathcal{X} \times \mathcal{R}_1 \times \mathcal{W}$.

Thanks to assumption (A5) and the bound (5.1), the first marginals of $(Q^N)_{N \in \mathbb{N}}$ are tight as random measures. The next lemma states that tightness of $(Q^N)_{N \in \mathbb{N}}$ as random measures follows. Thus we are asserting tightness of the measures $\gamma^N \in \mathcal{P}(\mathcal{P}(\mathcal{Z}))$ defined by $\gamma^N(A) = \mathbf{P}(Q^N \in A)$, $A \in \mathcal{B}(\mathcal{P}(\mathcal{Z}))$.

LEMMA 5.1. *The family $(Q^N)_{N \in \mathbb{N}}$ of $\mathcal{P}(\mathcal{Z})$ -valued random variables is tight.*

PROOF. The first marginals of $(Q^N)_{N \in \mathbb{N}}$ are tight by assumption (A5) and (5.1). Since the third marginals are obviously tight, we need only prove tightness of the second marginals. Observe that

$$g(r) \doteq \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt)$$

is a tightness function on \mathcal{R}_1 , that is, it is bounded from below and has compact level sets. To verify the last property take $c \in (0, \infty)$ and let $R_c \doteq \{r \in \mathcal{R}_1 : g(r) \leq c\}$. By Chebyshev's inequality, for all $M > 0$,

$$(*) \quad \sup_{r \in R_c} r(\{y \in \mathbb{R}^{d_1} : |y| > M\} \times [0, T]) \leq \frac{c}{M^2}.$$

Hence R_c is tight and thus relatively compact as a subset of \mathcal{R} . Consequently, any sequence in R_c has a weakly convergent subsequence with limit in \mathcal{R} . Let $(r_n) \subset R_c$ be such that (r_n) converges weakly to r_* for some $r_* \in \mathcal{R}$. It remains to show that r_* has finite first moment and that the first moments of (r_n) converge to that of r_* . By Hölder's inequality and a version of Fatou's lemma (cf. Theorem A.3.12 in [16], page 307),

$$\sqrt{T \cdot c} \geq \liminf_{n \rightarrow \infty} \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_n(dy \times dt) \geq \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_*(dy \times dt).$$

Let $M > 0$. By (*) and Hölder's inequality we have for all $r \in R_c$,

$$\int_{\{y \in \mathbb{R}^{d_1} : |y| > M\} \times [0, T]} |y| r(dy \times dt) \leq \frac{c}{M}.$$

Therefore, using weak convergence,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_n(dy \times dt) &\leq \frac{c}{M} + \int_{\{y \in \mathbb{R}^{d_1} : |y| \leq M\} \times [0, T]} |y| r_*(dy \times dt) \\ &\leq \frac{c}{M} + \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_*(dy \times dt). \end{aligned}$$

Since $M > 0$ may be arbitrarily big, it follows that

$$\lim_{n \rightarrow \infty} \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_n(dy \times dt) = \int_{\mathbb{R}^{d_1} \times [0, T]} |y| r_*(dy \times dt).$$

We conclude that g is a tightness function on \mathcal{R}_1 . Now define a function $G : \mathcal{P}(\mathcal{Z}) \rightarrow [0, \infty]$ by

$$G(\Theta) \doteq \int_{\mathcal{Z}} g(r) \Theta(d\varphi \times dr \times dw).$$

Then G is a tightness function on second marginals in $\mathcal{P}(\mathcal{Z})$ (see Theorem A.3.17 in [16], page 309). Thus in order to prove tightness of the second marginals of $(Q^N)_{N \in \mathbb{N}}$ (as random measures) it is enough to show that

$$\sup_{N \in \mathbb{N}} \mathbf{E}[G(Q^N)] < \infty.$$

However, this follows directly from (5.1). \square

In the next lemma we identify the limit points of (Q^N) as being weak solutions of (2.7) with probability one. The proof is similar in spirit to that of Theorem 5.3.1 in [23], page 102.

LEMMA 5.2. *Let $(Q^{N_j})_{j \in \mathbb{N}}$ be a weakly convergent subsequence of $(Q^N)_{N \in \mathbb{N}}$. Let Q be a $\mathcal{P}(\mathcal{Z})$ -valued random variable defined on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}})$ such that $Q^{N_j} \xrightarrow{j \rightarrow \infty} Q$ in distribution. Then Q_ω corresponds to a weak solution of (2.7) for $\tilde{\mathbf{P}}$ -almost all $\omega \in \tilde{\Omega}$.*

PROOF. Set $I \doteq \{N_j, j \in \mathbb{N}\}$, and write $(Q^n)_{n \in I}$ for $(Q^{N_j})_{j \in \mathbb{N}}$. By hypothesis, $Q^n \rightarrow Q$ in distribution.

Recall from Lemma 4.1 in Section 4 that a probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ with $\Theta(\{(\varphi, r, w) \in \mathcal{Z} : w(0) = 0\}) = 1$ corresponds to a weak solution of (2.7) if (and only if), for all $f \in \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$, M_f^Θ is a local martingale under Θ with respect to the canonical filtration (\mathcal{G}_t) , where M_f^Θ is defined by (4.1). Moreover, the local martingale property has to be checked only for those M_f^Θ where the test function f is a monomial of first or second order.

In verifying the local martingale property of M_f^Θ when $\Theta = Q_\omega$ for some $\omega \in \tilde{\Omega}$, we will work with randomized stopping times. Those stopping times live on an extension $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$ of the measurable space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ and are adapted to a filtration $(\hat{\mathcal{G}}_t)$ in $\mathcal{B}(\hat{\mathcal{Z}})$, where

$$\hat{\mathcal{Z}} \doteq \mathcal{Z} \times [0, 1], \quad \hat{\mathcal{G}}_t \doteq \mathcal{G}_t \times \mathcal{B}([0, 1]), \quad t \in [0, T],$$

and (\mathcal{G}_t) is the canonical filtration in $\mathcal{B}(\mathcal{Z})$. Any random object defined on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ also lives on $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$, and no notational distinction will be made.

Let λ denote the uniform distribution on $\mathcal{B}([0, 1])$. Any probability measure Θ on $\mathcal{B}(\mathcal{Z})$ induces a probability measure on $\mathcal{B}(\hat{\mathcal{Z}})$ given by $\hat{\Theta} \doteq \Theta \times \lambda$. For each $k \in \mathbb{N}$, define a stopping time τ_k on $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$ with respect to the filtration $(\hat{\mathcal{G}}_t)$ by setting, for $(z, a) \in \mathcal{Z} \times [0, 1]$,

$$\tau_k(z, a) \doteq \inf\{t \in [0, T] : v(z, t) \geq k + a\},$$

where

$$v((\varphi, r, w), t) \doteq \int_{\mathbb{R}^{d_1} \times [0, t]} |y| r(dy \times ds) + \sup_{s \in [0, t]} |\varphi(s)| + \sup_{s \in [0, t]} |w(s)|.$$

Note that the mapping $t \mapsto v((\varphi, r, w), t)$ is monotonic for all $(\varphi, r, w) \in \mathcal{Z}$. Hence the stopping times have the following properties. The boundedness of φ and w (being continuous functions on a compact interval) and the boundedness of $\int_{\mathbb{R}^{d_1} \times [0, T]} |y|r(dy \times ds)$ imply that $\tau_k \nearrow T$ as $k \rightarrow \infty$ with probability one under $\hat{\Theta}$. The second property of note is that the mapping

$$\mathcal{Z} \times [0, 1] \ni (z, a) \mapsto \tau_k(z, a) \in [0, T]$$

is continuous with probability one under $\hat{\Theta}$. To see this, note that for every $z \in \mathcal{Z}$ the set

$$A_z \doteq \{c \in \mathbb{R}_+ : v(z, s) = c \text{ for all } s \in [t, t + \delta], \text{ some } t \in [0, T], \text{ some } \delta > 0\}$$

is at most countable. However, $\hat{z} \mapsto \tau_k(\hat{z})$ fails to be continuous at (z, a) only when $k + a \in A_z$. Therefore, by Fubini's theorem,

$$\begin{aligned} & \hat{\Theta}(\{(z, a) \in \hat{\mathcal{Z}} : \tau_k \text{ discontinuous at } (z, a)\}) \\ &= \int_{\hat{\mathcal{Z}}} \mathbf{1}_{A_z}(k + a) \hat{\Theta}(dz \times da) \\ &= \int_{\mathcal{Z}} \int_{[0, 1]} \mathbf{1}_{A_z}(k + a) \lambda(da) \Theta(dz) \\ &= 0. \end{aligned}$$

Notice that if M_f^\ominus is a local martingale with respect to $(\hat{\mathcal{G}}_t)$ under $\hat{\Theta} = \Theta \times \lambda$ with localizing sequence of stopping times $(\tau_k)_{k \in \mathbb{N}}$, then M_f^\ominus is also a local martingale with respect to (\mathcal{G}_t) under Θ with localizing sequence of stopping times $(\tau_k(\cdot, 0))_{k \in \mathbb{N}}$; see [Appendix](#). Thus it suffices to prove the martingale property of M_f^\ominus up till time τ_k with respect to filtration $(\hat{\mathcal{G}}_t)$ and probability measure $\hat{\Theta}$.

Clearly, the process $M_f^\ominus(\cdot \wedge \tau_k)$ is a $(\hat{\mathcal{G}}_t)$ -martingale under $\hat{\Theta}$ if and only if

$$(5.4) \quad \mathbf{E}_{\Theta \times \lambda}[\Psi \cdot (M_f^\ominus(t_1 \wedge \tau_k) - M_f^\ominus(t_0 \wedge \tau_k))] = 0$$

for all $t_0, t_1 \in [0, T]$ with $t_0 \leq t_1$, and $\hat{\mathcal{G}}_{t_0}$ -measurable $\Psi \in \mathbf{C}_b(\hat{\mathcal{Z}})$.

To verify the martingale property of $M_f^\ominus(\cdot \wedge \tau_k)$ it is enough to check that (5.4) holds for any countable collection of times t_0, t_1 which is dense in $[0, T]$ and any countable collection of functions $\Psi \in \mathbf{C}_b(\hat{\mathcal{Z}})$ that generates the (countably many) σ -algebras $\hat{\mathcal{G}}_{t_0}$. Recall that the collection of test functions f for which a martingale property must be verified consists of just monomials of degree one or two, and hence is finite. Thus, there is a countable collection $\mathcal{T} \subset \mathbb{N} \times [0, T]^2 \times \mathbf{C}_b(\hat{\mathcal{Z}}) \times \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$ of test parameters such that if (5.4) holds for all $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$, then Θ corresponds to a weak solution of (2.7).

Let $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$. Define a mapping $\Phi = \Phi_{(k, t_0, t_1, \Psi, f)}$ by

$$\mathcal{P}(\mathcal{Z}) \ni \Theta \mapsto \Phi(\Theta) \doteq \mathbf{E}_{\Theta \times \lambda}[\Psi \cdot (M_f^\Theta(t_1 \wedge \tau_k) - M_f^\Theta(t_0 \wedge \tau_k))].$$

We claim that the mapping Φ is continuous in the topology of weak convergence on $\mathcal{P}(\mathcal{Z})$. To check this, take $\Theta \in \mathcal{P}(\mathcal{Z})$ and any sequence $(\Theta_l)_{l \in \mathbb{N}} \subset \mathcal{P}(\mathcal{Z})$ that converges to Θ . Recall the definitions (4.1) and (4.2). As a consequence of assumption (A2) and by construction of the stopping time τ_k , the integrand in (5.4) is bounded; thanks to assumption (A2) and the almost sure continuity of τ_k , it is continuous with probability one under $\hat{\Theta} \doteq \Theta \times \lambda$. By weak convergence and the mapping theorem [4], page 21, it follows that

$$(5.5) \quad \begin{aligned} & \mathbf{E}_{\Theta_l \times \lambda}[\Psi \cdot (M_f^{\Theta_l}(t_1 \wedge \tau_k) - M_f^{\Theta_l}(t_0 \wedge \tau_k))] \\ & \xrightarrow{l \rightarrow \infty} \mathbf{E}_{\Theta \times \lambda}[\Psi \cdot (M_f^\Theta(t_1 \wedge \tau_k) - M_f^\Theta(t_0 \wedge \tau_k))]. \end{aligned}$$

Since the sequence (Θ_l) converges to Θ , the set $\{\Theta_l : l \in \mathbb{N}\} \cup \{\Theta\}$ is compact in $\mathcal{P}(\mathcal{Z})$. Recalling (2.6), we find that the set of probability measures $\{\nu_{\Theta_l}(t) : l \in \mathbb{N}, t \in [0, T]\} \cup \{\nu_\Theta(t) : t \in [0, T]\}$ has compact closure in $\mathcal{P}(\mathbb{R}^d)$. We claim that together with assumption (A2) and the construction of τ_k , this implies that

$$\sup_{t \in [0, T], \hat{z} \in \hat{\mathcal{Z}}} |M_f^{\Theta_l}(t \wedge \tau_k(\hat{z}), \hat{z}) - M_f^\Theta(t \wedge \tau_k(\hat{z}), \hat{z})| \xrightarrow{l \rightarrow \infty} 0.$$

To see this, we consider, for example, the integral corresponding to the first term in the drift, which is

$$\int_0^{t \wedge \tau_k(\hat{z})} \langle b(\varphi(s), \nu_{\Theta_l}(s)), \nabla_x f(\varphi(s), w(s)) \rangle ds.$$

By the assumed continuity properties of b this converges uniformly in $t \in [0, T], \hat{z} \in \hat{\mathcal{Z}}$ to

$$\int_0^{t \wedge \tau_k(\hat{z})} \langle b(\varphi(s), \nu_\Theta(s)), \nabla_x f(\varphi(s), w(s)) \rangle ds,$$

and a similar result holds for each of the other terms. Since Ψ is bounded, it follows that

$$\begin{aligned} & |\mathbf{E}_{\Theta_l \times \lambda}[\Psi \cdot (M_f^{\Theta_l}(t_1 \wedge \tau_k) - M_f^{\Theta_l}(t_0 \wedge \tau_k))] \\ & - \mathbf{E}_{\Theta \times \lambda}[\Psi \cdot (M_f^\Theta(t_1 \wedge \tau_k) - M_f^\Theta(t_0 \wedge \tau_k))]| \xrightarrow{l \rightarrow \infty} 0. \end{aligned}$$

In combination with (5.5) this implies $\Phi(\Theta_l) \rightarrow \Phi(\Theta)$.

By hypothesis, the sequence $(Q^n)_{n \in I}$ of $\mathcal{P}(\mathcal{Z})$ -valued random variables converges to Q in distribution. Hence the mapping theorem and the continuity of Φ imply that $\Phi(Q^n) \rightarrow \Phi(Q)$ in distribution.

Let $n \in I$. By construction of Q^n and Fubini's theorem, for $\omega \in \Omega$,

$$\begin{aligned} \Phi(Q_\omega^n) &= \mathbf{E}_{Q_\omega^n \times \lambda}[\Psi \cdot (M_f^{Q_\omega^n}(t_1 \wedge \tau_k) - M_f^{Q_\omega^n}(t_0 \wedge \tau_k))] \\ &= \frac{1}{n} \sum_{i=1}^n \int_0^1 \Psi((\bar{X}^{i,n}(\cdot, \omega), \rho_\omega^{i,n}, W^i(\cdot, \omega)), a) \\ &\quad \times \left(f(\bar{X}^{i,n}(t_1 \wedge \bar{\tau}_k^{i,n}, \omega), W^i(t_1 \wedge \bar{\tau}_k^{i,n}, \omega)) \right. \\ &\quad \left. - f(\bar{X}^{i,n}(t_0 \wedge \bar{\tau}_k^{i,n}, \omega), W^i(t_0 \wedge \bar{\tau}_k^{i,n}, \omega)) \right. \\ &\quad \left. - \int_{t_0 \wedge \bar{\tau}_k^{i,n}}^{t_1 \wedge \bar{\tau}_k^{i,n}} \mathcal{A}_s^{\bar{\mu}_\omega^n}(f)(\bar{X}^{i,n}(s, \omega), u_i^n(s, \omega), \right. \\ &\quad \left. W^i(s, \omega)) ds \right) da, \end{aligned}$$

where $\mathcal{A}^{\bar{\mu}_\omega^n}$ is defined according to (4.2) with $\bar{\mu}_\omega^n$ in place of ν_Θ , and $\bar{\tau}_k^{i,n} = \bar{\tau}_k^{i,n}(\omega, a)$ is defined like $\tau_k((\varphi, r, w), a)$ with φ replaced by $\bar{X}^{i,n}(\cdot, \omega)$, r replaced by $\rho_\omega^{i,n}$, the relaxed control corresponding to $u_i^n(\cdot, \omega)$, and w replaced by $W^i(\cdot, \omega)$.

For all $a \in [0, 1]$, by Itô's formula, it holds \mathbf{P} -almost surely that

$$\begin{aligned} &f(\bar{X}^{i,n}(t_1 \wedge \bar{\tau}_k^{i,n}), W^i(t_1 \wedge \bar{\tau}_k^{i,n})) \\ &\quad - f(\bar{X}^{i,n}(t_0 \wedge \bar{\tau}_k^{i,n}), W^i(t_0 \wedge \bar{\tau}_k^{i,n})) \\ &\quad - \int_{t_0 \wedge \bar{\tau}_k^{i,n}}^{t_1 \wedge \bar{\tau}_k^{i,n}} \mathcal{A}_s^{\bar{\mu}_\omega^n}(f)(\bar{X}^{i,n}(s), u_i^n(s), W^i(s)) ds \\ &= \int_{t_0 \wedge \bar{\tau}_k^{i,n}}^{t_1 \wedge \bar{\tau}_k^{i,n}} \nabla_x f^\top(\bar{X}^{i,n}(s), W^i(s)) \sigma(X^{i,n}(s), \bar{\mu}^n(s)) dW^i(s) \\ &\quad + \int_{t_0 \wedge \bar{\tau}_k^{i,n}}^{t_1 \wedge \bar{\tau}_k^{i,n}} \nabla_z f^\top(\bar{X}^{i,n}(s), W^i(s)) dW^i(s), \end{aligned}$$

where $\bar{\tau}_k^{i,n} = \bar{\tau}_k^{i,n}(\cdot, a)$ and $\bar{\tau}_k^{i,n}, \bar{\mu}^n, \bar{X}^{i,n}, u_i^n$, are random objects on (Ω, \mathcal{F}) .

By Fubini's theorem and Jensen's inequality, we have

$$\begin{aligned} &\mathbf{E}[\Phi(Q^n)^2] \\ &\leq \int_0^1 \mathbf{E}[\mathbf{E}_{Q_\omega^n}[\Psi(\cdot, a) \cdot (M_f^{Q_\omega^n}(t_1 \wedge \tau_k(\cdot, a)) \\ &\quad - M_f^{Q_\omega^n}(t_0 \wedge \tau_k(\cdot, a)))]^2] da. \end{aligned}$$

For all $a \in [0, 1]$, by the Itô isometry and because $\Psi(\cdot, a)$ is \mathcal{G}_{t_0} -measurable, and $\tau_k(\cdot, a)$ is a stopping time with respect to (\mathcal{G}_t) , it holds that

$$\begin{aligned}
& \mathbf{E}[\mathbf{E}_{Q_\omega^n}[\Psi(\cdot, a) \cdot (M_f^{Q_\omega^n}(t_1 \wedge \tau_k(\cdot, a)) - M_f^{Q_\omega^n}(t_0 \wedge \tau_k(\cdot, a)))]^2] \\
&= \mathbf{E}[\mathbf{E}_{Q_\omega^n}[\Psi(\cdot, a) \cdot \mathbf{1}_{\{\tau_k(\cdot, a) \geq t_0\}} \\
&\quad \times (M_f^{Q_\omega^n}(t_1 \wedge \tau_k(\cdot, a)) - M_f^{Q_\omega^n}(t_0 \wedge \tau_k(\cdot, a)))]^2] \\
&= \mathbf{E}\left[\left(\frac{1}{n} \sum_{i=1}^n \int_{t_0 \wedge \bar{\tau}_k^{i,n}(\cdot, a)}^{t_1 \wedge \bar{\tau}_k^{i,n}(\cdot, a)} \Psi(\cdot, a) \cdot \mathbf{1}_{\{\bar{\tau}_k^{i,n}(\cdot, a) \geq t_0\}} \right. \right. \\
&\quad \times (\nabla_z f^\top(\bar{X}^{i,n}(s), W^i(s)) \\
&\quad \quad \quad \left. \left. + \nabla_x f^\top(\bar{X}^{i,n}(s), W^i(s)) \right. \right. \\
&\quad \quad \quad \left. \left. \times \sigma(X^{i,n}(s), \bar{\mu}^n(s)) dW^i(s) \right)^2\right] \\
&= \frac{1}{n^2} \sum_{i=1}^n \mathbf{E}\left[\int_{t_0 \wedge \bar{\tau}_k^{i,n}(\cdot, a)}^{t_1 \wedge \bar{\tau}_k^{i,n}(\cdot, a)} |\Psi(\cdot, a) \cdot \mathbf{1}_{\{\bar{\tau}_k^{i,n}(\cdot, a) \geq t_0\}} \right. \\
&\quad \times (\nabla_z f^\top(\bar{X}^{i,n}(s), W^i(s)) \\
&\quad \quad \quad \left. + \nabla_x f^\top(\bar{X}^{i,n}(s), W^i(s)) \right. \\
&\quad \quad \quad \left. \times \sigma(X^{i,n}(s), \bar{\mu}^n(s))|^2 ds\right] \\
&\xrightarrow{n \rightarrow \infty} 0.
\end{aligned}$$

It follows that for each $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$ there is a set $Z_{(k, t_0, t_1, \Psi, f)} \in \tilde{\mathcal{F}}$ such that $\tilde{\mathbf{P}}(Z_{(k, t_0, t_1, \Psi, f)}) = 0$ and

$$\Phi_{(k, t_0, t_1, \Psi, f)}(Q_\omega) = 0 \quad \text{for all } \omega \in \tilde{\Omega} \setminus Z_{(k, t_0, t_1, \Psi, f)}.$$

Let Z be the union of all sets $Z_{(k, t_0, t_1, \Psi, f)}$, $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$. Since \mathcal{T} is countable, we have $Z \in \tilde{\mathcal{F}}$, $\tilde{\mathbf{P}}(Z) = 0$ and

$$\Phi_{(k, t_0, t_1, \Psi, f)}(Q_\omega) = 0 \quad \text{for all } \omega \in \tilde{\Omega} \setminus Z, (k, t_0, t_1, \Psi, f) \in \mathcal{T}.$$

It follows that Q_ω corresponds to a weak solution of (2.7) for $\tilde{\mathbf{P}}$ -almost all $\omega \in \tilde{\Omega}$. \square

The function F in (3.4) is bounded and continuous. The variational lower bound now follows from (5.3), Lemmas 5.1 and 5.2, Fatou's lemma and the definition of I .

6. Variational upper bound. Let $\Theta \in \mathcal{P}_\infty$. We will construct a sequence $(u^N)_{N \in \mathbb{N}}$ with $u^N \in \mathcal{U}_N$ on a common stochastic basis such that (3.5) holds

$$\begin{aligned} & \limsup_{N \rightarrow \infty} \left\{ \frac{1}{2} \mathbf{E} \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^N(t)|^2 dt \right] + \mathbf{E}[F(\bar{\mu}^N)] \right\} \\ & \leq \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr) + F(\Theta_{\mathcal{X}}). \end{aligned}$$

Let (\bar{X}, ρ, W) be the canonical process on \mathcal{Z} (cf. end of Section 2). Then $((\mathcal{Z}, \mathcal{B}(\mathcal{Z}), \Theta), (\tilde{\mathcal{G}}_{t+}^\Theta), (\bar{X}, \rho, W))$ is a weak solution of (2.7). The filtration $(\tilde{\mathcal{G}}_{t+}^\Theta)$ satisfies the usual conditions, where $(\tilde{\mathcal{G}}_t^\Theta)$ denotes the Θ -augmentation of the canonical filtration (\mathcal{G}_t) (cf. Section 4).

Since the relaxed control process ρ appears linearly in (2.7), it corresponds, as far as the dynamics are concerned, to an ordinary (\mathcal{G}_t) -adapted process u , namely

$$u(t, \omega) \doteq \int_{\mathbb{R}^{d_1}} y \rho_{\omega, t}(dy), \quad t \in [0, T], \omega \in \mathcal{Z},$$

where $\rho_{\omega, t}$ is the derivative measure of ρ_ω at time t . For the associated costs, by Jensen's inequality,

$$\begin{aligned} \mathbf{E} \left[\int_0^T |u(t)|^2 dt \right] &= \mathbf{E} \left[\int_0^T \left| \int_{\mathbb{R}^{d_1}} y \rho_t(dy) \right|^2 dt \right] \\ &\leq \mathbf{E} \left[\int_0^T \int_{\mathbb{R}^{d_1}} |y|^2 \rho_t(dy) dt \right] \\ &= \mathbf{E} \left[\int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 \rho(dy \times dt) \right], \end{aligned}$$

whence u performs at least as well as ρ . Let $\tilde{\rho}$ be the relaxed control random variable corresponding to u according to (2.4). In general, $\tilde{\rho} \neq \rho$. However, since both (\bar{X}, ρ, W) and $(\bar{X}, \tilde{\rho}, W)$ are solutions of (2.7) under Θ and since the costs associated with u and thus $\tilde{\rho}$ never exceed the costs associated with ρ , we may and will assume that $\rho = \tilde{\rho}$.

Define a probability space $(\Omega_\infty, \mathcal{F}^\infty, \mathbf{P}_\infty)$ together with a filtration (\mathcal{F}_t^∞) as the countably infinite product of $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}), \Theta)$ and $(\tilde{\mathcal{G}}_{t+}^\Theta)$, respectively. For a typical element of Ω_∞ let us write $\omega = (\omega_1, \omega_2, \dots)$. For $i \in \mathbb{N}$ define

$$W^{i, \infty}(t, \omega) \doteq W(t, \omega_i), \quad u_i^\infty(t, \omega_i) \doteq u(t, \omega_i), \quad \omega \in \Omega_\infty, t \in [0, T].$$

Let $\rho^{i, \infty}$ be the relaxed control random variable corresponding to u_i^∞ . By construction, $(\rho^{i, \infty}, W^{i, \infty})$, $i \in \mathbb{N}$, are independent and identically distributed with common distribution the same as that of (ρ, W) . In particular, $W^{i, \infty}$, $i \in \mathbb{N}$, are independent d_1 -dimensional standard Wiener processes.

For $N \in \mathbb{N}$, let $\tilde{X}^{1,N}, \dots, \tilde{X}^{N,N}$ be the solution to the system of SDEs

$$\begin{aligned} d\tilde{X}^{i,N}(t) &= b(\tilde{X}^{i,N}(t), \tilde{\mu}^N(t)) dt + \sigma(\tilde{X}^{i,N}(t), \tilde{\mu}^N(t)) u_i^\infty(t) dt \\ &\quad + \sigma(\tilde{X}^{i,N}(t), \tilde{\mu}^N(t)) dW^{i,\infty}(t), \quad \tilde{X}^{i,N}(0) = x^{i,N}, \end{aligned}$$

where $\tilde{\mu}^N(t)$ is the empirical measure of $\tilde{X}^{1,N}, \dots, \tilde{X}^{N,N}$ at time t . Thus, $\tilde{X}^{i,N}$ solves (2.3) with the same deterministic initial condition as before, but on a different stochastic basis.

For each $N \in \mathbb{N}$ define, in analogy with (5.2), a $\mathcal{P}(\mathcal{Z})$ -valued random variable according to

$$\tilde{Q}_\omega^N(B \times R \times D) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}^{i,N}(\cdot, \omega)}(B) \cdot \delta_{\rho_\omega^{i,\infty}}(R) \cdot \delta_{W^{i,\infty}(\cdot, \omega)}(D),$$

$B \times R \times D \in \mathcal{B}(\mathcal{Z})$, $\omega \in \Omega_\infty$. In analogy with (5.3) we have

$$\begin{aligned} &\frac{1}{2} \mathbf{E}_\infty \left[\frac{1}{N} \sum_{i=1}^N \int_0^T |u_i^\infty(t)|^2 dt \right] + \mathbf{E}_\infty[F(\tilde{\mu}^N)] \\ (6.1) \quad &= \int_{\Omega_\infty} \left[\int_{\mathcal{R}_1} \left(\frac{1}{2} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \right) \tilde{Q}_{\omega, \mathcal{R}}^N(dr) \right. \\ &\quad \left. + F(\tilde{Q}_{\omega, \mathcal{X}}^N) \right] \mathbf{P}_\infty(d\omega). \end{aligned}$$

Since $(\tilde{\rho}^{i,\infty}, W^{i,\infty})$, $i \in \mathbb{N}$, are i.i.d., the second and third component of $(\tilde{Q}^N)_{N \in \mathbb{N}}$ are tight. Tightness of the first component is an immediate consequence of assumption (A5). Thus, $(\tilde{Q}^N)_{N \in \mathbb{N}}$ is tight as a family of $\mathcal{P}(\mathcal{Z})$ -valued random variables.

Let \tilde{Q} be any limit point of $(\tilde{Q}^N)_{N \in \mathbb{N}}$ defined on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbf{P}})$. By Lemma 5.2 and its proof, it follows that, for $\tilde{\mathbf{P}}$ -almost all $\omega \in \tilde{\Omega}$, \tilde{Q}_ω corresponds to a weak solution of (2.7). Moreover, since $(\rho^{i,\infty}, W^{i,\infty})$, $i \in \mathbb{N}$, are i.i.d. with common distribution (under \mathbf{P}_∞), the same as that of (ρ, W) (under Θ), Varadarajan's theorem [15], page 399, implies that, for $\tilde{\mathbf{P}}$ -almost all $\omega \in \tilde{\Omega}$,

$$\tilde{Q}_{\omega|\mathcal{B}(\mathcal{R}_1 \times \mathcal{W})} = \Theta \circ (\rho, W)^{-1};$$

that is, the joint distribution of the second and third component of the canonical process on \mathcal{Z} under a typical \tilde{Q}_ω equals the joint distribution of the control and Wiener process with which we started.

By assumption (A4), weak sense uniqueness holds for (2.7). Therefore, for $\tilde{\mathbf{P}}$ -almost all $\omega \in \tilde{\Omega}$,

$$\tilde{Q}_\omega = \Theta \circ (\bar{X}, \rho, W)^{-1}.$$

In view of (6.1), the above identification of the limit points establishes (3.5), the variational upper bound.

7. Remarks and extensions. A feature of the weak convergence approach to large deviations is its flexibility. To illustrate this point we show in Section 7.2 how to extend the Laplace principle established in Theorem 3.1 to weakly interacting systems described by stochastic delay (or functional) differential equations. Before, in Section 7.1, we compare our result to the classical large deviation principle (LDP) established in [10].

7.1. Comparison with existing results. In this subsection we compare our results with the now classical work [10]. One of the main assumptions in the latter work is the nondegeneracy of the diffusion coefficient σ . Although the expression for the rate function is well-defined even if the diffusion matrix $\sigma\sigma^T$ is not invertible, the assumption of nondegeneracy is important in the proof of the LDP. Additionally, weak interaction is allowed only through the drift term. Proofs proceed by first establishing a local version of the LDP which is then lifted to a global result using careful exponential probability estimates.

The approach taken in the current paper does not require any exponential estimates and proofs cover the setting of a degenerate σ and models with weak interactions in both the drift and diffusion coefficient. The significant additional assumption made in the current work over [10] is (A3); we require strong existence and uniqueness of solutions to (2.1) whereas the cited paper only assumes weak existence and uniqueness.

Of somewhat lesser significance is the difference in the topology considered on $\mathcal{P}(\mathbb{R}^d)$ and the space over which the LDP is formulated. In particular, in [10] the drift coefficient b need not be continuous on the entire product space $\mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$, where $\mathcal{P}(\mathbb{R}^d)$ is equipped with the topology of weak convergence, but only on $\mathbb{R}^d \times \mathcal{M}_\infty$, where \mathcal{M}_∞ is a set of probability measures on $\mathcal{B}(\mathbb{R}^d)$ which satisfy certain moment bounds in terms of a ‘‘Lyapunov function’’ $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}$. The set \mathcal{M}_∞ is equipped with the ‘‘inductive’’ topology induced by φ [10], Section 5.1. Additional assumptions in terms of this Lyapunov function are imposed which, in particular, ensure that $(\mu^N(t))_{0 \leq t \leq T}$ is a \mathcal{M}_∞ -valued process with continuous sample paths (see (B.2)–(B.4) in [10], Section 5.1). With some additional work, we can relax assumption (A2) on the continuity of b , σ in their second argument and, under Lyapunov function conditions analogous to (B.2)–(B.4), obtain an LDP in a space similar to the one used by [10], namely $C([0, T], \mathcal{M}_\infty)$. A minor difficulty, with the approach taken here, in working with \mathcal{M}_∞ is that the inductive topology is not metrizable. However, one can proceed as follows. Let $\mathcal{P}_\lambda(\mathbb{R}^d)$ be the set of all probability measures $\nu \in \mathcal{P}(\mathbb{R}^d)$ such that $\int \lambda(x)\nu(dx) < \infty$, where $\lambda(x) = |x|k_0(|x|, |x|)$ for some (suitable) symmetric, continuous, nonnegative and nondecreasing function k_0 cf. [28], page 123. The topology of λ -weak convergence, that is, weak convergence plus convergence of λ -moments, makes $\mathcal{P}_\lambda(\mathbb{R}^d)$ a Polish space; cf. Theorems 6.3.1 and 6.3.3 in [28], pages 130–134. Instead of (A2), we would assume that b , σ are continuous as functions defined on $\mathbb{R}^d \times \mathcal{P}_\lambda(\mathbb{R}^d)$ with $\mathcal{P}_\lambda(\mathbb{R}^d)$ carrying the topology of λ -weak convergence. The function λ plays

the role of the Lyapunov function φ used in [10], Section 5.1. The only further modification would regard assumption (A5). In addition to tightness of the sequences of empirical measures $(\bar{\mu}^N)$, one would have to guarantee that the time marginals $\bar{\mu}^N(t)$ stay in $\mathcal{P}_\lambda(\mathbb{R}^d)$. An appropriate condition (which would be analogous to conditions (B.2)–(B.4) in [10], Section 5.1) could be formulated in terms of the Lyapunov function.

The expression for the rate function given in equation (1.5) in [10] is different from the one given in Theorem 3.1 of this paper. The integrand in particular involves the maximization over a class of smooth test functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ at each time point $t \in [0, T]$. In the case where the diffusion coefficient σ is the identity matrix, test functions $f(t, \cdot)$, $t \in [0, T]$, induce feedback controls for (3.2) through $u(t, \omega) \doteq \nabla_x f(t, \bar{X}(t, \omega))$, cf. Remark 3.2. In this way one can see, at least formally, the equivalence of our expression for the rate function and the expression derived in [10].

7.2. Processes with delay. Our approach allows one to treat more general Itô equations than those of diffusion type with very little additional effort. A good example are SDEs whose coefficients are allowed to depend on the entire past of the state trajectories. Let us make this more precise. Suppose that the coefficients b, σ are progressive functionals defined on $[0, T] \times \mathcal{X} \times \mathcal{P}(\mathbb{R}^d)$, where we recall that $\mathcal{X} = \mathbf{C}([0, T], \mathbb{R}^d)$; that is, b, σ are Borel measurable and for each $t \in [0, T]$, b, σ restricted to $[0, t] \times \mathcal{X} \times \mathcal{P}(\mathbb{R}^d)$ is measurable with respect to $\mathcal{B}([0, t]) \times \mathcal{G}_t^{\mathcal{X}} \times \mathcal{B}(\mathcal{P}(\mathbb{R}^d))$ where $\mathcal{G}_t^{\mathcal{X}}$ is the σ -algebra generated by the coordinate process on \mathcal{X} . Equation (2.1), the prelimit equation for an individual particle (the i th out of N), takes the form

$$(7.1) \quad dX^{i,N}(t) = b(t, X^{i,N}, \mu^N(t)) dt + \sigma(t, X^{i,N}, \mu^N(t)) dW^i(t).$$

The system of N equations given by (7.1) is a system of stochastic functional differential equations or stochastic delay differential equations (SFDEs or SDDEs). The corresponding uncontrolled limit equation reads

$$(7.2) \quad dX(t) = b(t, X, \text{Law}(X(t))) dt + \sigma(t, X, \text{Law}(X(t))) dW(t),$$

while the controlled versions of (7.1) and (7.2) will be

$$(7.3) \quad \begin{aligned} d\bar{X}^{i,N}(t) &= b(t, \bar{X}^{i,N}, \bar{\mu}^N(t)) dt + \sigma(t, \bar{X}^{i,N}, \bar{\mu}^N(t)) u_i(t) dt \\ &\quad + \sigma(t, \bar{X}^{i,N}, \bar{\mu}^N(t)) u_i(t) dW^i(t), \end{aligned}$$

$$(7.4) \quad \begin{aligned} d\bar{X}(t) &= b(t, \bar{X}, \text{Law}(\bar{X}(t))) dt + \left(\int_{\mathbb{R}^{d_1}} \sigma(t, \bar{X}, \text{Law}(\bar{X}(t))) y \rho_t(dy) \right) dt \\ &\quad + \sigma(t, \bar{X}, \text{Law}(\bar{X}(t))) u(t) dW(t), \end{aligned}$$

respectively. In (7.3) u_i is the i th component of $u = (u_1, \dots, u_N)$ for some $u \in \mathcal{U}_N$, while ρ in (7.4) is an adapted \mathcal{R}_1 -valued random variable as in (2.7).

The Laplace principle can now be established in the same way as above except for two points which need modification. Those are the formulation of the local martingale problem in Section 4 and the continuity assumption (A3')–(A5') the analogues of assumptions (A3)–(A5), which are obtained by replacing all references to (2.1), (2.2), (2.3), (2.7) with (7.1), (7.2), (7.3), (7.4), respectively.

As to the martingale problem, we have to redefine the processes M_f^Θ and the “generators” $\mathcal{A}_s^\Theta(f)$ according to

$$M_f^\Theta(t, (\varphi, r, w)) \doteq f(\varphi(t), w(t)) - f(\varphi(0), 0) - \int_0^t \int_{\mathbb{R}^{d_1}} \mathcal{A}_s^\Theta(f)(\varphi, y, w(s)) r_s(dy) ds,$$

where for $s \in [0, T]$, $\varphi \in \mathcal{X}$, $y, z \in \mathbb{R}^{d_1}$,

$$\begin{aligned} \mathcal{A}_s^\Theta(f)(\varphi, y, z) &\doteq \langle b(s, \varphi, \nu_\Theta(s)) + \sigma(s, \varphi, \nu_\Theta(s))y, \nabla_x f(\varphi(s), z) \rangle \\ &+ \frac{1}{2} \sum_{j,k=1}^d (\sigma \sigma^\top)_{jk}(s, \varphi, \nu_\Theta(s)) \frac{\partial^2 f}{\partial x_j \partial x_k}(\varphi(s), z) \\ &+ \frac{1}{2} \sum_{l=1}^{d_1} \frac{\partial^2 f}{\partial z_l \partial z_l}(\varphi(s), z) \\ &+ \sum_{k=1}^d \sum_{l=1}^{d_1} \sigma_{kl}(s, \varphi, \nu_\Theta(s)) \frac{\partial^2 f}{\partial x_k \partial z_l}(\varphi(s), z). \end{aligned}$$

Notice that the test functions f are still elements of $\mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$. With these redefinitions, Lemma 4.1 continues to hold.

Assumption (A2) about the continuity of b , σ has to be modified in order to account for the time dependence and be supplemented by a condition of uniform continuity and boundedness, which is automatically satisfied in the diffusion case.

(A2') The functions $b(t, \cdot, \cdot)$, $\sigma(t, \cdot, \cdot)$ are continuous, and uniformly continuous and bounded on sets $B \times P$ whenever $B \subset \mathcal{X}$ is bounded and $P \subset \mathcal{P}(\mathbb{R}^d)$ is compact, uniformly in $t \in [0, T]$.

Define the set \mathcal{P}_∞^* of probability measures on $\mathcal{B}(\mathcal{Z})$ as the set \mathcal{P}_∞ in Section 3, replacing reference to (2.7) with (7.4). Then the following large deviation (or Laplace) principle holds.

THEOREM 7.1. *Grant assumptions (A1), (A2')–(A5'). Then the family of empirical measures $\{\mu^N, N \in \mathbb{N}\}$ associated with (7.1) satisfies the Laplace principle with rate function*

$$\tilde{I}(\theta) = \inf_{\Theta \in \mathcal{P}_\infty^* : \Theta_{\mathcal{X}} = \theta} \frac{1}{2} \int_{\mathcal{R}} \int_{\mathbb{R}^{d_1} \times [0, T]} |y|^2 r(dy \times dt) \Theta_{\mathcal{R}}(dr).$$

Note that there is also a simpler-looking form of the rate function as in Remark 3.2. The proof of Theorem 7.1 is completely analogous to that of Theorem 3.1 given in Sections 5 and 6. The proof of Lemma 5.2, in particular, and specifically the use of the local martingale problem and randomized stopping times there was tailored to fit not only the diffusion case, but the case of dynamics with delay as well.

Finally, note that we could further generalize our model to include the case of coefficients b, σ which also depend on the past of the empirical process. In this case, b, σ would be progressive functionals defined on $[0, T] \times \mathcal{X} \times \mathcal{P}(\mathcal{X})$, and a Laplace principle could be established in the same way as before.

APPENDIX: LOCAL MARTINGALES WITH RESPECT TO $(\hat{\mathcal{G}}_t)$ AND (\mathcal{G}_t)

Let the notation be that of the proof of Lemma 5.2 in Section 5. Let $\Theta \in \mathcal{P}(Z)$, $f \in \mathbf{C}^2(\mathbb{R}^d)$, and set $M(t) \doteq M_f^\Theta(t)$, $t \in [0, T]$. Notice that M is a random object defined on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ with values in $\mathcal{X} = \mathbf{C}([0, T], \mathbb{R}^d)$, which can be identified with the random object living on $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$ given by

$$\mathcal{Z} \times [0, 1] \ni (z, s) \mapsto (M(t, z))_{t \in [0, T]} \in \mathcal{X}.$$

Let $k \in \mathbb{N}$. Suppose that $M(\cdot \wedge \tau_k)$ is a martingale under $\hat{\Theta} = \Theta \times \lambda$ with respect to the canonical filtration $(\hat{\mathcal{G}}_t)$ in $\mathcal{B}(\hat{\mathcal{Z}})$. Set

$$\tau_k^\circ(z) \doteq \tau_k(z, 0), \quad z \in \mathcal{Z}.$$

We claim that $M(\cdot \wedge \tau_k^\circ)$ is a martingale under Θ with respect to the canonical filtration (\mathcal{G}_t) in $\mathcal{B}(\mathcal{Z})$.

PROOF OF THE MARTINGALE PROPERTY. Since τ_k is a $(\hat{\mathcal{G}}_t)$ -stopping time and $\hat{\mathcal{G}}_t = \mathcal{G}_t \times \mathcal{B}([0, 1])$, $t \in [0, T]$, it follows that τ_k° is a (\mathcal{G}_t) -stopping time. Moreover, τ_k° is also a $(\hat{\mathcal{G}}_t)$ -stopping time, because \mathcal{G}_t can be identified with $\mathcal{G}_t \times \{\emptyset, [0, 1]\}$, $t \in [0, T]$, and $(\mathcal{G}_t \times \{\emptyset, [0, 1]\})$ is a subfiltration of $(\hat{\mathcal{G}}_t)$.

Let $s, t \in [0, T]$, $s \leq t$. We have to show that

$$\mathbf{E}_\Theta[M(t \wedge \tau_k^\circ) \cdot \mathbf{1}_Z] = \mathbf{E}_\Theta[M(s \wedge \tau_k^\circ) \cdot \mathbf{1}_Z] \quad \text{for all } Z \in \mathcal{G}_s.$$

Since $M(\cdot \wedge \tau_k)$ is a martingale under $\hat{\Theta}$ with respect to $(\hat{\mathcal{G}}_t)$ and τ_k° is also a $(\hat{\mathcal{G}}_t)$ -stopping time, it follows that $M(\cdot \wedge \tau_k \wedge \tau_k^\circ)$ is a martingale under $\hat{\Theta}$ with respect to $(\hat{\mathcal{G}}_t)$. Yet for all $(z, t) \in \hat{\mathcal{Z}}$,

$$(\tau_k \wedge \tau_k^\circ)(z, t) = \tau_k(z, t) \wedge \tau_k(z, 0) = \tau_k(z, 0) = \tau_k^\circ(z)$$

by construction of τ_k and definition of τ_k° . Hence we know that

$$\mathbf{E}_{\hat{\Theta}}[M(t \wedge \tau_k^\circ) \cdot \mathbf{1}_{\hat{Z}}] = \mathbf{E}_{\hat{\Theta}}[M(s \wedge \tau_k^\circ) \cdot \mathbf{1}_{\hat{Z}}] \quad \text{for all } \hat{Z} \in \hat{\mathcal{G}}_s.$$

Let $Z \in \mathcal{G}_s$. Then $Z \times [0, 1] \in \hat{\mathcal{G}}_s$ and, by Fubini's theorem,

$$\begin{aligned}
 \mathbf{E}_\Theta[M(t \wedge \tau_k^\circ) \cdot \mathbf{1}_Z] &= \int_{\mathcal{Z}} M(t \wedge \tau_k^\circ(z)) \cdot \mathbf{1}_Z(z) \Theta(dz) \\
 &= \int_{[0,1]} \int_{\mathcal{Z}} M(t \wedge \tau_k^\circ(z)) \cdot \mathbf{1}_{Z \times [0,1]}(z, a) \Theta(dz) \lambda(da) \\
 &= \int_{\mathcal{Z} \times [0,1]} M(t \wedge \tau_k^\circ(z)) \cdot \mathbf{1}_{Z \times [0,1]}(z, a) \hat{\Theta}(dz \times da) \\
 &= \mathbf{E}_{\hat{\Theta}}[M(t \wedge \tau_k^\circ) \cdot \mathbf{1}_{Z \times [0,1]}] \\
 &= \mathbf{E}_{\hat{\Theta}}[M(s \wedge \tau_k^\circ) \cdot \mathbf{1}_{Z \times [0,1]}] \\
 &= \int_{\mathcal{Z} \times [0,1]} M(s \wedge \tau_k^\circ(z)) \cdot \mathbf{1}_{Z \times [0,1]}(z, a) \hat{\Theta}(dz \times da) \\
 &= \mathbf{E}_\Theta[M(s \wedge \tau_k^\circ) \cdot \mathbf{1}_Z]. \quad \square
 \end{aligned}$$

Acknowledgments. The authors thank the Editor and an anonymous referee for their critique and helpful comments and suggestions.

REFERENCES

- [1] AROUS, G. B. and GUIONNET, A. (1995). Large deviations for Langevin spin glass dynamics. *Probab. Theory Related Fields* **102** 455–509. [MR1346262](#)
- [2] BEN AROUS, G. and BRUNAUD, M. (1990). Méthode de Laplace: Étude variationnelle des fluctuations de diffusions de type “champ moyen.” *Stochastics Stochastics Rep.* **31** 79–144. [MR1080535](#)
- [3] BEN AROUS, G. and ZEITOUNI, O. (1999). Increasing propagation of chaos for mean field models. *Ann. Inst. H. Poincaré Probab. Statist.* **35** 85–102. [MR1669916](#)
- [4] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd ed. Wiley, New York. [MR1700749](#)
- [5] BOUÉ, M. and DUPUIS, P. (1998). A variational representation for certain functionals of Brownian motion. *Ann. Probab.* **26** 1641–1659. [MR1675051](#)
- [6] BUDHIRAJA, A. and DUPUIS, P. (2000). A variational representation for positive functionals of infinite dimensional Brownian motion. *Probab. Math. Statist.* **20** 39–61. [MR1785237](#)
- [7] BUDHIRAJA, A., DUPUIS, P. and MAROULAS, V. (2008). Large deviations for infinite dimensional stochastic dynamical systems. *Ann. Probab.* **36** 1390–1420. [MR2435853](#)
- [8] BUDHIRAJA, A., DUPUIS, P. and MAROULAS, V. (2011). Variational representations for continuous time processes. *Ann. Inst. Henri Poincaré Probab. Stat.* **47** 725–747. [MR2841073](#)
- [9] DAWSON, D. A. and DEL MORAL, P. (2005). Large deviations for interacting processes in the strong topology. In *Statistical Modeling and Analysis for Complex Data Problems. GERAD 25th Anniv. Ser.* **1** 179–208. Springer, New York. [MR2189537](#)
- [10] DAWSON, D. A. and GÄRTNER, J. (1987). Large deviations from the McKean–Vlasov limit for weakly interacting diffusions. *Stochastics* **20** 247–308. [MR0885876](#)
- [11] DAWSON, D. A. and GÄRTNER, J. (1994). Multilevel large deviations and interacting diffusions. *Probab. Theory Related Fields* **98** 423–487. [MR1271106](#)
- [12] DEL MORAL, P. and GUIONNET, A. (1998). Large deviations for interacting particle systems: Applications to non-linear filtering. *Stochastic Process. Appl.* **78** 69–95. [MR1653296](#)
- [13] DJEHICHE, B. and SCHIED, A. (1998). Large deviations for hierarchical systems of interacting jump processes. *J. Theoret. Probab.* **11** 1–24. [MR1607396](#)

- [14] DUAN, J. and MILLET, A. (2009). Large deviations for the Boussinesq equations under random influences. *Stochastic Process. Appl.* **119** 2052–2081. [MR2519356](#)
- [15] DUDLEY, R. M. (2002). *Real Analysis and Probability*. *Cambridge Studies in Advanced Mathematics* **74**. Cambridge Univ. Press, Cambridge. [MR1932358](#)
- [16] DUPUIS, P. and ELLIS, R. S. (1997). *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley, New York. [MR1431744](#)
- [17] FENG, J. and KURTZ, T. G. (2006). *Large Deviations for Stochastic Processes*. *Mathematical Surveys and Monographs* **131**. Amer. Math. Soc., Providence, RI. [MR2260560](#)
- [18] FONTBONA, J. (2004). Uniqueness for a weak nonlinear evolution equation and large deviations for diffusing particles with electrostatic repulsion. *Stochastic Process. Appl.* **112** 119–144. [MR2062570](#)
- [19] FUNAKI, T. (1984). A certain class of diffusion processes associated with nonlinear parabolic equations. *Z. Wahrsch. Verw. Gebiete* **67** 331–348. [MR0762085](#)
- [20] GÄRTNER, J. (1988). On the McKean–Vlasov limit for interacting diffusions. *Math. Nachr.* **137** 197–248. [MR0968996](#)
- [21] HERRMANN, S., IMKELLER, P. and PEITHMANN, D. (2008). Large deviations and a Kramers’ type law for self-stabilizing diffusions. *Ann. Appl. Probab.* **18** 1379–1423. [MR2434175](#)
- [22] KARATZAS, I. and SHREVE, S. E. (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. *Graduate Texts in Mathematics* **113**. Springer, New York. [MR1121940](#)
- [23] KUSHNER, H. J. (1990). *Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems*. *Systems & Control: Foundations & Applications* **3**. Birkhäuser, Boston, MA. [MR1102242](#)
- [24] LÉONARD, C. (1995). Large deviations for long range interacting particle systems with jumps. *Ann. Inst. H. Poincaré Probab. Statist.* **31** 289–323. [MR1324810](#)
- [25] LÉONARD, C. (1995). On large deviations for particle systems associated with spatially homogeneous Boltzmann type equations. *Probab. Theory Related Fields* **101** 1–44. [MR1314173](#)
- [26] MÉLÉARD, S. (1996). Asymptotic behaviour of some interacting particle systems; McKean–Vlasov and Boltzmann models. In *Probabilistic Models for Nonlinear Partial Differential Equations (Montecatini Terme, 1995)*. *Lecture Notes in Math.* **1627** 42–95. Springer, Berlin. [MR1431299](#)
- [27] OELSCHLÄGER, K. (1984). A martingale approach to the law of large numbers for weakly interacting stochastic processes. *Ann. Probab.* **12** 458–479. [MR0735849](#)
- [28] RACHEV, S. T. (1991). *Probability Metrics and the Stability of Stochastic Models*. Wiley, Chichester. [MR1105086](#)
- [29] REN, J. and ZHANG, X. (2005). Schilder theorem for the Brownian motion on the diffeomorphism group of the circle. *J. Funct. Anal.* **224** 107–133. [MR2139106](#)
- [30] RÖCKNER, M., ZHANG, T. and ZHANG, X. (2010). Large deviations for stochastic tamed 3D Navier-Stokes equations. *Appl. Math. Optim.* **61** 267–285. [MR2585144](#)
- [31] SRITHARAN, S. S. and SUNDAR, P. (2006). Large deviations for the two-dimensional Navier-Stokes equations with multiplicative noise. *Stochastic Process. Appl.* **116** 1636–1659. [MR2269220](#)
- [32] STROOCK, D. W. and VARADHAN, S. R. S. (1979). *Multidimensional Diffusion Processes*. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **233**. Springer, Berlin. [MR0532498](#)
- [33] TANAKA, H. (1984). Limit theorems for certain diffusion processes with interaction. In *Stochastic Analysis (Katata/Kyoto, 1982)*. *North-Holland Mathematical Library* **32** 469–488. North-Holland, Amsterdam. [MR0780770](#)

- [34] YONG, J. and ZHOU, X. Y. (1999). *Stochastic Controls: Hamiltonian systems and HJB equations. Applications of Mathematics (New York)* **43**. Springer, New York. [MR1696772](#)

A. BUDHIRAJA
DEPARTMENT OF STATISTICS
AND OPERATIONS RESEARCH
UNIVERSITY OF NORTH CAROLINA
AT CHAPEL HILL
CHAPEL HILL, NORTH CAROLINA 27599
USA
E-MAIL: budhiraj@email.unc.edu

P. DUPUIS
M. FISCHER
LEFSCHETZ CENTER FOR DYNAMICAL SYSTEMS
DIVISION OF APPLIED MATHEMATICS
BROWN UNIVERSITY
PROVIDENCE, RHODE ISLAND 02912
USA
E-MAIL: dupuis@dam.brown.edu
fischer@statlab.uni-heidelberg.de