A basic introduction to large deviations: Theory, applications, simulations[*]

Hugo Touchette

School of Mathematical Sciences, Queen Mary, University of London
London E1 4NS, United Kingdom

March 1, 2012

Abstract

The theory of large deviations deals with the probabilities of rare events (or fluctuations) that are exponentially small as a function of some parameter, e.g., the number of random components of a system, the time over which a stochastic system is observed, the amplitude of the noise perturbing a dynamical system or the temperature of a chemical reaction. The theory has applications in many different scientific fields, ranging from queuing theory to statistics and from finance to engineering. It is also increasingly used in statistical physics for studying both equilibrium and nonequilibrium systems. In this context, deep analogies can be made between familiar concepts of statistical physics, such as the entropy and the free energy, and concepts of large deviation theory having more technical names, such as the rate function and the scaled cumulant generating function.

The first part of these notes introduces the basic elements of large deviation theory at a level appropriate for advanced undergraduate and graduate students in physics, engineering, chemistry, and mathematics. The focus there is on the simple but powerful ideas behind large deviation theory, stated in non-technical terms, and on the application of these ideas in simple stochastic processes, such as sums of independent and identically distributed random variables and Markov processes. Some physical applications of these processes are covered in exercises contained at the end of each section.

In the second part, the problem of numerically evaluating large deviation probabilities is treated at a basic level. The fundamental idea of importance sampling is introduced there together with its sister idea, the exponential change of measure. Other numerical methods based on sample means and generating functions, with applications to Markov processes, are also covered.

Contents

1   Introduction

The goal of these lecture notes, as the title says, is to give a basic introduction to the theory of large deviations at three levels: theory, applications and simulations. The notes follow closely my recent review paper on large deviations and their applications in statistical mechanics [48], but are, in a way, both narrower and wider in scope than this review paper.

They are narrower, in the sense that the mathematical notations have been cut down to a minimum in order for the theory to be understood by advanced undergraduate and graduate students in science and engineering, having a basic background in probability theory and stochastic processes (see, e.g., [27]). The simplification of the mathematics amounts essentially to two things: i) to focus on random variables taking values in $\mathbb{R}$ or $\mathbb{R}^D$, and ii) to state all the results in terms of probability densities, and so to assume that probability densities always exist, if only in a weak sense. These simplifications are justified for most applications and are convenient for conveying the essential ideas of the theory in a clear way, without the hindrance of technical notations.

These notes also go beyond the review paper [48], in that they cover subjects not contained in that paper, in particular the subject of numerical estimation or *simulation* of large deviation probabilities. This is an important subject that I intend to cover in more depth in the future.

Sections 4 and 5 of these notes are a first and somewhat brief attempt in this direction. Far from covering all the methods that have been developed to simulate large deviations and rare events, they concentrate on the central idea of large deviation simulations, namely that of exponential change of measure, and they elaborate from there on certain simulation techniques that are easily applicable to sums of independent random variables, Markov chains, stochastic differential equations and continuous-time Markov processes in general.

Many of these applications are covered in the exercises contained at the end of each section. The level of difficulty of these exercises is quite varied: some are there to practice the material presented, while others go beyond that material and may take hours, if not days or weeks, to solve completely. For convenience, I have rated them according to Knuth's logarithmic scale.[1]

In closing this introduction, let me emphasize again that these notes are not meant to be complete in any way. For one thing, they lack the rigorous

---

[1] 00 = Immediate; 10 = Simple; 20 = Medium; 30 = Moderately hard; 40 = Term project; 50 = Research problem. See the superscript attached to each exercise.

notation needed for handling large deviations in a precise mathematical way, and only give a hint of the vast subject that is large deviation and rare event simulation. On the simulation side, they also skip over the subject of error estimates, which would fill an entire section in itself. In spite of this, I hope that the material covered will have the effect of enticing readers to learn more about large deviations and of conveying a sense of the wide applicability, depth and beauty of this subject, both at the theoretical and computational levels.

2 Basic elements of large deviation theory

2.1 Examples of large deviations

We start our study of large deviation theory by considering a sum of real random variables (RV for short) having the form

$$S_n = \frac{1}{n} \sum_{i=1}^{n} X_i. \tag{2.1}$$

Such a sum is often referred to in mathematics or statistics as a **sample mean**. We are interested in computing the **probability density function** $p_{S_n}(s)$ (pdf for short) of S_n in the simple case where the n RVs are mutually **independent and identically distributed** (IID for short).[2] This means that the joint pdf of $X_1, \ldots, X_n$ factorizes as follows:

$$p(X_1, \ldots, X_n) = \prod_{i=1}^{n} p(X_i), \tag{2.2}$$

with $p(X_i)$ a fixed pdf for each of the X_i's.[3]

We compare next two cases for $p(X_i)$:

- **Gaussian pdf:**

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/(2\sigma^2)}, \quad x \in \mathbb{R}, \tag{2.3}$$

where $\mu = E[X]$ is the **mean** of X and $\sigma^2 = E[(X-\mu)^2]$ its **variance**.[4]

---

[2] The pdf of S_n will be denoted by $p_{S_n}(s)$ or more simply by $p(S_n)$ when no confusion arises.

[3] To avoid confusion, we should write $p_{X_i}(x)$, but in order not to overload the notation, we will simply write $p(X_i)$. The context should make clear what RV we are referring to.

[4] The symbol $E[X]$ denotes the expectation or expected value of X, which is often denoted by $\langle X \rangle$ in physics.

- **Exponential pdf:**

$$p(x) = \frac{1}{\mu}\mathrm{e}^{-x/\mu}, \quad x \in [0, \infty), \tag{2.4}$$

with mean $E[X] = \mu > 0$.

What is the form of $p_{S_n}(s)$ for each pdf?

To find out, we write the pdf associated with the event $S_n = s$ by summing the pdf of all the values or **realizations** $(x_1, \ldots, x_n) \in \mathbb{R}^n$ of $X_1, \ldots, X_n$ such that $S_n = s$.[5] In terms of Dirac's delta function $\delta(x)$, this is written as

$$p_{S_n}(s) = \int_{\mathbb{R}} \mathrm{d}x_1 \cdots \int_{\mathbb{R}} \mathrm{d}x_n \; \delta\left(\sum_{i=1}^{n} x_i - ns\right) p(x_1, \ldots, x_n). \tag{2.5}$$

From this expression, we can then obtain an explicit expression for $p_{S_n}(s)$ by using the method of generating functions (see Exercise 2.7.1) or by substituting the Fourier representation of $\delta(x)$ above and by explicitly evaluating the $n + 1$ resulting integrals. The result obtained for both the Gaussian and the exponential densities has the general form

$$p_{S_n}(s) \approx \mathrm{e}^{-nI(s)}, \tag{2.6}$$

where

$$I(s) = \frac{(x - \mu)^2}{2\sigma^2}, \qquad s \in \mathbb{R} \tag{2.7}$$

for the Gaussian pdf, whereas

$$I(s) = \frac{s}{\mu} - 1 - \ln \frac{s}{\mu}, \qquad s \geq 0 \tag{2.8}$$

for the exponential pdf.

We will come to understand the meaning of the approximation sign ($\approx$) more clearly in the next subsection. For now we just take it as meaning that the dominant behaviour of $p(S_n)$ as a function of n is a decaying exponential in n. Other terms in n that may appear in the exact expression of $p(S_n)$ are sub-exponential in n.

The picture of the behaviour of $p(S_n)$ that we obtain from the result of (2.6) is that $p(S_n)$ decays to 0 exponentially fast with n for all values $S_n = s$ for which the function $I(s)$, which controls the rate of decay of

---

[5]Whenever possible, random variables will be denoted by uppercase letters, while their values or realizations will be denoted by lowercase letters. This follows the convention used in probability theory. Thus, we will write $X = x$ to mean that the RV X takes the value x.
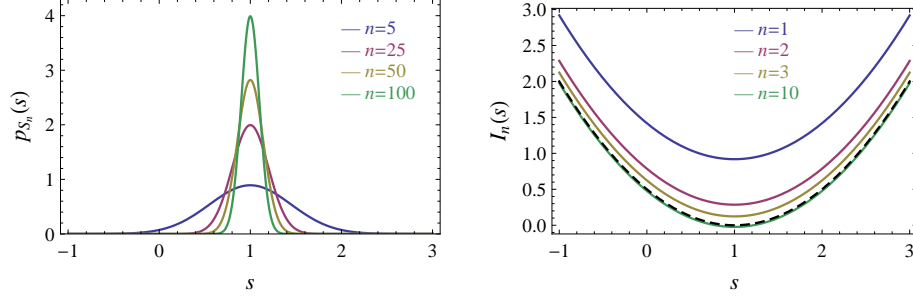
Figure 2.1: (Left) pdf $p_{S_n}(s)$ of the Gaussian sample mean S_n for $\mu = \sigma = 1$. (Right) $I_n(s) = -\frac{1}{n} \ln p_{S_n}(s)$ for different values of n demonstrating a rapid convergence towards the rate function $I(s)$ (dashed line).

$p(S_n)$, is positive. But notice that $I(s) \geq 0$ for both the Gaussian and exponential densities, and that $I(s) = 0$ in both cases only for $s = \mu = E[X_i]$. Therefore, since the pdf of S_n is normalized, it must get more and more concentrated around the value $s = \mu$ as $n \to \infty$ (see Fig. 2.1), which means that $p_{S_n}(s) \to \delta(s - \mu)$ in this limit. We say in this case that S_n converges in probability or in density to its mean.

As a variation of the Gaussian and exponential samples means, let us now consider a sum of **discrete random variables** having a probability distribution $P(X_i)$ instead of continuous random variables with a pdf $p(X_i)$.[6] To be specific, consider the case of IID **Bernoulli RVs** $X_1, \ldots, X_n$ taking values in the set $\{0, 1\}$ with probability $P(X_i = 0) = 1 - \alpha$ and $P(X_i = 1) = \alpha$. What is the form of the probability distribution $P(S_n)$ of S_n in this case?

Notice that we are speaking of a probability distribution because S_n is now a discrete variable taking values in the set $\{0, 1/n, 2/n, \ldots, (n-1)/n, 1\}$. In the previous Gaussian and exponential examples, S_n was a continuous variable characterized by its pdf $p(S_n)$.

With this in mind, we can obtain the exact expression of $P(S_n)$ using methods similar to those used to obtain $p(S_n)$ (see Exercise 2.7.4). The result is different from that found for the Gaussian and exponential densities, but what is remarkable is that the distribution of the Bernoulli sample mean

---

[6]Following the convention in probability theory, we use the lowercase p for continuous probability densities and the uppercase P for discrete probability distributions and probability assignments in general. Moreover, following the notation used before, we will denote the distribution of a discrete S_n by $P_{S_n}(s)$ or simply $P(S_n)$.

also contains a dominant exponential term having the form

$$P_{S_n}(s) \approx \mathrm{e}^{-nI(s)}, \tag{2.9}$$

where $I(s)$ is now given by

$$I(s) = s \ln \frac{s}{\alpha} + (1-s) \ln \frac{1-s}{1-\alpha}, \qquad s \in [0,1]. \tag{2.10}$$

The behaviour of the exact expression of $P(S_n)$ as n grows is shown in Fig. 2.2 together with the plot of $I(s)$ as given by Eq. (2.10). Notice how $P(S_n)$ concentrates around its mean $\mu = E[X_i] = \alpha$ as a result of the fact that $I(s) \geq 0$ and that $s = \mu$ is the only value of S_n for which $I = 0$. Notice also how the support of $P(S_n)$ becomes "denser" as $n \to \infty$, and compare this property with the fact that $I(s)$ is a continuous function despite S_n being discrete. Should $I(s)$ not be defined for discrete values if S_n is a discrete RV? We address this question next.

2.2 The large deviation principle

The general exponential form $\mathrm{e}^{-nI(s)}$ that we found for our three previous sample means (Gaussian, exponential and Bernoulli) is the founding result or property of large deviation theory, referred to as the large deviation principle. The reason why a whole theory can be built on such a seemingly simple result is that it arises in the context of many stochastic processes, not just IID sample means, as we will come to see in the next sections, and as can be seen from other contributions to this volume; see, e.g., Engel's.[7]

The rigorous, mathematical definition of the large deviation principle involves many concepts of topology and measure theory that are too technical to be presented here (see Sec. 3.1 and Appendix B of [48]). For simplicity, we will say here that a random variable S_n or its pdf $p(S_n)$ satisfies a **large deviation principle** (LDP) if the following limit exists:

$$\lim_{n \to \infty} -\frac{1}{n} \ln p_{S_n}(s) = I(s) \tag{2.11}$$

and gives rise to a function $I(s)$, called the **rate function**, which is not everywhere zero.

The relation between this definition and the approximation notation used earlier should be clear: the fact that the behaviour of $p_{S_n}(s)$ is dominated

---

[7]The volume referred to is the volume of lecture notes produced for the summer school; see http://www.mcs.uni-oldenburg.de/
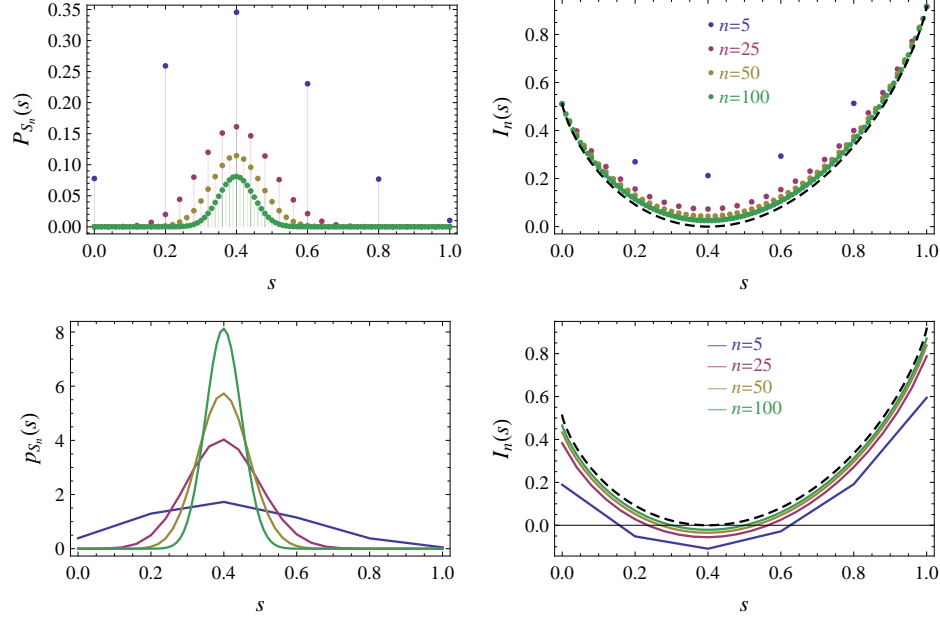
Figure 2.2: (Top left) Discrete probability distribution $P_{S_n}(s)$ of the Bernoulli sample mean for $\alpha = 0.4$ and different values of n. (Top right) Finite-n rate function $I_n(s) = -\frac{1}{n} \ln P_{S_n}(s)$. The rate function $I(s)$ is the dashed line. (Bottom left). Coarse-grained pdf $p_{S_n}(s)$ for the Bernoulli sample mean. (Bottom right) $I_n(s) = -\frac{1}{n} \ln p_{S_n}(s)$ as obtained from the coarse-grained pdf.

for large n by a decaying exponential means that the exact pdf of S_n can be written as

$$p_{S_n}(s) = \mathrm{e}^{-nI(s)+o(n)} \tag{2.12}$$

where $o(n)$ stands for any correction term that is sub-linear in n. Taking the large deviation limit of (2.11) then yields

$$\lim_{n\to\infty} -\frac{1}{n} \ln p_{S_n}(s) = I(s) - \lim_{n\to\infty} \frac{o(n)}{n} = I(s), \tag{2.13}$$

since $o(n)/n \to 0$. We see therefore that the large deviation limit of Eq. (2.11) is the limit needed to retain the dominant exponential term in $p(S_n)$ while discarding any other sub-exponential terms. For this reason, large deviation theory is often said to be concerned with estimates of probabilities on the logarithmic scale.

This point is illustrated in Fig. 2.2. There we see that the function

$$I_n(s) = -\frac{1}{n} \ln p_{S_n}(s) \tag{2.14}$$

is not quite equal to the limiting rate function $I(s)$, because of terms of order $o(n)/n = o(1)$; however, it does converge to $I(s)$ as n gets larger. Plotting $p(S_n)$ on this scale therefore reveals the rate function in the limit of large n. This convergence will be encountered repeatedly in the sections on the numerical evaluation of large deviation probabilities.

It should be emphasized again that the definition of the LDP given above is a simplification of the rigorous definition used in mathematics, due to the mathematician Srinivas Varadhan.[8] The real, mathematical definition is expressed in terms of probability measures of certain sets rather than in terms of probability densities, and involves upper and lower bounds on these probabilities rather than a simple limit (see Sec. 3.1 and Appendix B of [48]). The mathematical definition also applies a priori to any RVs, not just continuous RVs with a pdf.

In these notes, we will simplify the mathematics by assuming that the random variables or stochastic processes that we study have a pdf. In fact, we will often assume that pdfs exist even for RVs that are not continuous but "look" continuous at some scale.

To illustrate this point, consider again the Bernoulli sample mean. We noticed that S_n in this case is not a continuous RV, and so does not have a pdf. However, we also noticed that the values that S_n can take become dense in the interval $[0, 1]$ as $n \to \infty$, which means that the support of the discrete probability distribution $P(S_n)$ becomes dense in this limit, as shown in Fig. 2.2. From a practical point of view, it therefore makes sense to treat S_n as a continuous RV for large n by interpolating $P(S_n)$ to a continuous function $p(S_n)$ representing the "probability density" of S_n.

This pdf is obtained in general simply by considering the probability $P(S_n \in [s, s + \Delta s])$ that S_n takes a value in a tiny interval surrounding the value s, and by then dividing this probability by the "size" Δs of that interval:

$$p_{S_n}(s) = \frac{P(S_n \in [s, s + \Delta s])}{\Delta s}. \tag{2.15}$$

The pdf obtained in this way is referred to as a **smoothed**, **coarse-grained**

---

[8] Recipient of the 2007 Abel Prize for his "fundamental contributions to probability theory and in particular for creating a unified theory of large deviations."

or **weak density**.[9] In the case of the Bernoulli sample mean, it is simply given by

$$p_{S_n}(s) = nP_{S_n}(s) \tag{2.16}$$

if the spacing Δs between the values of S_n is chosen to be $1/n$.

This process of replacing a discrete variable by a continuous one in some limit or at some scale is known in physics as the **continuous** or **continuum limit** or as the **macroscopic limit**. In mathematics, this limit is expressed via the notion of weak convergence (see [16]).

2.3 The Gärtner-Ellis Theorem

Our goal in the next sections will be to study random variables and stochastic processes that satisfy an LDP and to find analytical as well as numerical ways to obtain the corresponding rate function. There are many ways whereby a random variable, say S_n, can be shown to satisfy an LDP:

- **Direct method:** Find the expression of $p(S_n)$ and show that it has the form of the LDP;

- **Indirect method:** Calculate certain functions of S_n, such as generating functions, whose properties can be used to infer that S_n satisfies an LDP;

- **Contraction method:** Relate S_n to another random variable, say A_n, which is known to satisfy an LDP and derive from this an LDP for S_n.

We have used the first method when discussing the Gaussian, exponential and Bernoulli sample means.

The main result of large deviation theory that we will use in these notes to obtain LDPs along the indirect method is called the **Gärtner-Ellis Theorem** (GE Theorem for short), and is based on the calculation of the following function:

$$\lambda(k) = \lim_{n \to \infty} \frac{1}{n} \ln E[\mathrm{e}^{nkS_n}], \tag{2.17}$$

---

[9]To be more precise, we should make clear that the coarse-grained pdf depends on the spacing Δs by writing, say, $p_{S_n, \Delta s}(s)$. However, to keep the notation to a minimum, we will use the same lowercase p to denote a coarse-grained pdf and a "true" continuous pdf. The context should make clear which of the two is used.

known as the **scaled cumulant generating function**[10] (SCGF for short). In this expression $E[\cdot]$ denotes the expected value, k is a real parameter, and S_n is an arbitrary RV; it is not necessarily an IID sample mean or even a sample mean.

The point of the GE Theorem is to be able to calculate $\lambda(k)$ without knowing $p(S_n)$. We will see later that this is possible. Given $\lambda(k)$, the GE Theorem then says that, if $\lambda(k)$ is differentiable,[11] then

- S_n satisfies an LDP, i.e.,

$$\lim_{n\to\infty} -\frac{1}{n}\ln p_{S_n}(s) = I(s); \tag{2.18}$$

- The rate function $I(s)$ is given by the **Legendre-Fenchel transform** of $\lambda(k)$:

$$I(s) = \sup_{k\in\mathbb{R}}\{ks - \lambda(k)\}, \tag{2.19}$$

where "sup" stands for the **supremum**.[12]

We will come to calculate $\lambda(k)$ and its Legendre-Fenchel transform for specific RVs in the next section. It will also be seen in there that $\lambda(k)$ does not always exist (this typically happens when the pdf of a random variable does not admit an LDP). Some exercises at the end of this section illustrate many useful properties of $\lambda(k)$ when this function exists and is twice differentiable.

2.4 Varadhan's Theorem

The rigorous proof of the GE Theorem is too technical to be presented here. However, there is a way to justify this result by deriving in a heuristic way another result known as Varadhan's Theorem.

The latter theorem is concerned with the evaluation of a functional[13] expectation of the form

$$W_n[f] = E[\mathrm{e}^{nf(S_n)}] = \int_{\mathbb{R}} p_{S_n}(s)\,\mathrm{e}^{nf(s)}\,\mathrm{d}s, \tag{2.20}$$

---

[10]The function $E[\mathrm{e}^{kX}]$ for k real is known as the **generating function** of the RV X; $\ln E[\mathrm{e}^{kX}]$ is known as the **log-generating function** or **cumulant generating function**. The word "scaled" comes from the extra factor $1/n$.

[11]The statement of the GE Theorem given here is a simplified version of the full result, which is essentially the result proved by Gärtner [22]. See Sec. 3.3 of [48] for a more complete presentation of GE Theorem and [18] for a rigorous account of it.

[12]In these notes, a sup can be taken to mean the same as a max.

[13]A function of a function is called a **functional**.

where f is some function of S_n, which is taken to be a real RV for simplicity. Assuming that S_n satisfies an LDP with rate function $I(s)$, we can write

$$W_n[f] \approx \int_{\mathbb{R}} e^{n[f(s)-I(s)]} \, ds. \qquad (2.21)$$

with sub-exponential corrections in n. This integral has the form of a so-called **Laplace integral**, which is known to be dominated for large n by its largest integrand when it is unique. Assuming this is the case, we can proceed to approximate the whole integral as

$$W_n[f] \approx e^{n \sup_s [f(s)-I(s)]}. \qquad (2.22)$$

Such an approximation is referred to as a **Laplace approximation**, the **Laplace principle** or a **saddle-point approximation** (see Chap. 6 of [3]), and is justified in the context of large deviation theory because the corrections to this approximation are sub-exponential in n, as are those of the LDP. By defining the following functional:

$$\lambda[f] = \lim_{n \to \infty} \frac{1}{n} \ln W_n(f) \qquad (2.23)$$

using a limit similar to the limit defining the LDP, we then obtain

$$\lambda[f] = \sup_{s \in \mathbb{R}} \{f(s) - I(s)\}. \qquad (2.24)$$

The result above is what is referred to as **Varadhan Theorem** [52, 48]. The contribution of Varadhan was to prove this result for a large class of RVs, which includes not just IID sample means but also random vectors and even random functions, and to rigorously handle all the heuristic approximations used above. As such, Varadhan's Theorem can be considered as a rigorous and general expression of the Laplace principle.

To connect Varadhan's Theorem with the GE Theorem, consider the special case $f(s) = ks$ with $k \in \mathbb{R}$.[14] Then Eq. (2.24) becomes

$$\lambda(k) = \sup_{s \in \mathbb{R}} \{ks - I(s)\}, \qquad (2.25)$$

where $\lambda(k)$ is the same function as the one defined in Eq. (2.17). Thus we see that if S_n satisfies an LDP with rate function $I(s)$, then the SCGF $\lambda(k)$ of

---

[14]Varadhan's Theorem holds in its original form for bounded functions f, but another stronger version can be applied to unbounded functions and in particular to $f(x) = kx$; see Theorem 1.3.4 of [16].

S_n is the Legendre-Fenchel transform of $I(s)$. This in a sense is the inverse of the GE Theorem, so an obvious question is, can we invert the equation above to obtain $I(s)$ in terms of $\lambda(k)$? The answer provided by the GE Theorem is that, a sufficient condition for $I(s)$ to be the Legendre-Fenchel transform of $\lambda(k)$ is for the latter function to be differentiable.[15]

This is the closest that we can get to the GE Theorem without proving it. It is important to note, to fully appreciate the importance of that theorem, that it is actually more just than an inversion of Varadhan's Theorem because the existence of $\lambda(k)$ implies the existence of an LDP for S_n. In our heuristic inversion of Varadhan's Theorem we assumed that an LDP exists.

2.5 The contraction principle

We mentioned before that LDPs can be derived by a contraction method. The basis of this method is the following: let A_n be a random variable A_n known to have an LDP with rate function $I_A(a)$, and consider another random variable B_n which is a function of the first $B_n = f(A_n)$. We want to know whether B_n satisfies an LDP and, if so, to find its rate function.

To find the answer, write the pdf of B_n in terms of the pdf of A_n:

$$p_{B_n}(b) = \int_{\{a:f(a)=b\}} p_{A_n}(a)\, \mathrm{d}a \tag{2.26}$$

and use the LDP for A_n to write

$$p_{B_n}(b) \approx \int_{\{a:f(a)=b\}} \mathrm{e}^{-nI_A(a)}\, \mathrm{d}a. \tag{2.27}$$

Then apply the Laplace principle to approximate the integral above by its largest term, which corresponds to the minimum of $I(a)$ for a such that $b = f(a)$. Therefore,

$$p_{B_n}(b) \approx \exp\left(-n \inf_{\{a:f(a)=b\}} I_A(a)\right). \tag{2.28}$$

This shows that $p(B_n)$ also satisfies an LDP with rate function $I_B(b)$ given by

$$I_B(b) = \inf_{\{a:f(a)=b\}} I_A(a). \tag{2.29}$$

---

[15]The differentiability condition, though sufficient, is not necessary: $I(s)$ in some cases can be the Legendre-Fenchel transform of $\lambda(k)$ even when the latter is not differentiable; see Sec. 4.4 of [48].

This formula is called the **contraction principle** because f can be many-to-one, i.e., there might be many a's such that $b = f(a)$, in which case we are "contracting" information about the rate function of A_n down to B_n. In physical terms, this formula is interpreted by saying that an improbable fluctuation[16] of B_n is brought about by the most probable of all improbable fluctuations of A_n.

The contraction principle has many applications in statistical physics. In particular, the maximum entropy and minimum free energy principles, which are used to find equilibrium states in the microcanonical and canonical ensembles, respectively, can be seen as deriving from the contraction principle (see Sec. 5 of [48]).

2.6   From small to large deviations

An LDP for a random variable, say S_n again, gives us a lot of information about its pdf. First, we know that $p(S_n)$ concentrates on certain points corresponding to the zeros of the rate function $I(s)$. These points correspond to the most probable or **typical values** of S_n in the limit $n \to \infty$ and can be shown to be related mathematically to the **Law of Large Numbers** (LLN for short). In fact, an LDP always implies some form of LLN (see Sec. 3.5.7 of [48]).

Often it is not enough to know that S_n converges in probability to some values; we may also want to determine the likelihood that S_n takes a value away but close to its typical value(s). Consider one such typical value s^* and assume that $I(s)$ admits a Taylor series around s^*:

$$I(s) = I(s^*) + I'(s^*)(s - s^*) + \frac{I''(s^*)}{2}(s - s^*)^2 + \cdots . \qquad (2.30)$$

Since s^* must correspond to a zero of $I(s)$, the first two terms in this series vanish, and so we are left with the prediction that the **small deviations** of S_n around its typical value are Gaussian-distributed:

$$p_{S_n}(s) \approx \mathrm{e}^{-nI''(s^*)(s - s^*)^2/2}. \qquad (2.31)$$

In this sense, large deviation theory contains the **Central Limit Theorem** (see Sec. 3.5.8 of [48]). At the same time, large deviation theory can be seen as an extension of the Central Limit Theorem because it gives information not only about the small deviations of S_n, but also about its

---

[16]In statistical physics, a deviation of a random variable away from its typical value is referred to as a **fluctuation**.

large deviations far away from its typical value(s); hence the name of the theory.[17]

2.7   Exercises

2.7.1.[10] (Generating functions) Let A_n be a sum of n IID RVs:

$$A_n = \sum_{i=1}^{n} X_i. \tag{2.32}$$

Show that the generating function of A_n, defined as

$$W_{A_n}(k) = E[e^{kA_n}], \tag{2.33}$$

satisfies the following factorization property:

$$W_{A_n}(k) = \prod_{i=1}^{n} W_{X_i}(k) = W_{X_i}(k)^n, \tag{2.34}$$

$W_{X_i}(k)$ being the generating function of X_i.

2.7.2.[12] (Gaussian sample mean) Find the expression of $W_X(k)$ for X distributed according to a Gaussian pdf with mean μ and variance σ^2. From this result, find $W_{S_n}(k)$ following the previous exercise and $p(S_n)$ by inverse Laplace transform.

2.7.3.[15] (Exponential sample mean) Repeat the previous exercise for the exponential pdf shown in Eq. (2.4). Obtain from your result the approximation shown in (2.8).

2.7.4.[10] (Bernoulli sample mean) Show that the probability distribution of the Bernoulli sample mean is the binomial distribution:

$$P_{S_n}(s) = \frac{n!}{(sn)![(1-s)n]!}\, \alpha^{sn}(1-\alpha)^{(1-s)n}. \tag{2.35}$$

Use Stirling's approximation to put this result in the form of (2.9) with $I(s)$ given by Eq. (2.10).

2.7.5.[12] (Multinomial distribution) Repeat the previous exercise for IID RVs taking values in the set $\{0, 1, \ldots, q-1\}$ instead of $\{0, 1\}$. Use Stirling's approximation to arrive at an exponential approximation similar to (2.9). (See solution in [20].)

---

[17]A large deviation is also called a large fluctuation or a **rare and extreme event**.

2.7.6.[10] (SCGF at the origin) Let $\lambda(k)$ be the SCGF of an IID sample mean S_n. Prove the following properties of $\lambda(k)$ at $k = 0$:

- $\lambda(0) = 0$
- $\lambda'(0) = E[S_n] = E[X_i]$
- $\lambda''(0) = \text{var}(X_i)$

Which of these properties remain valid if S_n is not an IID sample mean, but is some general RV?

2.7.7.[12] (Convexity of SCGFs) Show that $\lambda(k)$ is a convex function of k.

2.7.8.[12] (Legendre transform) Show that, when $\lambda(k)$ is everywhere differentiable and is strictly convex (i.e., has no affine or linear parts), the Legendre-Fenchel transform shown in Eq. (2.19) reduces to the **Legendre transform**:

$$I(s) = k(s)s - \lambda(k(s)), \tag{2.36}$$

where $k(s)$ is the unique of $\lambda'(k) = s$. Explain why the latter equation has a unique root.

2.7.9.[20] (Convexity of rate functions) Prove that rate functions obtained from the GE Theorem are strictly convex.

2.7.10.[12] (Varadhan's Theorem) Verify Varadhan's Theorem for the Gaussian and exponential sample means by explicitly calculating the Legendre-Fenchel transform of $I(s)$ obtained for these sample means. Compare your results with the expressions of $\lambda(k)$ obtained from its definition.

3 Applications of large deviation theory

We study in this section examples of random variables and stochastic processes with interesting large deviation properties. We start by revisiting the simplest application of large deviation theory, namely, IID sample means, and then move to Markov processes, which are often used for modeling physical and man-made systems. More information about applications of large deviation theory in statistical physics can be found in Secs 5 and 6 of [48], as well as in the contribution of Engel contained in this volume.

3.1 IID sample means

Consider again the sample mean

$$S_n = \frac{1}{n} \sum_{i=1}^{n} X_i \tag{3.1}$$

involving n real IID variables $X_1, \ldots, X_n$ with common pdf $p(x)$. To determine whether S_n satisfies an LDP and obtain its rate function, we can follow the GE Theorem and calculate the SCGF $\lambda(k)$ defined in Eq. (2.17). Because of the IID nature of S_n, $\lambda(k)$ takes a simple form:

$$\lambda(k) = \ln E[e^{kX}]. \tag{3.2}$$

In this expression, X is any of the X_i's (remember they are IID). Thus to obtain an LDP for S_n, we only have to calculate the log-generating function of X, check that it is differentiable and, if so, calculate its Legendre-Fenchel transform (or in this case, its Legendre transform; see Exercise 2.7.8). Exercise 3.6.1 considers different distributions $p(x)$ for which this calculation can be carried out, including the Gaussian, exponential and Bernoulli distributions studied before.

As an extra exercise, let us attempt to find the rate function of a special sample mean defined by

$$L_{n,j} = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_i, j} \tag{3.3}$$

for a sequence $X_1, \ldots, X_n$ of n discrete IID RVs with finite state space $\mathcal{X} = \{1, 2, \ldots, q\}$. This sample mean is called the **empirical distribution** of the X_i's as it "counts" the number of times the value or symbol $j \in \mathcal{X}$ appears in a given realization of $X_1, \ldots, X_n$. This number is normalized by the total number n of RVs, so what we have is the **empirical frequency** for the appearance of the symbol j in realizations of $X_1, \ldots, X_n$.

The values of $L_{n,j}$ for all $j \in \mathcal{X}$ can be put into a vector $\mathbf{L}_n$ called the **empirical vector**. For the Bernoulli sample mean, for example, $\mathcal{X} = \{0, 1\}$ and so $\mathbf{L}_n$ is a two-dimensional vector $\mathbf{L}_n = (L_{n,0}, L_{n,1})$ containing the empirical frequency of 0's and 1's appearing in a realization $x_1, \ldots, x_n$ of $X_1, \ldots, X_n$:

$$L_{n,0} = \frac{\# \ 0\text{'s in } x_1, \ldots, x_n}{n}, \quad L_{n,1} = \frac{\# \ 1\text{'s in } x_1, \ldots, x_n}{n}. \tag{3.4}$$

To find the rate function associated with the random vector $\mathbf{L}_n$, we apply the GE Theorem but adapt it to the case of random vectors by replacing k

in $\lambda(k)$ by a vector $\mathbf{k}$ having the same dimension as $\mathbf{L}_n$. The calculation of $\lambda(\mathbf{k})$ is left as an exercise (see Exercise 3.6.8). The result is

$$\lambda(\mathbf{k}) = \ln \sum_{j \in \mathcal{X}} P_j \, \mathrm{e}^{k_j} \tag{3.5}$$

where $P_j = P(X_i = j)$. It is easily checked that $\lambda(\mathbf{k})$ is differentiable in the vector sense, so we can use the GE Theorem to conclude that $\mathbf{L}_n$ satisfies an LDP, which we write as

$$p(\mathbf{L}_n = \boldsymbol{\mu}) \approx \mathrm{e}^{-nI(\boldsymbol{\mu})}, \tag{3.6}$$

with a vector rate function $I(\boldsymbol{\mu})$ given by the Legendre-Fenchel transform of $\lambda(\mathbf{k})$. The calculation of this transform is the subject of Exercise 3.6.8. The final result is

$$I(\boldsymbol{\mu}) = \sum_{j \in \mathcal{X}} \mu_j \ln \frac{\mu_j}{P_j}. \tag{3.7}$$

This rate function is called the **relative entropy** or **Kullback-Leibler divergence** [8]. The full LDP for $\mathbf{L}_n$ is referred to as **Sanov's Theorem** (see [44] and Sec. 4.2 of [48]).

It can be checked that $I(\boldsymbol{\mu})$ is a convex function of $\boldsymbol{\mu}$ and that it has a unique minimum and zero at $\boldsymbol{\mu} = \mathbf{P}$, i.e., $\mu_j = P_j$ for all $j \in \mathcal{X}$. As seen before, this property is an expression of the LLN, which says here that the empirical frequencies $L_{n,j}$ converge to the probabilities P_j as $n \to \infty$. The LDP goes beyond this result by describing the fluctuations of $\mathbf{L}_n$ around the typical value $\boldsymbol{\mu} = \mathbf{P}$.

3.2 Markov chains

Instead of assuming that the sample mean S_n arises from a sum of IID RVs $X_1, \ldots, X_n$, consider the case where the X_i's form a **Markov chain**. This means that the joint pdf $p(X_1, \ldots, X_n)$ has the form

$$p(X_1, \ldots, X_n) = p(X_1) \prod_{i=1}^{n-1} \pi(X_{i+1}|X_i), \tag{3.8}$$

where $p(X_i)$ is some initial pdf for X_1 and $\pi(X_{i+1}|X_i)$ is the **transition probability density** that X_{i+1} follows X_i in the Markov sequence $X_1 \to X_2 \to \cdots \to X_n$ (see [27] for background information on Markov chains).

The GE Theorem can still be applied in this case, but the expression of $\lambda(k)$ is more complicated. Skipping the details of the calculation (see Sec. 4.3

18

of [48]), we arrive at the following result: provided that the Markov chain is homogeneous and ergodic (see [27] for a definition of ergodicity), the SCGF of S_n is given by

$$\lambda(k) = \ln \zeta(\tilde{\Pi}_k), \tag{3.9}$$

where $\zeta(\tilde{\Pi}_k)$ is the **dominant eigenvalue** (i.e., with largest magnitude) of the matrix $\tilde{\Pi}_k$ whose elements $\tilde{\pi}_k(x, x')$ are given by $\tilde{\pi}_k(x, x') = \pi(x'|x)e^{kx'}$. We call the matrix $\tilde{\Pi}_k$ the **tilted matrix** associated with S_n. If the Markov chain is finite, it can be proved furthermore that $\lambda(k)$ is analytic and so differentiable. From the GE Theorem, we then conclude that S_n has an LDP with rate function

$$I(s) = \sup_{k \in \mathbb{R}} \{ks - \ln \zeta(\tilde{\Pi}_k)\}. \tag{3.10}$$

Some care must be taken if the Markov chain has an infinite number of states. In this case, $\lambda(k)$ is not necessarily analytic and may not even exist (see, e.g., [29, 30, 38] for illustrations).

An application of the above result for a simple Bernoulli Markov chain is considered in Exercise 3.6.10. In Exercises 3.6.12 and 3.6.13, the case of random variables having the form

$$Q_n = \frac{1}{n} \sum_{i=1}^{n-1} q(x_i, x_{i+1}) \tag{3.11}$$

is covered. This type of RV arises in statistical physics when studying currents in stochastic models of interacting particles (see [30]). The tilted matrix $\tilde{\Pi}_k$ associated with Q_n has the form $\tilde{\pi}_k(x, x') = \pi(x'|x)e^{kq(x,x')}$.

3.3 Continuous-time Markov processes

The preceding subsection can be generalized to ergodic Markov processes evolving continuously in time by taking a continuous time limit of discrete-time ergodic Markov chains.

To illustrate this limit, consider an ergodic continuous-time process described by the state $X(t)$ for $0 \leq t \leq T$. For an infinitesimal **time-step** Δt, time-discretize this process into a sequence of $n+1$ states $X_0, X_1, \ldots, X_n$ with $n = T/\Delta t$ and $X_i = X(i\Delta t)$, $i = 0, 1, \ldots, n$. The sequence $X_0, X_1, \ldots, X_n$ is an ergodic Markov chain with infinitesimal transition matrix $\Pi(\Delta t)$ given by

$$\Pi(\Delta t) = \pi(x(t + \Delta t)|x(t)) = e^{G\Delta t} = I + G\Delta t + o(\Delta t), \tag{3.12}$$

where I is the identity matrix and G the **generator** of $X(t)$. With this discretization, it is now possible to study LDPs for processes involving $X(t)$

by studying these processes in discrete time at the level of the Markov chain $X_0, X_1, \ldots, X_n$, and transfer these LDPs into continuous time by taking the limits $\Delta t \to 0$, $n \to \infty$.

As a general application of this procedure, consider a so-called **additive process** defined by

$$S_T = \frac{1}{T} \int_0^T X(t) \, dt. \tag{3.13}$$

The discrete-time version of this RV is the sample mean

$$S_n = \frac{1}{n\Delta t} \sum_{i=0}^n X_i \, \Delta t = \frac{1}{n} \sum_{i=0}^n X_i. \tag{3.14}$$

From this association, we find that the SCGF of S_T, defined by the limit

$$\lambda(k) = \lim_{T \to \infty} \frac{1}{T} \ln E[e^{TkS_T}], \qquad k \in \mathbb{R}, \tag{3.15}$$

is given by

$$\lambda(k) = \lim_{\Delta t \to 0} \lim_{n \to \infty} \frac{1}{n\Delta t} \ln E[e^{k\Delta t \sum_{i=0}^n X_i}] \tag{3.16}$$

at the level of the discrete-time Markov chain.

According to the previous subsection, the limit above reduces to $\ln \zeta(\tilde{\Pi}_k(\Delta t))$, where $\tilde{\Pi}_k(\Delta t)$ is the matrix (or operator) corresponding to

$$
\begin{aligned}
\tilde{\Pi}_k(\Delta t) &= e^{k\Delta t x'} \Pi(\Delta t) \\
&= \left(1 + kx'\Delta t + o(\Delta t)\right)\left(I + G'\Delta t + o(\Delta t)\right) \\
&= I + \tilde{G}_k\Delta t + o(\Delta t), \tag{3.17}
\end{aligned}
$$

where

$$\tilde{G}_k = G + kx'\delta_{x,x'}. \tag{3.18}$$

For the continuous-time process $X(t)$, we must therefore have

$$\lambda(k) = \zeta(\tilde{G}_k), \tag{3.19}$$

where $\zeta(\tilde{G}_k)$ is the dominant eigenvalue of the **tilted generator** $\tilde{G}_k$.[18] Note that the reason why the logarithmic does not appear in the expression of $\lambda(k)$ for the continuous-time process[19] is because ζ is now the dominant

---

[18] As for Markov chains, this result holds for continuous, ergodic processes with finite state-space. For infinite-state or continuous-space processes, a similar result holds provided G_k has an isolated dominant eigenvalue.

[19] Compare with Eq. (3.9).

eigenvalue of the generator of the infinitesimal transition matrix, which is itself the exponential of the generator.

With the knowledge of $\lambda(k)$, we can obtain an LDP for S_T using the GE Theorem: If the dominant eigenvalue is differentiable in k, then S_T satisfies an LDP in the long-time limit, $T \to \infty$, which we write as $p_{S_T}(s) \approx e^{-TI(s)}$, where $I(s)$ is the Legendre-Fenchel transform $\lambda(k)$. Some applications of this result are presented in Exercises 3.6.11 and 3.6.13. Note that in the case of current-type processes having the form of Eq. (3.11), the tilted generator is not simply given by $\tilde{G}_k = G + kx'\delta_{x,x'}$; see Exercise 3.6.13.

3.4 Paths large deviations

We complete our tour of mathematical applications of large deviation theory by studying a different large deviation limit, namely, the low-noise limit of the following **stochastic differential equation** (SDE for short):

$$\dot{x}(t) = f(x(t)) + \sqrt{\varepsilon}\,\xi(t), \qquad x(0) = 0, \tag{3.20}$$

which involves a **force** $f(x)$ and a **Gaussian white noise** $\xi(t)$ with the properties $E[\xi(t)] = 0$ and $E[\xi(t')\xi(t)] = \delta(t' - t)$; see Chap. XVI of [49] for background information on SDEs.

We are interested in studying for this process the pdf of a given **random path** $\{x(t)\}_{t=0}^T$ of duration T in the limit where the **noise power** ε vanishes. This abstract pdf can be defined heuristically using path integral methods (see Sec. 6.1 of [48]). In the following, we denote this pdf by the functional notation $p[x]$ as a shorthand for $p(\{x(t)\}_{t=0}^T)$.

The idea behind seeing the low-noise limit as a large deviation limit is that, as $\varepsilon \to 0$, the random path arising from the SDE above should converge in probability to the **deterministic path** $x(t)$ solving the ordinary differential equation

$$\dot{x}(t) = f(x(t)), \qquad x(0) = 0. \tag{3.21}$$

This convergence is a LLN-type result, and so in the spirit of large deviation theory we are interested in quantifying the likelihood that a random path $\{x(t)\}_{t=0}^T$ ventures away from the deterministic path in the limit $\varepsilon \to 0$. The functional LDP that characterizes these path fluctuations has the form

$$p[x] \approx e^{-I[x]/\varepsilon}, \tag{3.22}$$

where

$$I[x] = \int_0^T [\dot{x}(t) - f(x(t))]^2 \, dt. \tag{3.23}$$

See [21] and Sec. 6.1 of [48] for historical sources on this LDP.

The rate functional $I[x]$ is called the **action**, **Lagrangian** or **entropy** of the path $\{x(t)\}_{t=0}^T$. The names "action" and "Lagrangian" come from an analogy with the action of quantum trajectories in the path integral approach of quantum mechanics (see Sec. 6.1 of [48]). There is also a close analogy between the low-noise limit of SDEs and the semi-classical or WKB approximation of quantum mechanics [48].

The path LDP above can be generalized to higher-dimensional SDEs as well as SDEs involving state-dependent noise and correlated noises (see Sec. 6.1 of [48]). In all cases, the minimum and zero of the rate functional is the trajectory of the deterministic system obtained in the zero-noise limit. This is verified for the 1D system considered above: $I[x] \geq 0$ for all trajectories and $I[x] = 0$ for the unique trajectory solving Eq. (3.21).

Functional LDPs are the most refined LDPs that can be derived for SDEs as they characterize the probability of complete trajectories. Other "coarser" LDPs can be derived from these by contraction. For example, we might be interested to determine the pdf $p(x, T)$ of the state $x(T)$ reached after a time T. The contraction in this case is obvious: $p(x, T)$ must have the large deviation form $p(x, T) \approx e^{-V(x,T)/\varepsilon}$ with

$$V(x, T) = \inf_{x(t): x(0)=0, x(T)=x} I[x]. \tag{3.24}$$

That is, the probability of reaching $x(T) = x$ from $x(0) = 0$ is determined by the path connecting these two endpoints having the largest probability. We call this path the **optimal path**, **maximum likelihood path** or **instanton**. Using variational calculus techniques, often used in classical mechanics, it can be proved that this path satisfies an Euler-Lagrange-type equation as well as a Hamilton-type equation (see Sec. 6.1 of [48]). Applications of these equations are covered in Exercises 3.6.14–3.6.17.

Quantities similar to the additive process S_T considered in the previous subsection can also be defined for SDEs (see Exercises 3.6.11 and 3.6.12). An interesting aspect of these quantities is that their LDPs can involve the noise power ε if the low-noise limit is taken, as well as the integration time T, which arises because of the additive (in the sense of sample mean) nature of these quantities. In this case, the limit $T \to 0$ must be taken before the low-noise limit $\varepsilon \to 0$. If $\varepsilon \to 0$ is taken first, the system considered is no longer random, which means that there can be no LDP in time.

3.5 Physical applications

Some applications of the large deviation results seen so far are covered in the contribution of Engel to this volume. The following list gives an indication of these applications and some key references for learning more about them. A more complete presentation of the applications of large deviation theory in statistical physics can be found in Secs. 5 and 6 of [48].

- **Equilibrium systems:** Equilibrium statistical mechanics, as embodied by the ensemble theory of Boltzmann and Gibbs, can be seen with hindsight as a large deviation theory of many-body systems at equilibrium. This becomes evident by realizing that the thermodynamic limit is a large deviation limit, that the entropy is the equivalent of a rate function and the free energy the equivalent of a SCGF. Moreover, the Legendre transform of thermodynamics connecting the entropy and free energy is nothing but the Legendre-Fenchel transform connecting the rate function and the SCGF in the GE Theorem and in Varadhan's Theorem. For a more complete explanation of these analogies, and historical sources on the development of large deviation theory in relation to equilibrium statistical mechanics, see the book by Ellis [19] and Sec. 3.7 of [48].

- **Chaotic systems and multifractals:** The so-called **thermodynamic formalism** of dynamical systems, developed by Ruelle [40] and others, can also be re-interpreted with hindsight as an application of large deviation theory for the study of chaotic systems.[20] There are two quantities in this theory playing the role of the SCGF, namely, the **topological pressure** and the **structure function**. The Legendre transform appearing in this theory is also an analogue of the one encountered in large deviation theory. References to these analogies can be found in Secs. 7.1 and 7.2 of [48].

- **Nonequilibrium systems:** Large deviation theory is becoming the standard formalism used in studies of nonequilibrium systems modelled by SDEs and Markov processes in general. In fact, large deviation theory is currently experiencing a sort of revival in physics and mathematics as a result of its growing application in nonequilibrium statistical mechanics. Many LDPs have been derived in this context: LDPs for the current fluctuations or the occupation of interacting particle models,

---

[20]There is a reference to this connection in a note of Ruelle's popular book, Chance and Chaos [39] (see Note 2 of Chap. 19).

such as the exclusion process, the zero-range process (see [28]) and their many variants [45], as well as LDPs for work-related and entropy-production-related quantities for nonequilibrium systems modelled with SDEs [6, 7]. Good entry points in this vast field of research are [13] and [30].

- **Fluctuation relations:** Several LDPs have come to be studied in recent years under the name of fluctuation relations. To illustrate these results, consider the additive process S_T considered earlier and assume that this process admits an LDP with rate function $I(s)$. In many cases, it is interesting not only to know that S_T admits an LDP but to know how probable positive fluctuations of S_T are compared to negative fluctuations. For this purpose, it is common to study the ratio

$$\frac{p_{S_T}(s)}{p_{S_T}(-s)}, \tag{3.25}$$

which reduces to

$$\frac{p_{S_T}(s)}{p_{S_T}(-s)} \approx e^{T[I(-s)-I(s)]} \tag{3.26}$$

if we assume an LDP for S_T. In many cases, the difference $I(-s) - I(s)$ is linear in s, and one then says that S_T satisfies a **conventional fluctuation relation**, whereas if it nonlinear in s, then one says that S_T satisfies an **extended fluctuation relation**. Other types of fluctuation relations have come to be defined in addition to these; for more information, the reader is referred to Sec. 6.3 of [48]. A list of the many physical systems for which fluctuation relations have been derived or observed can also be found in this reference.

3.6 Exercises

3.6.1.[12] (IID sample means) Use the GE Theorem to find the rate function of the IID sample mean S_n for the following probability distribution and densities:

- Gaussian:
$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/(2\sigma^2)}, \qquad x \in \mathbb{R}. \tag{3.27}$$

- Bernoulli: $X_i \in \{0, 1\}$, $P(X_i = 0) = 1 - \alpha$, $P(X_i = 1) = \alpha$.
- Exponential:
$$p(x) = \frac{1}{\mu} e^{-x/\mu}, \qquad x \in [0, \infty), \quad \mu > 0. \tag{3.28}$$

24

- Uniform:

$$p(x) = \begin{cases} \frac{1}{a} & x \in [0, a] \\ 0 & \text{otherwise.} \end{cases} \tag{3.29}$$

- Cauchy:

$$p(x) = \frac{\sigma}{\pi(x^2 + \sigma^2)}, \qquad x \in \mathbb{R}, \quad \sigma > 0. \tag{3.30}$$

3.6.2.[15] (Nonconvex rate functions) Rate functions obtained from the GE Theorem are necessarily convex (strictly convex, in fact; see Exercise 2.7.9), but rate functions in general need not be convex. Consider, as an example, the process defined by

$$S_n = Y + \frac{1}{n} \sum_{i=1}^{n} X_i, \tag{3.31}$$

where $Y = \pm 1$ with probability $\frac{1}{2}$ and the X_i's are Gaussian IID RVs. Find the rate function for S_n assuming that Y is independent of the X_i's. Then find the corresponding SCGF. What is the relation between the Legendre-Fenchel transform of $\lambda(k)$ and the rate function $I(s)$? How is the nonconvexity of $I(s)$ related to the differentiability of $\lambda(k)$? (See Example 4.7 of [48] for the solution.)

3.6.3.[15] (Exponential mixture) Repeat the previous exercise by replacing the Bernoulli Y with $Y = Z/n$, where Z is an exponential RV with mean 1. (See Example 4.8 of [48] for the solution.)

3.6.4.[12] (Product process) Find the rate function of

$$Z_n = \left(\prod_{i=1}^{n} X_i \right)^{1/n} \tag{3.32}$$

where $X_i \in \{1, 2\}$ with $p(X_i = 1) = \alpha$ and $p(X_i = 2) = 1 - \alpha$, $0 < \alpha < 1$.

3.6.5.[15] (Self-process) Consider a sequence of IID Gaussian RVs $X_1, \ldots, X_n$. Find the rate function of the so-called **self-process** defined by

$$S_n = -\frac{1}{n} \ln p(X_1, \ldots, X_n). \tag{3.33}$$

(See Sec. 4.5 of [48] for information about this process.) Repeat the problem for other choices of pdf for the RVs.

3.6.6.[20] (Iterated large deviations [35]) Consider the sample mean

$$S_{m,n} = \frac{1}{m} \sum_{i=1}^{m} X_i^{(n)} \tag{3.34}$$

involving m IID copies of a random variable $X^{(n)}$. Show that if $X^{(n)}$ satisfies an LDP of the form

$$p_{X^{(n)}}(x) \approx e^{-nI(x)}, \tag{3.35}$$

then $S_{m,n}$ satisfies an LDP having the form

$$p_{S_{m,n}}(s) \approx e^{-mnI(x)}. \tag{3.36}$$

3.6.7.[40] (Persistent random walk) Use the result of the previous exercise to find the rate function of

$$S_n = \frac{1}{n} \sum_{i=1}^{n} X_i, \tag{3.37}$$

where the X_i's are independent Bernoulli RVs with non-identical distribution $P(X_i = 0) = \alpha^i$ and $P(X_i = 1) = 1 - \alpha^i$ with $0 < \alpha < 1$.

3.6.8.[15] (Sanov's Theorem) Consider the empirical vector $\mathbf{L}_n$ with components defined in Eq. (3.3). Derive the expression found in Eq. (3.5) for

$$\lambda(\mathbf{k}) = \lim_{n \to \infty} \frac{1}{n} \ln E[e^{n \mathbf{k} \cdot \mathbf{L}_n}], \tag{3.38}$$

where $\mathbf{k} \cdot \mathbf{L}_n$ denotes the scalar product of $\mathbf{k}$ and $\mathbf{L}_n$. Then obtain the rate function found in (3.7) by calculating the Legendre transform of $\lambda(\mathbf{k})$. Check explicitly that the rate function is convex and has a single minimum and zero.

3.6.9.[20] (Contraction principle) Repeat the first exercise of this section by using the contraction principle. That is, use the rate function of the empirical vector $\mathbf{L}_n$ to obtain the rate function of S_n. What is the mapping from $\mathbf{L}_n$ to S_n? Is this mapping many-to-one or one-to-one?

3.6.10.[17] (Bernoulli Markov chain) Consider the Bernoulli sample mean

$$S_n = \frac{1}{n} \sum_{i=1}^{n} X_i, \qquad X_i \in \{0, 1\}. \tag{3.39}$$

Find the expression of $\lambda(k)$ and $I(s)$ for this process assuming that the X_i's form a Markov process with symmetric transition matrix

$$\pi(x'|x) = \begin{cases} 1 - \alpha & x' = x \\ \alpha & x' \neq x \end{cases} \tag{3.40}$$

with $0 \leq \alpha \leq 1$. (See Example 4.4 of [48] for the solution.)

3.6.11.[20] (Residence time) Consider a Markov process with state space $\mathcal{X} = \{0, 1\}$ and generator

$$G = \begin{pmatrix} -\alpha & \beta \\ \alpha & -\beta \end{pmatrix}. \tag{3.41}$$

Find for this process the rate function of the random variable

$$L_T = \frac{1}{T} \int_0^T \delta_{X(t),0} \, dt, \tag{3.42}$$

which represents the fraction of time the state $X(t)$ of the Markov chain spends in the state $X = 0$ over a period of time T.

3.6.12.[20] (Current fluctuations in discrete time) Consider the Markov chain of Exercise 3.6.10. Find the rate function of the **mean current** Q_n defined as

$$Q_n = \frac{1}{n} \sum_{i=1}^{n} f(x_{i+1}, x_i), \tag{3.43}$$

$$f(x', x) = \begin{cases} 0 & x' = x \\ 1 & x' \neq x. \end{cases} \tag{3.44}$$

Q_n represents the mean number of jumps between the states 0 and 1: at each transition of the Markov chain, the current is incremented by 1 whenever a jump between the states 0 and 1 occurs.

3.6.13.[22] (Current fluctuations in continuous time) Repeat the previous exercise for the Markov process of Exercise 3.6.11 and the current Q_T defined as

$$Q_T = \lim_{\Delta t \to 0} \frac{1}{T} \sum_{i=0}^{n-1} f(x_{i+1}, x_i), \tag{3.45}$$

where $x_i = x(i\Delta t)$ ad $f(x', x)$ as above, i.e., $f(x_{i+1}, x_i) = 1 - \delta_{x_i, x_{i+1}}$. Show for this process that the tilted generator is

$$\tilde{G}_k = \begin{pmatrix} -\alpha & \beta \, e^k \\ \alpha \, e^k & -\beta \end{pmatrix}. \tag{3.46}$$

3.6.14.[17] (Ornstein-Uhlenbeck process) Consider the following linear SDE:

$$\dot{x}(t) = -\gamma x(t) + \sqrt{\varepsilon}\,\xi(t), \tag{3.47}$$

often referred to as the **Langevin equation** or **Ornstein-Uhlenbeck process**. Find for this SDE the solution of the optimal path connecting the initial point $x(0) = 0$ with the fluctuation $x(T) = x$. From the solution, find the pdf $p(x, T)$ as well as the stationary pdf obtained in the limit $T \to \infty$. Assume $\varepsilon \to 0$ throughout but then verify that the results obtained are valid for all $\varepsilon > 0$.

3.6.15.[20] (Conservative system) Show by large deviation techniques that the stationary pdf of the SDE

$$\dot{\mathbf{x}}(t) = -\nabla U(\mathbf{x}(t)) + \sqrt{\varepsilon}\,\xi(t) \tag{3.48}$$

admits the LDP $p(\mathbf{x}) \approx \mathrm{e}^{-V(\mathbf{x})/\varepsilon}$ with rate function $V(\mathbf{x}) = 2U(\mathbf{x})$. The rate function $V(\mathbf{x})$ is called the **quasi-potential**.

3.6.16.[25] (Transversal force) Show that the LDP found in the previous exercise also holds for the SDE

$$\dot{\mathbf{x}}(t) = -\nabla U(\mathbf{x}(t)) + \mathbf{A}(\mathbf{x}) + \sqrt{\varepsilon}\,\xi(t) \tag{3.49}$$

if $\nabla U \cdot \mathbf{A} = 0$. (See Sec. 4.3 of [21] for the solution.)

3.6.17.[25] (Noisy Van der Pol oscillator) The following set of SDEs describes a noisy version of the well-known Van der Pol equation:

$$\begin{aligned} \dot{x} &= v \\ \dot{v} &= -x + v(\alpha - x^2 - v^2) + \sqrt{\varepsilon}\,\xi. \end{aligned} \tag{3.50}$$

The parameter $\alpha \in \mathbb{R}$ controls a bifurcation of the zero-noise system: for $\alpha \leq 0$, the system with $\xi = 0$ has a single attracting point at the origin, whereas for $\alpha > 0$, it has a stable limit cycle centered on the origin. Show that the stationary distribution of the noisy oscillator has the large deviation form $p(x, v) \approx \mathrm{e}^{-U(x,v)/\varepsilon}$ as $\varepsilon \to 0$ with rate function

$$U(x, v) = -\alpha(x^2 + v^2) + \frac{1}{2}(x^2 + v^2)^2. \tag{3.51}$$

Find a different set of SDEs that has the same rate function. (See [26] for the solution.)

3.6.18.[30] (Dragged Brownian particle [50]) The stochastic dynamics of a tiny glass bead immersed in water and pulled with laser tweezers at constant velocity can be modelled in the overdamped limit by the following reduced Langevin equation:

$$\dot{x} = -(x(t) - \nu t) + \sqrt{\varepsilon}\,\xi(t), \tag{3.52}$$

where $x(t)$ is the position of the glass bead, ν the pulling velocity, $\xi(t)$ a Gaussian white noise modeling the influence of the surrounding fluid, and ε the noise power related to the temperature of the fluid. Show that the **mean work** done by the laser as it pulls the glass bead over a time T, which is defined as

$$W_T = -\frac{\nu}{T} \int_0^T (x(t) - \nu t)\,\mathrm{d}t, \tag{3.53}$$

satisfies an LDP in the limit $T \to \infty$. Derive this LDP by assuming first that $\varepsilon \to 0$, and then derive the LDP without this assumption. Finally, determine whether W_T satisfies a conventional or extended fluctuation relation, and study the limit $\nu \to 0$. (See Example 6.9 of [48] for more references on this model.)

4 Numerical estimation of large deviation probabilities

The previous sections might give the false impression that rate functions can be calculated explicitly for many stochastic processes. In fact, exact and explicit expressions of rate functions can be found only in few simple cases. For most stochastic processes of scientific interest (e.g,, noisy nonlinear dynamical systems, chemical reactions, queues, etc.), we have to rely on analytical approximations and numerical methods to evaluate rate functions.

The rest of these notes attempts to give an overview of some of these numerical methods and to illustrate them with the simple processes (IID sample means and Markov processes) treated before. The goal is to give a flavor of the general ideas behind large deviation simulations rather than a complete survey, so only a subset of the many methods that have come to be devised for numerically estimating rate functions are covered in what follows. Other methods not treated here, such as the transition path method discussed by Dellago in this volume, are mentioned at the end of the next section.

4.1 Direct sampling

The problem addressed in this section is to obtain a numerical estimate of the pdf $p_{S_n}(s)$ for a real random variable S_n satisfying an LDP, and to extract from this an estimate of the rate function $I(s)$.[21] To be general, we take S_n to be a function of n RVs $X_1, \ldots, X_n$, which at this point are not necessarily IID. To simplify the notation, we will use the shorthand $\boldsymbol{\omega} = X_1, \ldots, X_n$. Thus we write S_n as the function $S_n(\boldsymbol{\omega})$ and denote by $p(\boldsymbol{\omega})$ the joint pdf of the X_i's.[22]

Numerically, we cannot of course obtain $p_{S_n}(s)$ or $I(s)$ for all $s \in \mathbb{R}$, but only for a finite number of values s, which we take for simplicity to be equally spaced with a small step Δs. Following our discussion of the LDP, we thus attempt to estimate the coarse-grained pdf

$$p_{S_n}(s) = \frac{P(S_n \in [s, s + \Delta s])}{\Delta s} = \frac{P(S_n \in \Delta_s)}{\Delta s}, \qquad (4.1)$$

where Δ_s denotes the small interval $[s, s + \Delta s]$ anchored at the value s.[23]

To construct this estimate, we follow the **statistical sampling** or **Monte Carlo method**, which we broke down into the following steps (see the contribution of Katzgraber in this volume for more details):

1. Generate a **sample** $\{\boldsymbol{\omega}^{(j)}\}_{j=1}^{L}$ of L copies or realizations of the sequence $\boldsymbol{\omega}$ from its pdf $p(\boldsymbol{\omega})$.

2. Obtain from this sample a sample $\{s^{(j)}\}_{j=1}^{L}$ of values or realizations for S_n:
$$s^{(j)} = S_n(\boldsymbol{\omega}^{(j)}), \qquad j = 1, \ldots, L. \qquad (4.2)$$

3. Estimate $P(S_n \in \Delta_s)$ by calculating the sample mean

$$\hat{P}_L(\Delta_s) = \frac{1}{L} \sum_{j=1}^{L} \mathbf{1}_{\Delta_s}(s^{(j)}), \qquad (4.3)$$

where $\mathbf{1}_A(x)$ denotes the **indicator function** for the set A, which is equal to 1 if $x \in A$ and 0 otherwise.

---

[21] The literature on large deviation simulation usually considers the estimation of $P(S_n \in A)$ for some set A rather than the estimation of the whole pdf of S_n.

[22] For simplicity, we consider the X_i's to be real RVs. The case of discrete RVs follow with slight changes of notations.

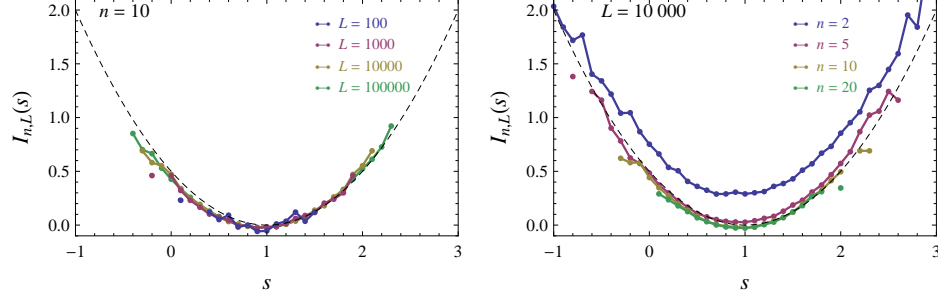[23] Though not explicitly noted, the coarse-grained pdf depends on Δs; see Footnote 9.

Figure 4.1: (Left) Naive sampling of the Gaussian IID sample mean ($\mu = \sigma = 1$) for $n = 10$ and different sample sizes L. The dashed line is the exact rate function. As L grows, a larger range of $I(s)$ is sampled. (Right) Naive sampling for a fixed sample size $L = 10\,000$ and various values of n. As n increases, $I_{n,L}(s)$ approaches the expected rate function but the sampling becomes inefficient as it becomes restricted to a narrow domain.

4. Turn the **estimator** $\hat{P}_L(\Delta_s)$ of the probability $P(S_n \in \Delta_s)$ into an estimator $\hat{p}_L(s)$ of the probability density $p_{S_n}(s)$:

$$\hat{p}_L(s) = \frac{\hat{P}_L(\Delta_s)}{\Delta s} = \frac{1}{L\Delta s} \sum_{j=1}^{L} \mathbf{1}_{\Delta_s}(s^{(j)}). \tag{4.4}$$

The result of these steps is illustrated in Fig. 4.1 for the case of the IID Gaussian sample mean. Note that $\hat{p}_L(s)$ above is nothing but an empirical vector for S_n (see Sec. 3.1) or a **density histogram** of the sample $\{s^{(j)}\}_{j=1}^{L}$.

The reason for choosing $\hat{p}(s)$ as our estimator of $p_{S_n}(s)$ is that it is an **unbiased estimator** in the sense that

$$E[\hat{p}_L(s)] = p_{S_n}(s) \tag{4.5}$$

for all L. Moreover, we know from the LLN that $\hat{p}_L(s)$ converges in probability to its mean $p_{S_n}(s)$ as $L \to \infty$. Therefore, the larger our sample, the closer we should get to a valid estimation of $p_{S_n}(s)$.

To extract a rate function $I(s)$ from $\hat{p}_L(s)$, we simply compute

$$I_{n,L}(s) = -\frac{1}{n} \ln \hat{p}_L(s) \tag{4.6}$$

and repeat the whole process for larger and larger integer values of n and L until $I_{n,L}(s)$ converges to some desired level of accuracy.[24]

---

[24]Note that $\hat{P}_L(\Delta_s)$ and $\hat{p}_L(s)$ differ only by the factor Δs. $I_{n,L}(s)$ can therefore be

4.2   Importance sampling

A basic rule of thumb in statistical sampling, suggested by the LLN, is that an event with probability P will appear in a sample of size L roughly LP times. Thus to get at least one instance of that event in the sample, we must have $L > 1/P$ as an approximate lower bound for the size of our sample (see Exercise 4.7.5 for a more precise derivation of this estimate).

Applying this result to $p_{S_n}(s)$, we see that, if this pdf satisfies an LDP of the form $p_{S_n}(s) \approx \mathrm{e}^{-nI(s)}$, then we need to have $L > \mathrm{e}^{nI(s)}$ to get at least one instance of the event $S_n \in \Delta_s$ in our sample. In other words, our sample must be exponentially large with n in order to see any large deviations (see Fig. 4.1 and Exercise 4.7.2).

This is a severe limitation of the sampling scheme outlined earlier, and we call it for this reason **crude Monte Carlo** or **naive sampling**. The way around this limitation is to use **importance sampling** (IS for short) which works basically as follows (see the contribution of Katzgraber for more details):

1. Instead of sampling the X_i's according to the joint pdf $p(\boldsymbol{\omega})$, sample them according to a new pdf $q(\boldsymbol{\omega})$;[25]

2. Calculate instead of $\hat{p}_L(s)$ the estimator

$$\hat{q}_L(s) = \frac{1}{L\Delta s} \sum_{j=1}^{L} \mathbf{1}_{\Delta_s}\big(S_n(\boldsymbol{\omega}^{(j)})\big)\, R(\boldsymbol{\omega}^{(j)}), \qquad (4.7)$$

where

$$R(\boldsymbol{\omega}) = \frac{p(\boldsymbol{\omega})}{q(\boldsymbol{\omega})} \qquad (4.8)$$

is the called the **likelihood ratio**. Mathematically, R also corresponds to the **Radon-Nikodym derivative** of the measures associated with p and q.[26]

The new estimator $\hat{q}_L(s)$ is also an unbiased estimator of $p_{S_n}(s)$ because (see Exercise 4.7.3)

$$E_q[\hat{q}_L(s)] = E_p[\hat{p}_L(s)]. \qquad (4.9)$$

---

computed from either estimator with a difference $(\ln \Delta s)/n$ that vanishes as $n \to \infty$.

[25]The pdf q must have a support at least as large as that of p, i.e., $q(\boldsymbol{\omega}) > 0$ if $p(\boldsymbol{\omega}) > 0$, otherwise the ratio R_n is ill-defined. In this case, we say that q is **relatively continuous** with respect to p and write $q \gg p$.

[26]The Radon-Nikodym derivative of two measures μ and ν such that $\nu \gg \mu$ is defined as $d\mu/d\nu$. If these measures have densities p and q, respectively, then $d\mu/d\nu = p/q$.

However, there is a reason for choosing $\hat{q}_L(s)$ as our new estimator: we might be able to come up with a suitable choice for the pdf q such that $\hat{q}_L(s)$ has a smaller variance than $\hat{p}_L(s)$. This is in fact the goal of IS: select $q(\boldsymbol{\omega})$ so as to minimize the variance

$$\text{var}_q(\hat{q}_L(s)) = E_q[(\hat{q}_L(s) - p(s))^2]. \tag{4.10}$$

If we can choose a q such that $\text{var}_q(\hat{q}_L(s)) < \text{var}_p(\hat{p}_L(s))$, then $\hat{q}_L(s)$ will converge faster to $p_{S_n}(s)$ than $\hat{p}_L(s)$ as we increase L.

It can be proved (see Exercise 4.7.6) that in the class of all pdfs that are relatively continuous with respect to p there is a unique pdf q^* that minimizes the variance above. This **optimal** IS pdf has the form

$$q^*(\boldsymbol{\omega}) = \mathbf{1}_{\Delta_s}(S_n(\boldsymbol{\omega}))\frac{p(\boldsymbol{\omega})}{p_{S_n}(s)} = p(\boldsymbol{\omega}|S_n(\boldsymbol{\omega}) = s) \tag{4.11}$$

and has the desired property that $\text{var}_{q^*}(\hat{q}_L(s)) = 0$. It does seem therefore that our sampling problem is solved, until one realizes that q^* involves the unknown pdf that we want to estimate, namely, $p_{S_n}(s)$. Consequently, q^* cannot be used in practice as a sampling pdf. Other pdfs must be considered.

The next subsections present a more practical sampling pdf, which can be proved to be optimal in the large deviation limit $n \to \infty$. We will not attempt to justify the form of this pdf nor will we prove its optimality. Rather, we will attempt to illustrate how it works in the simplest way possible by applying it to IID sample means and Markov processes. For a more complete and precise treatment of IS of large deviation probabilities, see [4] and Chap. VI of [2]. For background information on IS, see Katzgraber (this volume) and Chap. V of [2].

4.3   Exponential change of measure

An important class of IS pdfs used for estimating $p_{S_n}(s)$ is given by

$$p_k(\boldsymbol{\omega}) = \frac{\mathrm{e}^{nkS_n(\boldsymbol{\omega})}}{W_n(k)}p(\boldsymbol{\omega}), \tag{4.12}$$

where $k \in \mathbb{R}$ and $W_n(k)$ is a normalization factor given by

$$W_n(k) = E_p[\mathrm{e}^{nkS_n}] = \int_{\mathbb{R}^n} \mathrm{e}^{nkS_n(\boldsymbol{\omega})}\, p(\boldsymbol{\omega})\, \mathrm{d}\boldsymbol{\omega}. \tag{4.13}$$

We call this class or family of pdfs parameterized by k the **exponential family**. In large deviation theory, p_k is also known as the **tilted pdf**

associated with p or the **exponential twisting** of p.[27] The likelihood ratio associated with this change of pdf is

$$R(\boldsymbol{\omega}) = \mathrm{e}^{-nkS_n(\boldsymbol{\omega})} W_n(k), \qquad (4.14)$$

which, for the purpose of estimating $I(s)$, can be approximated by

$$R(\boldsymbol{\omega}) \approx \mathrm{e}^{-n[kS_n(\boldsymbol{\omega})-\lambda(k)]} \qquad (4.15)$$

given the large deviation approximation $W_n(k) \approx \mathrm{e}^{n\lambda(k)}$.

It is important to note that a single pdf of the exponential family with parameter k cannot be used to efficiently sample the whole of the function $p_{S_n}(s)$ but only a particular point of that pdf. More precisely, if we want to sample $p_{S_N}(s)$ at the value s, we must choose k such that

$$E_{p_k}[\hat{q}_L(s)] = p_{S_n}(s) \qquad (4.16)$$

which is equivalent to solving $\lambda'(k) = s$ for k, where

$$\lambda(k) = \lim_{n\to\infty} \frac{1}{n} \ln E_p[\mathrm{e}^{nkS_n}] \qquad (4.17)$$

is the SCGF of S_n defined with respect to p (see Exercise 4.7.7). Let us denote the solution of these equations by $k(s)$. For many processes S_n, it can be proved that the tilted pdf $p_{k(s)}$ corresponding to $k(s)$ is asymptotically optimal in n, in the sense that the variance of $\hat{q}_L(s)$ under $p_{k(s)}$ goes asymptotically to 0 as $n \to \infty$. When this happens, the convergence of $\hat{q}_L(s)$ towards $p_{S_n}(s)$ is fast for increasing L and requires a sub-exponential number of samples to achieve a given accuracy for $I_{n,L}(s)$. To obtain the full rate function, we then simply need to repeat the sampling for other values of s, and so other values of k, and scale the whole process for larger values of n.

In practice, it is often easier to reverse the roles of k and s in the sampling. That is, instead of fixing s and selecting $k = k(s)$ for the numerical estimation of $I(s)$, we can fix k to obtain the rate function at $s(k) = \lambda'(k)$. This way, we can build a parametric representation of $I(s)$ over a certain range of s values by covering or "scanning" enough values of k.

At this point, there should be a suspicion that we will not achieve much with this **exponential change of measure** method (ECM method short) since the forms of p_k and $\hat{q}_L(s)$ presuppose the knowledge of $W_n(k)$ and in turn $\lambda(k)$. We know from the GE Theorem that $\lambda(k)$ is in many cases

---

[27] In statistics and actuarial mathematics, p_k is known as the **associated law** or **Esscher transform** of p.

sufficient to obtain $I(s)$, so why taking the trouble of sampling $\hat{q}_L(s)$ to get $I(s)$?[28]

The answer to this question is that $I(s)$ can be estimated from p_k without knowing $W_n(k)$ using two insights:

- Estimate $I(s)$ indirectly by sampling an estimator of $\lambda(k)$ or of S_n, which does not involve $W_n(k)$, instead of sampling the estimator $\hat{q}_L(s)$ of $p_{S_n}(s)$;

- Sample $\boldsymbol{\omega}$ according to the a priori pdf $p(\boldsymbol{\omega})$ or use the Metropolis algorithm, also known as the Markov chain Monte Carlo algorithm, to sample $\boldsymbol{\omega}$ according to $p_k(\boldsymbol{\omega})$ without knowing $W_n(k)$ (see Appendix A and the contribution of Katzgraber in this volume).

These points will be explained in the next section (see also Exercise 4.7.11).

For the rest of this subsection, we will illustrate the ECM method for a simple Gaussian IID sample mean, leaving aside the issue of having to know $W_n(k)$. For this process

$$p_k(\boldsymbol{\omega}) = p_k(x_1, \ldots, x_n) = \prod_{i=1}^{n} p_k(x_i), \tag{4.18}$$

where

$$p_k(x_i) = \frac{\mathrm{e}^{kx_i} \, p(x_i)}{W(k)}, \qquad W(k) = E_p[\mathrm{e}^{kX}], \tag{4.19}$$

with $p(x_i)$ a Gaussian pdf with mean μ and variance σ^2. We know from Exercise 3.6.1 the expression of the generating function $W(k)$ for this pdf:

$$W(k) = \mathrm{e}^{k\mu + \frac{\sigma^2}{2}k^2}. \tag{4.20}$$

The explicit expression of $p_k(x_i)$ is therefore

$$p_k(x_i) = \mathrm{e}^{k(x_i-\mu)-\frac{\sigma^2}{2}k^2} \frac{\mathrm{e}^{-(x_i-\mu)^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}} = \frac{\mathrm{e}^{-(x_i-\mu-\sigma^2 k)^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}}. \tag{4.21}$$

The estimated rate function obtained by sampling the X_i's according to this Gaussian pdf with mean $\mu + \sigma^2 k$ and variance σ^2 is shown in Fig. 4.2 for various values of n and L. This figure should be compared with Fig. 4.1 which illustrates the naive sampling method for the same sample mean. It is obvious from these two figures that the ECM method is more efficient than naive sampling.

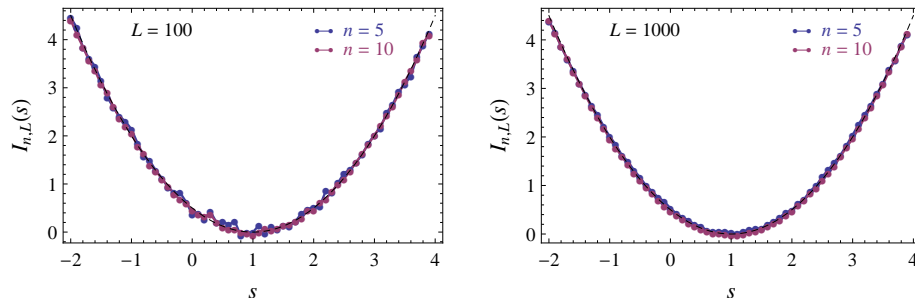[28]To make matters worst, the sampling of $\hat{q}_L(s)$ does involve $I(s)$ directly; see Exercise 4.7.10.

Figure 4.2: IS for the Gaussian IID sample mean ($\mu = \sigma = 1$) with the exponential change of measure.

As noted before, the rate function can be obtained in a parametric way by "scanning" k instead of "scanning" s and fixing k according to s. If we were to determine $k(s)$ for a given s, we would find

$$\lambda'(k) = \mu + \sigma^2 k = s, \tag{4.22}$$

so that $k(s) = (s - \mu)/\sigma^2$. Substituting this result in $p_k(x_i)$, then yields

$$p_{k(s)}(x_i) = \frac{e^{-(x_i - s)^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}}. \tag{4.23}$$

Thus to efficiently sample $p_{S_n}(s)$, we must sample the X_i's according to a Gaussian pdf with mean s instead of the a priori pdf with mean μ. The sampling is efficient in this case simply because S_n concentrates in the sense of the LLN at the value s instead of μ, which means that s has become the typical value of S_n under $p_{k(s)}$.

The idea of the ECM method is the same for processes other than IID sample means: the goal in general is to change the sampling pdf from p to p_k in such a way that an event that is a rare event under p becomes a typical under p_k.[29] For more information on the ECM method and IS in general; see [4] and Chaps. V and VI of [2]. For applications of these methods, see Exercise 4.7.9.

4.4 Applications to Markov chains

The application of the ECM method to Markov chains follows the IID case closely: we generate L IID realizations $x_1^{(j)}, \ldots, x_n^{(j)}$, $j = 1, \ldots, L$, of the

---

[29]Note that this change is not always possible: for certain processes, there is no k that makes certain events typical under p_k. For examples, see [1].

chain $\boldsymbol{\omega} = X_1, \ldots, X_n$ according to the tilted joint pdf $p_k(\boldsymbol{\omega})$, using either p_k directly or a Metropolis-type algorithm. The value $k(s)$ that achieves the efficient sampling of S_n at the value s is determined in the same way as in the IID case by solving $\lambda'(k) = s$, where $\lambda(k)$ is now given by Eq. (3.9). In this case, $S_n = s$ becomes the typical event of S_n in the limit $n \to \infty$. A simple application of this procedure is proposed in Exercise 4.7.14. For further reading on the sampling of Markov chains, see [4, 42, 43, 5, 41].

An important difference to notice between the IID and Markov cases is that the ECM does not preserve the factorization structure of $p(\boldsymbol{\omega})$ for the latter case. In other words, $p_k(\boldsymbol{\omega})$ does not describe a Markov chain, though, as we noticed, $p_k(\boldsymbol{\omega})$ is a product pdf when $p(\boldsymbol{\omega})$ is itself a product pdf. In the case of Markov chains, $p_k(\boldsymbol{\omega})$ does retain a product structure that looks like a Markov chain, and one may be tempted to define a tilted transition pdf as

$$\pi_k(x'|x) = \frac{\mathrm{e}^{kx'} \pi(x'|x)}{W(k|x)}, \qquad W(k|x) = \int_{\mathbb{R}} \mathrm{e}^{kx'} \pi(x'|x)\, \mathrm{d}x'. \qquad (4.24)$$

However, it is easy to see that the joint pdf of $\boldsymbol{\omega}$ obtained from this transition matrix does not reproduce the tilted pdf $p_k(\boldsymbol{\omega})$.

4.5 Applications to continuous-time Markov processes

The generalization of the results of the previous subsection to continuous-time Markov processes follows from our discussion of the continuous-time limit of Markov chains (see Sec. 3.2). In this case, we generate, according to the tilted pdf of the process, a sample of time-discretized trajectories. The choice of Δt and n used in the discretization will of course influence the precision of the estimates of the pdf and rate function, in addition to the sample size L.

Similarly to Markov chains, the tilted pdf of a continuous-time Markov process does not in general describe a new "tilted" Markov process having a generator related to the tilted generator $\tilde{G}_k$ defined earlier. In some cases, the tilting of a Markov process can be seen as generating a new Markov process, as will be discussed in the next subsection, but this is not true in general.

4.6 Applications to stochastic differential equations

Stochastic differential equations (SDEs) are covered by the continuous-time results of the previous subsection. However, for this type of stochastic processes the ECM can be expressed more explicitly in terms of the path

pdf $p[x]$ introduced in Sec. 3.4. The tilted version of this pdf, which is the functional analogue of $p_k(\boldsymbol{\omega})$, is written as

$$p_k[x] = \frac{\mathrm{e}^{TkS_T[x]} p[x]}{W_T(k)}, \tag{4.25}$$

where $S_T[x]$ is some functional of the trajectory $\{x(t)\}_{t=0}^{T}$ and

$$W_T(k) = E_p[\mathrm{e}^{TkS_T}]. \tag{4.26}$$

The likelihood ratio associated with this change of pdf is

$$R[x] = \frac{p[x]}{p_k[x]} = \mathrm{e}^{-TkS_T[x]} W_T(k) \approx \mathrm{e}^{-T[kS_T[x]-\lambda(k)]}. \tag{4.27}$$

As a simple application of these expressions, let us consider the SDE

$$\dot{x}(t) = \xi(t), \tag{4.28}$$

where $\xi(t)$ is a Gaussian white noise, and the following additive process:

$$D_T[x] = \frac{1}{T} \int_0^T \dot{x}(t)\,\mathrm{d}t, \tag{4.29}$$

which represents the average velocity or **drift** of the SDE. What is the form of $p_k[x]$ in this case? Moreover, how can we use this tilted pdf to estimate the rate function $I(d)$ of D_T in the long-time limit?

An answer to the first question is suggested by recalling the goal of the ECM. The LLN implies for the SDE above that $D_T \to 0$ in probability as $T \to \infty$, which means that $D_T = 0$ is the typical event of the "natural" dynamic of $x(t)$. The tilted dynamics realized by $p_{k(d)}[x]$ should change this typical event to the fluctuation $D_T = d$; that is, the typical event of D_T under $p_{k(d)}[x]$ should be $D_T = d$ rather than $D_T = 0$. One candidate SDE which leads to this behavior of D_T is

$$\dot{x}(t) = d + \xi(t), \tag{4.30}$$

so the obvious guess is that $p_{k(d)}[x]$ is the path pdf of this SDE.

Let us show that this is a correct guess. The form of $p_k[x]$ according to Eq. (4.25) is

$$p_k[x] \approx \mathrm{e}^{-TI_k[x]}, \qquad I_k[x] = I[x] - kD_T[x] + \lambda(k), \tag{4.31}$$

where

$$I[x] = \frac{1}{2T} \int_0^T \dot{x}^2 \, \mathrm{d}t \qquad (4.32)$$

is the action of our original SDE and

$$\lambda(k) = \lim_{T \to \infty} \frac{1}{T} \ln E[\mathrm{e}^{TkD_T}] \qquad (4.33)$$

is the SCGF of D_T. The expectation entering in the definition of $\lambda(k)$ can be calculated explicitly using path integrals with the result $\lambda(k) = k^2/2$. Therefore,

$$I_k[x] = I[x] - kD_T[x] - \frac{k^2}{2}. \qquad (4.34)$$

From this result, we find $p_{k(d)}[x]$ by solving $\lambda'(k(d)) = d$, which in this case simply yields $k(d) = d$. Hence

$$I_{k(d)}[x] = I[x] - dD_T[x] + \frac{d^2}{2} = \frac{1}{2T} \int_0^T (\dot{x} - d)^2 \, \mathrm{d}t, \qquad (4.35)$$

which is exactly the action of the **boosted SDE** of Eq. (4.30).

Since we have $\lambda(k)$, we can of course directly obtain the rate function $I(d)$ of D_T by Legendre transform:

$$I(d) = k(d)d - \lambda(k(d)) = \frac{d^2}{2}. \qquad (4.36)$$

This shows that the fluctuations of D_T are Gaussian, which is not surprising given that D_T is up to a factor $1/T$ a Brownian motion.[30] Note that the same result can be obtained following Exercise 4.7.10 by calculating the likelihood ratio $R[x]$ for a path $\{x_d(t)\}_{t=0}^T$ such that $D_T[x_d] = d$, i.e.,

$$R[x_d] = \frac{p[x_d]}{p_{k(d)}[x_d]} \approx \mathrm{e}^{-TI(d)}. \qquad (4.37)$$

These calculations give a representative illustration of the ECM method and the idea that the large deviations of an SDE can be obtained by replacing this SDE by a boosted or **effective SDE** whose typical events are large deviations of the former SDE. Exercises 4.7.16 and 4.7.17 show that effective SDEs need not correspond exactly to an ECM, but may in some cases reproduce such a change of measure in the large deviation limit $T \to \infty$.

---

[30]Brownian motion is defined as the integral of Gaussian white noise. By writing $\dot{x} = \xi(t)$, we therefore define $x(t)$ to be a Brownian motion.

The important point in all cases is to be able to calculate the likelihood ratio between a given SDE and a boosted version of it in order to obtain the desired rate function in the limit $T \to \infty$.[31]

For the particular drift process D_T studied above, the likelihood ratio happens to be equivalent to a mathematical result known as **Girsanov's formula**, which is often used in financial mathematics; see Exercise 4.7.15.

4.7 Exercises

4.7.1.[5] (Variance of Bernoulli RVs) Let X be a Bernoulli RV with $P(X = 1) = \alpha$ and $P(X = 0) = 1 - \alpha$, $0 \leq \alpha \leq 1$. Find the variance of X in terms of α.

4.7.2.[17] (Direct sampling of IID sample means) Generate on a computer a sample of size $L = 100$ of n IID Bernoulli RVs $X_1, \ldots, X_n$ with bias $\alpha = 0.5$, and numerically estimate the rate function $I(s)$ associated with the sample mean S_n of these RVs using the naive estimator $\hat{p}_L(s)$ and the finite-size rate function $I_{n,L}(s)$ defined in Eq. (4.6). Repeat for $L = 10^3, 10^4, 10^5$, and 10^6 and observe how $I_{n,L}(s)$ converges to the exact rate function given in Eq. (2.10). Repeat for exponential RVs.

4.7.3.[10] (Unbiased estimator) Prove that $\hat{q}_L(s)$ is an unbiased estimator of $p_{S_n}(s)$, i.e., prove Eq. (4.9).

4.7.4.[15] (Estimator variance) The estimator $\hat{p}_L(s)$ is actually a sample mean of Bernoulli RVs. Based on this, calculate the variance of this estimator with respect to $p(\omega)$ using the result of Exercise 4.7.1 above. Then calculate the variance of the importance estimator $\hat{q}_L(s)$ with respect to $q(\omega)$.

4.7.5.[15] (Relative error of estimators) The real measure of the quality of the estimator $\hat{p}_L(s)$ is not so much its variance $\text{var}(\hat{p}_L(s))$ but its relative error $\text{err}(\hat{p}_L(s))$ defined by

$$\text{err}(\hat{p}_L(s)) = \frac{\sqrt{\text{var}(\hat{p}_L(s))}}{p_{S_n}(s)}. \tag{4.38}$$

Use the previous exercise to find an approximation of $\text{err}(\hat{p}_L(s))$ in terms of L and n. From this result, show that the sample size L needed to achieve a certain relative error grows exponentially with n. (Source: Sec. V.I of [2].)

---

[31]The limit $\varepsilon \to 0$ can also be considered.

4.7.6.[17] (Optimal importance pdf) Show that the density q^* defined in (4.11) is optimal in the sense that

$$\text{var}_q(\hat{q}_L(s)) \geq \text{var}_{q^*}(\hat{q}_L(s)) = 0 \tag{4.39}$$

for all q relatively continuous to p. (See Chap. V of [2] for help.)

4.7.7.[15] (Exponential change of measure) Show that the value of k solving Eq. (4.16) is given by $\lambda'(k) = s$, where $\lambda(k)$ is the SCGF defined in Eq. (4.3).

4.7.8.[40] (LDP for estimators) We noticed before that $\hat{p}_L(s)$ is the empirical vector of the IID sample $\{S_n^{(j)}\}_{j=1}^L$. Based on this, we could study the LDP of $\hat{p}_L(s)$ rather than just its mean and variance, as usually done in sampling theory. What is the form of the LDP associated with $\hat{p}_L(s)$ with respect to $p(\boldsymbol{\omega})$? What is its rate function? What is the minimum and zero of that rate function? Answer the same questions for $\hat{q}_L(s)$ with respect to $q(\boldsymbol{\omega})$. (Hint: Use Sanov's Theorem and the results of Exercise 3.6.6.)

4.7.9.[25] (Importance sampling of IID sample means) Implement the ECM method to numerically estimate the rate function associated with the IID sample means of Exercise 3.6.1. Use a simple IID sampling based on the explicit expression of p_k in each case (i.e., assume that $W(k)$ is known). Study the convergence of the results as a function of n and L as in Fig. 4.2.

4.7.10.[15] (IS estimator) Show that $\hat{q}_L(s)$ for $p_{k(s)}$ has the form

$$\hat{q}_L(s) \approx \frac{1}{L\Delta s} \sum_{j=1}^{L} \mathbf{1}_{\Delta_s}(s^{(j)}) \, e^{-nI(s^{(j)})}. \tag{4.40}$$

Why is this estimator trivial for estimating $I(s)$?

4.7.11.[30] (Metropolis sampling of IID sample means) Repeat Exercise 4.7.9, but instead of using an IID sampling method, use the Metropolis algorithm to sample the X_i's in S_n. (For information on the Metropolis algorithm, see Appendix A.)

4.7.12.[15] (Tilted distribution from contraction) Let $X_1, \ldots, X_n$ be a sequence of real IID RVs with common pdf $p(x)$, and denote by S_n and $L_n(x)$ the sample mean and empirical functional, respectively, of these RVs. Show that the "value" of $L_n(x)$ solving the minimization problem associated with the contraction of the LDP of $L_n(x)$ down to the LDP of S_n is given by

$$\mu_k(x) = \frac{e^{kx}p(x)}{W(k)}, \qquad W(k) = E_p[e^{kX}], \tag{4.41}$$

where k is such that $\lambda'(k) = (\ln W(k))' = s$. To be more precise, show that this pdf is the solution of the minimization problem

$$\inf_{\mu:f(\mu)=s} I(\mu), \qquad (4.42)$$

where

$$I(\mu) = \int_{\mathbb{R}} \mathrm{d}x\, \mu(x) \ln \frac{\mu(x)}{p(x)} \qquad (4.43)$$

is the continuous analogue of the relative entropy defined in Eq. (3.7) and

$$f(\mu) = \int_{\mathbb{R}} x\, \mu(x)\, \mathrm{d}x \qquad (4.44)$$

is the contraction function. Assume that $W(k)$ exists. Why is μ_k the same as the tilted pdf used in IS?

4.7.13.[50] (Equivalence of ensembles) Comment on the meaning of the following statements: (i) The optimal pdf q^* is to the microcanonical ensemble what the tilted pdf p_k is to the canonical ensemble. (ii) Proving that p_k achieves a zero variance for the estimator $\hat{q}_L(s)$ in the limit $n \to \infty$ is the same as proving that the canonical ensemble becomes equivalent to the microcanonical ensemble in the thermodynamic limit.

4.7.14.[25] (Tilted Bernoulli Markov chain) Find the expression of the tilted joint pdf $p_k(\boldsymbol{\omega})$ for the Bernoulli Markov chain of Exercise 3.6.10. Then use the Metropolis algorithm (see Appendix A) to generate a large-enough sample of realizations of $\boldsymbol{\omega}$ in order to estimate $p_{S_n}(s)$ and $I(s)$. Study the converge of the estimation of $I(s)$ in terms of n and L towards the expression of $I(s)$ found in Exercise 3.6.10. Note that for a starting configuration $x_1, \ldots, x_n$ and a target configuration $x'_1, \ldots, x'_n$, the acceptance ratio is

$$\frac{p_k(x'_1, \ldots, x'_n)}{p_k(x_1, \ldots, x_n)} = \frac{\mathrm{e}^{kx'_1} p(x'_1) \prod_{i=2}^{n} \mathrm{e}^{kx'_i} \pi(x'_i | x'_{i-1})}{\mathrm{e}^{kx_1} p(x_1) \prod_{i=2}^{n} \mathrm{e}^{kx_i} \pi(x_i | x_{i-1})}, \qquad (4.45)$$

where $p(x_i)$ is some initial pdf for the first state of the Markov chain. What is the form of $\hat{q}_L(s)$ in this case?

4.7.15.[20] (Girsanov's formula) Denote by $p[x]$ the path pdf associated with the SDE

$$\dot{x}(t) = \sqrt{\varepsilon}\, \xi(t), \qquad (4.46)$$

where $\xi(t)$ is Gaussian white noise with unit noise power. Moreover, let $q[x]$ be the path pdf of the boosted SDE

$$\dot{x}(t) = \mu + \sqrt{\varepsilon}\,\xi(t). \tag{4.47}$$

Show that

$$R[x] = \frac{p[x]}{q[x]} = \exp\left(-\frac{\mu}{\varepsilon}\int_0^T \dot{x}\,\mathrm{d}t + \frac{\mu^2}{2\varepsilon}T\right) = \exp\left(-\frac{\mu}{\varepsilon}x(T) + \frac{\mu^2}{2\varepsilon}T\right). \tag{4.48}$$

4.7.16.[27] (Effective SDE) Consider the additive process,

$$S_T = \frac{1}{T}\int_0^T x(t)\,\mathrm{d}t, \tag{4.49}$$

with $x(t)$ evolving according to the simple Langevin equation of Exercise 3.6.14. Show that the tilted path pdf $p_{k(d)}[x]$ associated with this process is *asymptotically* equal to the path pdf of the following boosted SDE:

$$\dot{x}(t) = -\gamma(x(t) - s) + \xi(t). \tag{4.50}$$

More precisely, show that the action $I_{k(s)}[x]$ of $p_{k(s)}[x]$ and the action $J[x]$ of the boosted SDE differ by a boundary term which is of order $O(1/T)$.

4.7.17.[22] (Non-exponential change of measure) Although the path pdf of the boosted SDE of Eq. (4.50) is not exactly the tilted path pdf $p_k[x]$ obtained from the ECM, the likelihood ratio $R[x]$ of these two pdfs does yield the rate function of S_T. Verify this by obtaining $R[x]$ exactly and by using the result of Exercise 4.7.10.

5   Other numerical methods for large deviations

The use of the ECM method appears to be limited, as mentioned before, by the fact that the tilted pdf $p_k(\boldsymbol{\omega})$ involves $W_n(k)$. We show in this section how to circumvent this problem by considering estimators of $W_n(k)$ and $\lambda(k)$ instead of estimators of the pdf $p(S_n)$, and by sampling these new estimators according to the tilted pdf $p_k(\boldsymbol{\omega})$, but with the Metropolis algorithm in order not to rely on $W_n(k)$, or directly according to the a priori pdf $p(\boldsymbol{\omega})$. At the end of the section, we also briefly describe other important methods used in numerical simulations of large deviations.

5.1 Sample mean method

We noted in the previous subsection that the parameter k entering in the ECM had to be chosen by solving $\lambda'(k) = s$ in order for s to be the typical event of S_n under p_k. The reason for this is that

$$\lim_{n \to \infty} E_{p_k}[S_n] = \lambda'(k). \tag{5.1}$$

By the LLN we also have that $S_n \to \lambda'(k)$ in probability as $n \to \infty$ if S_n is sampled according to p_k.

These results suggest using

$$s_{L,n}(k) = \frac{1}{L} \sum_{j=1}^{L} S_n(\boldsymbol{\omega}^{(j)}) \tag{5.2}$$

as an estimator of $\lambda'(k)$ by sampling $\boldsymbol{\omega}$ according to $p_k(\boldsymbol{\omega})$, and then integrating this estimator with the boundary condition $\lambda(0) = 0$ to obtain an estimate of $\lambda(k)$. In other words, we can take our estimator for $\lambda(k)$ to be

$$\hat{\lambda}_{L,n}(k) = \int_0^k s_{L,n}(k') \, \mathrm{d}k'. \tag{5.3}$$

From this estimator, we then obtain a parametric estimation of $I(s)$ from the GE Theorem by Legendre-transforming $\hat{\lambda}_{L,n}(k)$:

$$I_{L,n}(s) = ks - \hat{\lambda}_{L,n}(k), \tag{5.4}$$

where $s = s_{L,n}(k)$ or, alternatively,

$$I_{L,n}(s) = k(s)\, s - \hat{\lambda}_{L,n}(k(s)), \tag{5.5}$$

where $k(s)$ is the root of $\hat{\lambda}'_{L,n}(k) = s$.[32]

The implementation of these estimators with the Metropolis algorithm is done explicitly via the following steps:

1. For a given $k \in \mathbb{R}$, generate an IID sample $\{\boldsymbol{\omega}^{(j)}\}_{j=1}^{L}$ of L configurations $\boldsymbol{\omega} = x_1, \ldots, x_n$ distributed according to $p_k(\boldsymbol{\omega})$ using the Metropolis algorithm, which only requires the computation of the ratio

$$\frac{p_k(\boldsymbol{\omega})}{p_k(\boldsymbol{\omega}')} = \frac{\mathrm{e}^{nkS_n(\boldsymbol{\omega})}\, p(\boldsymbol{\omega})}{\mathrm{e}^{nkS_n(\boldsymbol{\omega}')}\, p(\boldsymbol{\omega}')} \tag{5.6}$$

for any two configurations $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$; see Appendix A;

---

[32]In practice, we have to check numerically (using finite-scale analysis) that $\lambda_{n,L}(k)$ does not develop non-differentiable points in k as n and L are increased to ∞. If non-differentiable points arise, then $I(s)$ is recovered from the GE Theorem only in the range of λ'; see Sec. 4.4 of [48].
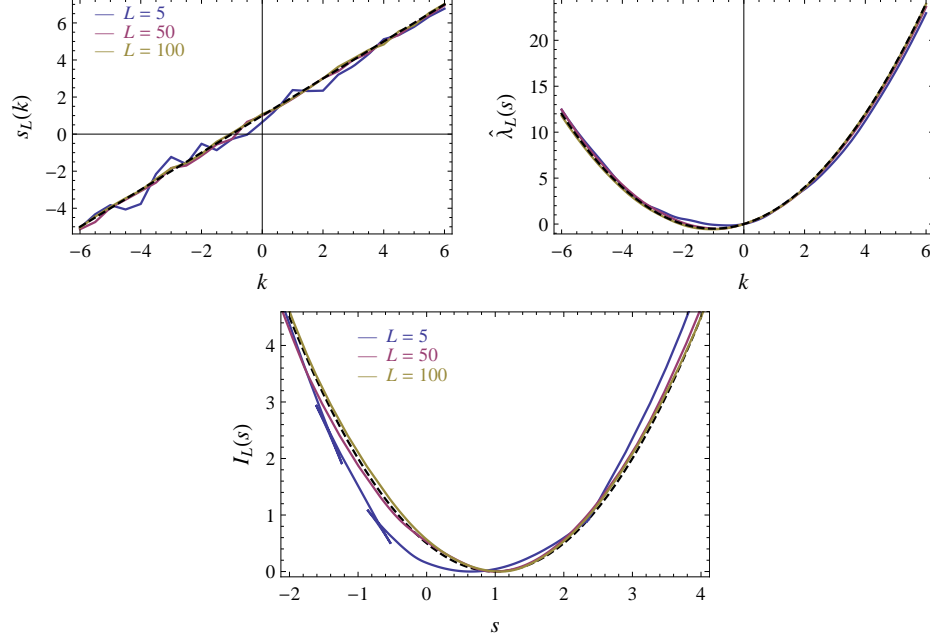
Figure 5.1: (Top left) Empirical sample mean $s_L(k)$ for the Gaussian IID sample mean ($\mu = \sigma = 1$). The spacing of the k values is $\Delta k = 0.5$. (Top right) SCGF $\hat{\lambda}_L(k)$ obtained by integrating $s_L(k)$. (Bottom) Estimate of $I(s)$ obtained from the Legendre transform of $\hat{\lambda}_L(k)$. The dashed line in each figure is the exact result. Note that the estimate of the rate function for $L = 5$ is multi-valued because the corresponding $\hat{\lambda}_L(k)$ is nonconvex.

2. Compute the estimator $s_{L,n}(k)$ of $\lambda'(k)$ for the generated sample;

3. Repeat the two previous steps for other values of k to obtain a numerical approximation of the function $\lambda'(k)$;

4. Numerically integrate the approximation of $\lambda'(k)$ over the mesh of k values starting from $k = 0$. The result of the integration is the estimator $\hat{\lambda}_{L,n}(k)$ of $\lambda(k)$;

5. Numerically compute the Legendre transform, Eq. (5.4), of the estimator of $\lambda(k)$ to obtain the estimator $I_{L,n}(s)$ of $I(s)$ at $s = s_{L,n}(k)$;

6. Repeat the previous steps for larger values of n and L until the results converge to some desired level of accuracy.

The result of these steps are illustrated in Fig. 5.1 for our test case of the IID Gaussian sample mean for which $\lambda'(k) = \mu + \sigma^2 k$. Exercise 5.4.1 covers other sample means studied before (e.g., exponential, binary, etc.).

Note that for IID sample means, one does not need to generate L realizations of the whole sequence $\boldsymbol{\omega}$, but only L realizations of one summand X_i according the tilted marginal pdf $p_k(x) = e^{kx}p(x)/W(k)$. In this case, L takes the role of the large deviation parameter n. For Markov chains, realizations of the whole sequence $\boldsymbol{\omega}$ must be generated one after the other, e.g., using the Metropolis algorithm.

5.2 Empirical generating functions

The last method that we cover is based on the estimation of the generating function

$$W_n(k) = E[e^{nkS_n}] = \int_{\mathbb{R}^n} p(\boldsymbol{\omega}) \, e^{nkS_n(\boldsymbol{\omega})} \, d\boldsymbol{\omega}. \tag{5.7}$$

Consider first the case where S_n is an IID sample mean. Then $\lambda(k) = \ln W(k)$, where $W(k) = E[e^{kX_i}]$, as we know from Sec. 3.1, and so the problem reduces to the estimation of the SCGF of the common pdf $p(X)$ of the X_i's. An obvious estimator for this function is

$$\hat{\lambda}_L(k) = \ln \frac{1}{L} \sum_{j=1}^{L} e^{kX^{(j)}}, \tag{5.8}$$

where $X^{(j)}$ are the IID samples drawn from $p(X)$. From this estimator, which we call the **empirical SCGF**, we build a parametric estimator of the rate function $I(s)$ of S_n using the Legendre transform of Eq. (5.4) with

$$s(k) = \hat{\lambda}'_L(k) = \frac{\sum_{j=1}^{L} X^{(j)} \, e^{kX^{(j)}}}{\sum_{j=1}^{L} e^{kX^{(j)}}}. \tag{5.9}$$

Fig. 5.2 shows the result of these estimators for the IID Gaussian sample mean. As can be seen, the convergence of $I_L(s)$ is fast in this case, but is limited to the central part of the rate function. This illustrates one limitation of the empirical SCGF method: for unbounded RVs, such as the Gaussian sample mean, $W(k)$ is correctly recovered for large $|k|$ only for large samples, i.e., large L, which implies that the tails of $I(s)$ are also recovered only for large L. For bounded RVs, such as the Bernoulli sample mean, this limitation does not arise; see [14, 15] and Exercise 5.4.2.

The application of the empirical SCGF method to sample means of ergodic Markov chains is relatively direct. In this case, although W_n no
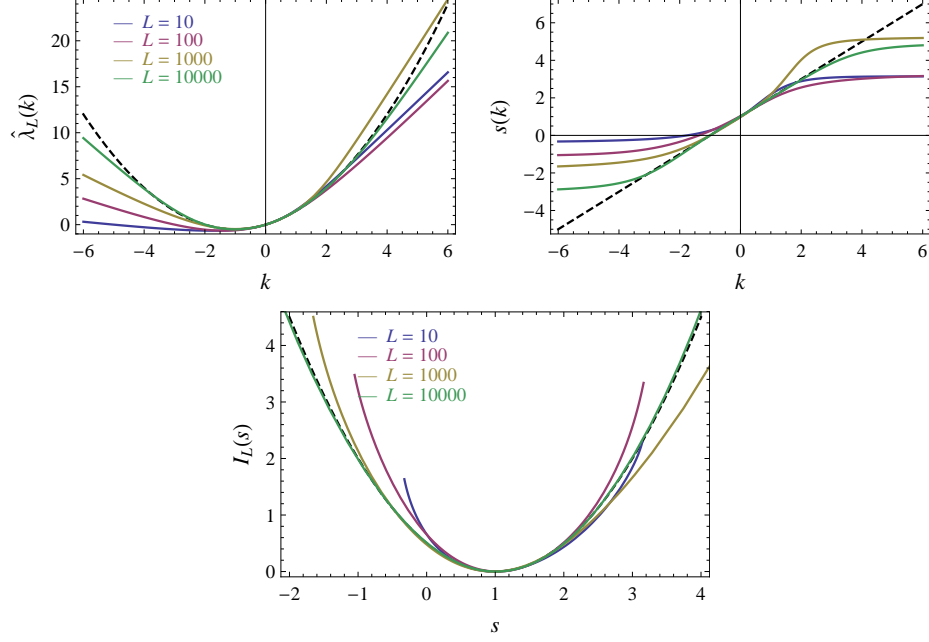
46

Figure 5.2: (Top left) Empirical SCGF $\hat{\lambda}_L(k)$ for the Gaussian IID sample mean ($\mu = \sigma = 1$). (Top right) $s(k) = \hat{\lambda}'_L(k)$. (Bottom) Estimate of $I(s)$ obtained from $\hat{\lambda}_L(k)$. The dashed line in each figure is the exact result.

longer factorizes into $W(k)^n$, it is possible to "group" the X_i's into blocks of b RVs as follows:

$$\underbrace{X_1 + \cdots + X_b}_{Y_1} + \underbrace{X_{b+1} + \cdots + X_{2b}}_{Y_2} + \cdots + \underbrace{X_{n-b+1} + \cdots + X_n}_{Y_m} \qquad (5.10)$$

where $m = n/b$, so as to rewrite the sample mean S_n as

$$S_n = \frac{1}{n} \sum_{i=1}^{n} X_i = \frac{1}{bm} \sum_{i=1}^{m} Y_i. \qquad (5.11)$$

If b is large enough, then the blocks Y_i can be treated as being independent. Moreover, if the Markov chain is ergodic, then the Y_i's should be identically distributed for i large enough. As a result, $W_n(k)$ can be approximated for large n and large b (but $b \ll n$) as

$$W_n(k) = E[e^{k \sum_{i=1}^{m} Y_i}] \approx E[e^{kY_i}]^m, \qquad (5.12)$$

so that
$$\lambda_n(k) = \frac{1}{n} \ln W_n(k) \approx \frac{m}{n} \ln E[\mathrm{e}^{kY_i}] = \frac{1}{b} \ln E[\mathrm{e}^{kY_i}]. \qquad (5.13)$$

We are thus back to our original IID problem of estimating a generating function; the only difference is that we must perform the estimation at the level of the Y_i's instead of the X_i's. This means that our estimator for $\lambda(k)$ is now

$$\hat{\lambda}_{L,n}(k) = \frac{1}{b} \ln \frac{1}{L} \sum_{j=1}^{L} \mathrm{e}^{kY^{(j)}}, \qquad (5.14)$$

where $Y^{(j)}$, $j = 1, \ldots, L$ are IID samples of the block RVs. Following the IID case, we then obtain an estimate of $I(s)$ in the usual way using the Legendre transform of Eq. (5.4) or (5.5).

The application of these results to Markov chains and SDEs is covered in Exercises 5.4.2 and 5.4.3. For more information on the empirical SCGF method, including a more detailed analysis of its convergence, see [14, 15].

5.3   Other methods

We briefly describe below a number of other important methods used for estimating large deviation probabilities, and give for each of them some pointer references for interested readers. The fundamental ideas and concepts related to sampling and ECM that we have treated before also arise in these other methods.

- **Splitting methods:** Splitting or cloning methods are the most used sampling methods after Monte Carlo algorithms based on the ECM. The idea of splitting is to "grow" a sample or "population" $\{\boldsymbol{\omega}\}$ of sequences or configurations in such a way that a given configuration $\boldsymbol{\omega}$ appears in the population in proportion to its weight $p_k(\boldsymbol{\omega})$. This is done in some iterative fashion by "cloning" and "killing" some configurations. The result of this process is essentially an estimate for $W_n(k)$, which is then used to obtain $I(s)$ by Legendre transform. For mathematically-oriented introductions to splitting methods, see [33, 34, 9, 37]; for more physical introductions, see [23, 32, 46, 47].

- **Optimal path method:** We briefly mentioned in Sec. 3.4 that rate functions related to SDEs can often be derived by contraction of the action functional $I[x]$. The optimization problem that results from this contraction is often not solvable analytically, but there are several techniques stemming from optimization theory, classical mechanics,

and dynamic programming that can be used to solve it. A typical optimization problem in this context is to find the optimal path that a noisy system follows with highest probability to reach a certain state considered to be a rare event or fluctuation. This optimal path is also called a **transition path** or a **reactive trajectory**, and can be found analytically only in certain special systems, such as linear SDEs and conservative systems; see Exercises 5.4.4–5.4.6. For nonlinear systems, the Lagrangian or Hamiltonian equations governing the dynamics of these paths must be simulated directly; see [26] and Sec. 6.1 of [48] for more details.

- **Transition path method:** The transition path method is essentially a Monte Carlo sampling method in the space of configurations $\boldsymbol{\omega}$ (for discrete-time processes) or trajectories $\{x(t)\}_{t=0}^{T}$ (for continuous-time processes and SDEs), which aims at finding, as the name suggests, transition paths or optimal paths. The contribution of Dellago in this volume covers this method in detail; see also [11, 12, 10, 36, 51]. A variant of the transition method is the so-called **string method** [17] which evolves paths in conservative systems towards optimal paths connecting different local minima of the potential or landscape function of these systems.

- **Eigenvalue method:** We have noted that the SCGF of an ergodic continuous-time process is given by the dominant eigenvalue of its tilted generator. From this connection, one can attempt to numerically evaluate SCGFs using numerical approximation methods for linear functional operators, combined with numerical algorithms for finding eigenvalues. A recent application of this approach, based on the renormalization group, can be found in [24, 25].

5.4   Exercises

5.4.1.[25] (Sample mean method) Implement the steps of the sample mean method to numerically estimate the rate functions of all the IID sample means studied in these notes, including the sample mean of the Bernoulli Markov chain.

5.4.2.[25] (Empirical SCGF method) Repeat the previous exercise using the empirical SCGF method instead of the sample mean method. Explain why the former method converges quickly when S_n is bounded, e.g., in the

Bernoulli case. Analyze in detail how the method converges when S_n is unbounded, e.g., for the exponential case. (Source: [14, 15].)

5.4.3.[30] (Sampling of SDEs) Generalize the sample mean and empirical mean methods to SDEs. Study the additive process of Exercise 4.7.16 as a test case.

5.4.4.[30] (Optimal paths for conservative systems) Consider the conservative system of Exercise 3.6.15 and assume, without loss of generality, that the global minimum of the potential $U(\mathbf{x})$ is at $\mathbf{x} = 0$. Show that the optimal path which brings the system from $\mathbf{x}(0) = 0$ to some fluctuation $\mathbf{x}(T) = \mathbf{x} \neq 0$ after a time T is the time-reversal of the solution of the noiseless dynamics $\dot{\mathbf{x}} = -\nabla U(\mathbf{x})$ under the terminal conditions $\mathbf{x}(0) = \mathbf{x}$ and $\mathbf{x}(T) = 0$. The latter path is often called the **decay path**.

5.4.5.[25] (Optimal path for additive processes) Consider the additive process S_T of Exercise 4.7.16 under the Langevin dynamics of Exercise 3.6.14. Show that the optimal path leading to the fluctuation $S_T = s$ is the constant path

$$x(t) = s, \quad t \in [0, T]. \tag{5.15}$$

Explain why this result is the same as the asymptotic solution of the boosted SDE of Eq. (4.50).

5.4.6.[25] (Optimal path for the dragged Brownian particle) Apply the previous exercise to Exercise 3.6.18. Discuss the physical interpretation of the results.

Acknowledgments

A   Metropolis-Hastings algorithm

The **Metropolis-Hastings algorithm**[33] or **Markov chain Monte Carlo algorithm** is a simple method used to generate a sequence $x_1, \ldots, x_n$ of variates whose empirical distribution $L_n(x)$ converges asymptotically to a

---

[33]Or, more precisely, the Metropolis–Rosenbluth–Rosenbluth–Teller–Teller–Hastings algorithm.

target pdf $p(x)$ as $n \to \infty$. The variates are generated sequentially in the manner of a Markov chain as follows:

1. Given the value $x_i = x$, choose another value x', called the **proposal** or **move**, according to some (fixed) conditional pdf $q(x'|x)$, called the **proposal pdf**;

2. Accept the move x', i.e., set $x_{i+1} = x'$ with probability $\min\{1, a\}$ where

$$a = \frac{p(x')\, q(x|x')}{p(x)\, q(x'|x)};$$

(A.1)

3. If the move is not accepted, set $x_{i+1} = x$.

This algorithm is very practical for sampling pdfs having an exponential form, such as the Boltzmann-Gibbs distribution or the tilted pdf p_k, because it requires only the computation of the **acceptance ratio** a for two values x and x', and so does not require the knowledge of the normalization constant that usually enters in these pdfs. For the tilted pdf $p_k(\boldsymbol{\omega})$, for example,

$$a = \frac{p_k(\boldsymbol{\omega}')\, q(\boldsymbol{\omega}|\boldsymbol{\omega}')}{p_k(\boldsymbol{\omega})\, q(\boldsymbol{\omega}'|\boldsymbol{\omega})} = \frac{\mathrm{e}^{nkS_n(\boldsymbol{\omega})}\, p(\boldsymbol{\omega})\, q(\boldsymbol{\omega}|\boldsymbol{\omega}')}{\mathrm{e}^{nkS_n(\boldsymbol{\omega})}\, p(\boldsymbol{\omega})\, q(\boldsymbol{\omega}'|\boldsymbol{\omega})}.$$

(A.2)

The original **Metropolis algorithm** commonly used in statistical physics is obtained by choosing $q(x'|x)$ to be symmetric, i.e., if $q(x'|x) = q(x|x')$ for all x and x', in which case $a = p(x')/p(x)$. If $q(x'|x)$ does not depend on x, the algorithm is referred to as the **independent chain Metropolis-Hastings algorithm**.

For background information on Monte Carlo algorithms, applications, and technical issues such as mixing, correlation and convergence, see [31] and Chap. 13 of [2].

References

[1] S. Asmussen, P. Dupuis, R. Rubinstein, and H. Wang. Importance sampling for rare event. In S. Gass and M. Fu, editors, *Encyclopedia of Operations Research and Management Sciences*. Kluwer, 3rd edition, 2011.

[2] S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Stochastic Modelling and Applied Probability. Springer, New York, 2007.

[3] C. M. Bender and S. A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers*. McGraw-Hill, New York, 1978.

[4] J. A. Bucklew. *Introduction to Rare Event Simulation*. Springer, New York, 2004.

[5] J. A. Bucklew, P. Ney, and J. S. Sadowsky. Monte Carlo simulation and large deviations theory for uniformly recurrent Markov chains. *J. Appl. Prob.*, 27(1):44–59, 1990.

[6] R. Chetrite, G. Falkovich, and K. Gawedzki. Fluctuation relations in simple examples of non-equilibrium steady states. *J. Stat. Mech.*, 2008(08):P08005, 2008.

[7] R. Chetrite and K. Gawedzki. Fluctuation relations for diffusion processes. *Comm. Math. Phys.*, 282(2):469–518, 2008.

[8] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley, New York, 1991.

[9] T. Dean and P. Dupuis. Splitting for rare event simulation: A large deviation approach to design and analysis. *Stoch. Proc. Appl.*, 119(2):562–587, 2009.

[10] C. Dellago and P. Bolhuis. Transition path sampling and other advanced simulation techniques for rare events. In C. Holm and K. Kremer, editors, *Advanced Computer Simulation Approaches for Soft Matter Sciences III*, volume 221 of *Advances in Polymer Science*, pages 167–233. Springer, Berlin, 2009.

[11] C. Dellago, P. G. Bolhuis, and P. L. Geissler. Transition path sampling. *Adv. Chem. Phys.*, 123:1–78, 2003.

[12] C. Dellago, P. G. Bolhuis, and P. L. Geissler. Transition path sampling methods. In M. Ferrario, G. Ciccotti, and K. Binder, editors, *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, volume 703 of *Lecture Notes in Physics*. Springer, New York, 2006.

[13] B. Derrida. Non-equilibrium steady states: Fluctuations and large deviations of the density and of the current. *J. Stat. Mech.*, 2007(07):P07023, 2007.

[14] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russell, and F. Toomey. Entropy of ATM traffic streams: A tool for estimating QoS parameters. *IEEE J. Select. Areas Comm.*, 13(6):981–990, 1995.

[15] K. Duffy and A. P. Metcalfe. The large deviations of estimating rate functions. *J. Appl. Prob.*, 42(1):267–274, 2005.

[16] P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics. John Wiley, New York, 1997.

[17] W. E, W. Ren, and E. Vanden-Eijnden. String method for the study of rare events. *Phys. Rev. B*, 66(5):052301, 2002.

[18] R. S. Ellis. Large deviations for a general class of random vectors. *Ann. Prob.*, 12(1):1–12, 1984.

[19] R. S. Ellis. *Entropy, Large Deviations, and Statistical Mechanics*. Springer, New York, 1985.

[20] R. S. Ellis. The theory of large deviations: From Boltzmann's 1877 calculation to equilibrium macrostates in 2D turbulence. *Physica D*, 133:106–136, 1999.

[21] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*, volume 260 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, New York, 1984.

[22] J. Gärtner. On large deviations from the invariant measure. *Th. Prob. Appl.*, 22:24–39, 1977.

[23] C. Giardina, J. Kurchan, and L. Peliti. Direct evaluation of large-deviation functions. *Phys. Rev. Lett.*, 96(12):120603, 2006.

[24] M. Gorissen, J. Hooyberghs, and C. Vanderzande. Density-matrix renormalization-group study of current and activity fluctuations near nonequilibrium phase transitions. *Phys. Rev. E*, 79:020101, 2009.

[25] M. Gorissen and C. Vanderzande. Finite size scaling of current fluctuations in the totally asymmetric exclusion process. *J. Phys. A: Math. Theor.*, 44(11):115005, 2011.

[26] R. Graham. Macroscopic potentials, bifurcations and noise in dissipative systems. In F. Moss and P. V. E. McClintock, editors, *Noise in*

Nonlinear Dynamical Systems, volume 1, pages 225–278, Cambridge, 1989. Cambridge University Press.

[27] G. Grimmett and D. Stirzaker. *Probability and Random Processes.* Oxford University Press, Oxford, 2001.

[28] R. J. Harris, A. Rákos, and G. M. Schütz. Current fluctuations in the zero-range process with open boundaries. *J. Stat. Mech.*, 2005(08):P08003, 2005.

[29] R. J. Harris, A. Rákos, and G. M. Schütz. Breakdown of Gallavotti-Cohen symmetry for stochastic dynamics. *Europhys. Lett.*, 75:227–233, 2006.

[30] R. J. Harris and G. M. Schütz. Fluctuation theorems for stochastic dynamics. *J. Stat. Mech.*, 2007(07):P07020, 2007.

[31] W. Krauth. *Statistical Mechanics: Algorithms and Computations.* Oxford Master Series in Statistical, Computational, and Theoretical Physics. Oxford University Press, Oxford, 2006.

[32] V. Lecomte and J. Tailleur. A numerical approach to large deviations in continuous time. *J. Stat. Mech.*, 2007(03):P03004, 2007.

[33] P. L'Ecuyer, V. Demers, and B. Tuffin. Rare events, splitting, and quasi-Monte Carlo. *ACM Trans. Model. Comput. Simul.*, 17(2):1049–3301, 2007.

[34] P. L'Ecuyer, F. Le Gland, P. Lezaud, and B. Tuffin. Splitting techniques. In G. Rubino and B. Tuffin, editors, *Rare Event Simulation Using Monte Carlo Methods*, pages 39–62. Wiley, New York, 2009.

[35] M. Löwe. Iterated large deviations. *Stat. Prob. Lett.*, 26(3):219–223, 1996.

[36] P. Metzner, C. Schutte, and E. Vanden Eijnden. Illustration of transition path theory on a collection of simple examples. *J. Chem. Phys.*, 125(8):084110, 2006.

[37] J. Morio, R. Pastel, and F. Le Gland. An overview of importance splitting for rare event simulation. *Eur. J. Phys.*, 31(5), 2010.

[38] A. Rákos and R. J. Harris. On the range of validity of the fluctuation theorem for stochastic Markovian dynamics. *J. Stat. Mech.*, 2008(05):P05005, 2008.

[39] D. Ruelle. *Chance and Chaos*. Penguin Books, London, 1993.

[40] D. Ruelle. *Thermodynamic Formalism*. Cambridge University Press, Cambridge, 2nd edition, 2004.

[41] J. S. Sadowsky. On Monte Carlo estimation of large deviations probabilities. *Ann. Appl. Prob.*, 6(2):399–422, 1996.

[42] J. S. Sadowsky and J. A. Bucklew. Large deviations theory techniques in Monte Carlo simulation. In E. A. MacNair, K. J. Musselman, and P. Heidelberger, editors, *Proceedings of the 1989 Winter Simulation Conference*, pages 505–513, New York, 1989. ACM.

[43] J. S. Sadowsky and J. A. Bucklew. On large deviations theory and asymptotically efficient Monte Carlo estimation. *IEEE Trans. Info. Th.*, 36:579–588, 1990.

[44] I. N. Sanov. On the probability of large deviations of random variables. In *Select. Transl. Math. Statist. and Probability, Vol. 1*, pages 213–244. Inst. Math. Statist. and Amer. Math. Soc., Providence, R.I., 1961.

[45] H. Spohn. *Large Scale Dynamics of Interacting Particles*. Springer Verlag, Heidelberg, 1991.

[46] J. Tailleur. *Grandes déviations, physique statistique et systèmes dynamiques*. PhD thesis, Université Pierre et Marie Curie, Paris, 2007.

[47] J. Tailleur and V. Lecomte. Simulation of large deviation functions using population dynamics. In J. Marro, P. L. Garrido, and P. I. Hurtado, editors, *Modeling and Simulation of New Materials: Proceedings of Modeling and Simulation of New Materials*, volume 1091, pages 212–219, Melville, NY, 2009. AIP.

[48] H. Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478(1-3):1–69, 2009.

[49] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland, Amsterdam, 1992.

[50] R. van Zon and E. G. D. Cohen. Stationary and transient work-fluctuation theorems for a dragged Brownian particle. *Phys. Rev. E*, 67:046102, 2003.

[51] E. Vanden-Eijnden. Transition path theory. In M. Ferrario, G. Ciccotti, and K. Binder, editors, *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, volume 703 of *Lecture Notes in Physics*, pages 453–493. Springer, 2006.

[52] S. R. S. Varadhan. Asymptotic probabilities and differential equations. *Comm. Pure Appl. Math.*, 19:261–286, 1966.

Errata

I list below some errors contained in the published version of these notes and in the versions (v1 and v2) previously posted on the arxiv.

The numbering of the equations in the published version is different from the version posted on the arxiv. The corrections listed here refer to the arxiv version.

I thank Bernhard Altaner and Andreas Engel for reporting errors. I also thank Vivien Lecomte for a discussion that led me to spot the errors of Sections 4.4 and 4.5.

Corrections to arxiv version 1

- p. 15, in Eq. (2.35): The correct result is $\alpha^{sn}(1-\alpha)^{(1-s)n}$. [Noted by B. Altaner]

- p. 16, in Ex. 2.7.8: "$\lambda(k)$ has a continuous derivative" added.

- p. 20, in Eq. (3.15): The limit $T \to 0$ should be $T \to \infty$. [Noted by B. Altaner]

- p. 25, in Ex. 3.6.2: $Y \in \{-1, +1\}$ with $P(Y = 1) = P(Y = -1) = 1/2$ instead of Bernoulli.

- p. 25, in Ex. 3.6.5: The RVs are now specified to be Gaussian RVs.

Corrections to arxiv version 2

- p. 1: In the abstract: missing 's' to 'introduce'.

- p. 16, in Ex. 2.7.8: The conditions of $\lambda(k)$ stated in the exercise are again wrong: $\lambda(k)$ must be differentiable and must not have affine parts to ensure that $\lambda'(k) = s$ has a unique solution.

- p. 20, in Eq. (3.17): The tilted generator is missing an identity matrix; it should read
$$\tilde{G}_k = G + kx'\delta_{x,x'}.$$
[Noted by A. Engel] Part of Section 3.3 and Eq. (3.17) was rewritten to emphasize where this identity matrix term is coming from.

- p. 27, in Eq. (3.40): The generator should have the off diagonal elements exchanged. Idem for the matrix shown in Eq. (3.45). [Noted by A. Engel]

These notes represent probability vectors as columns vectors. Therefore, the columns, not the rows, of stochastic matrices must sum to one, which means that the columns, not the rows, of stochastic generators must sum to zero.

- p. 27, in Ex. 3.6.13: It should be mentioned that $f(x_i, x_{i+1})$ is chosen as $f(x_i, x_{i+1}) = 1 - \delta_{x_i, x_{i+1}}$. [Noted by A. Engel]

- p. 36, in Section 4.4: Although the tilted pdf $p_k(\boldsymbol{\omega})$ of a Markov chain has a product structure reminiscent of a Markov chain, it cannot be described as a Markov chain, contrary to what is written. In particular, the matrix shown in Eq. (4.24) is not a proper stochastic matrix with columns summing to one. For that, we should define the matrix

$$\pi_k(x'|x) = \frac{e^{kx'}\pi(x'|x)}{W(k|x)}, \qquad W(k|x) = \int e^{kx'}\pi(x'|x)\,dx'.$$

However, the pdf of $\boldsymbol{\omega}$ constructed from this transition matrix is different from the tilted joint pdf $p_k(\boldsymbol{\omega})$. Section 4.4 was rewritten to correct this error.

This section is not entirely wrong: the Metropolis sampling can still be done at the level of p_k by generating IID sequences $\boldsymbol{\omega}^{(j)}$, as now described.

- p. 37, in Section 4.5: The error of Section 4.4 is repeated here at the level of generators: the G_k shown in Eq. (4.25) is not a proper stochastic generator with "columns" summing to zero. Section 4.5 was completely rewritten to correct this error.

- p. 42, Ex. 4.7.14: This exercise refers to the tilted joint pdf $p_k(\boldsymbol{\omega})$; there is no tilted transition matrix here. The exercise was merged with the next one to correct this error.