Stochastic Tools for Mathematics and Science

Alexandre J. Chorin and Ole H. Hald

Prefaces

Preface to the Second Edition

In preparing the second edition we have tried to improve and clarify the presentation, guided in part by the many comments we have received, and also to make the various arguments more precise, as far as we could while keeping this book short and introductory.

There are many dozens of small changes and corrections. The more substantial changes from the first edition include: a completely rewritten discussion of renormalization, and significant revisions of the sections on prediction for stationary processes, Markov chain Monte Carlo, turbulence, and branching random motion. We have added a discussion of Feynman diagrams to the section on Wiener integrals, a discussion of fixed points to the section on the central limit theorem, a discussion of perfect gases and the equivalence of ensembles to the section on entropy and equilibrium. There are new figures, new exercises, and new references.

We are grateful to the many people who have talked with us or written to us with comments and suggestions for improvement. We are also grateful to Valerie Heatlie for her patient help in putting the revised manuscript together.

Alexandre J. Chorin
Ole H. Hald
Berkeley, California
March, 2009

Preface to the First Edition

This book started out as a set of lecture notes for a first-year graduate course on the "stochastic methods of applied mathematics" at the Department of Mathematics of the University of California at Berkeley. The course was started when the department asked a group of its former students who had gone into nonacademic jobs, in national labs and industry, what they actually did in their jobs, and found that most of them did stochastic things that had not appeared anywhere in our graduate course lineup; over the years the course changed as a result of the comments and requests of the students, who have turned out to be a mix of mathematics students and students from the sciences and engineering. The course has not endeavored to present a full, rigorous theory of probability and its applications, but rather to provide mathematics students with some inkling of the many beautiful applications of probability, as well as introduce the nonmathematical students to the general ideas behind methods and tools they already use. We hope that the book too can accomplish these tasks.

We have simplified the mathematical explanations as much as we could everywhere we could. On the other hand, we have not tried to present applications in any detail either. The book is meant to be an introduction, hopefully an easily accessible one, to the topics on which it touches.

The chapters in the book cover some background material on least squares and Fourier series, basic probability (with Monte Carlo methods, Bayes' theorem, and some ideas about estimation), some applications of Brownian motion, stationary stochastic processes (the Khinchin theorem, an application to turbulence, prediction for time series and data assimilation), equilibrium statistical mechanics (including Markov chain Monte Carlo), and time-dependent statistical mechanics (including optimal prediction). The leitmotif of the book is conditional expectation (introduced in a drastically simplified way) and its uses in approximation, prediction, and renormalization. All topics touched upon come with immediate applications; there is an unusual emphasis on time-dependent statistical mechanics and the Mori-Zwanzig formalism, in accordance with our interests and as well as our convictions. Each chapter is followed by references; it is, of course, hopeless to try to provide a full bibliography of all the topics included here; the bibliographies are simply lists of books and papers we have actually used in preparing notes and should be seen as acknowledgments as well as suggestions for further reading in the spirit of the text.

We thank Dr. David Bernstein, Dr. Maria Kourkina-Cameron, and Professor Panagiotis Stinis, who wrote down and corrected the notes on which this book is based and then edited the result; the book would not have existed without them. We are profoundly indebted to many wonderful collaborators on the topics covered in this book, in particular Professor G.I. Barenblatt, Dr. Anton Kast, Professor Raz Kupferman, and Professor Panagiotis Stinis, as well as Dr. John Barber, Dr. Alexander Gottlieb, Dr. Peter Graf, Dr. Eugene Ingerman, Dr. Paul Krause, Professor Doron Levy, Professor Kevin Lin, Dr. Paul Okunev, Dr. Benjamin Seibold, and Professor Mayya Tokman; we have learned from all of them (but obviously not enough) and greatly enjoyed their friendly collaboration. We also thank the students in the Math 220 classes at the University of California, Berkeley, and Math 280 at the University of California, Davis, for their comments, corrections, and patience, and in particular Ms. K. Schwarz, who corrected errors and obscurities. We are deeply grateful to Ms. Valerie Heatlie, who performed the nearly-Sisyphean task of preparing the various typescripts with unflagging attention and good will. Finally, we are thankful to the US Department of Energy and the National Science Foundation for their generous support of our endeavors over the years.

Alexandre J. Chorin
Ole H. Hald
Berkeley, California
September, 2005

Contents

Prefaces v

Contents ix

Chapter 1. Preliminaries 1
1.1. Least Squares Approximation 1
1.2. Orthonormal Bases 6
1.3. Fourier Series 9
1.4. Fourier Transform 12
1.5. Exercises 16
1.6. Bibliography 17

Chapter 2. Probability 19
2.1. Definitions 19
2.2. Expected Values and Moments 22
2.3. Monte Carlo Methods 28
2.4. Parametric Estimation 32
2.5. The Central Limit Theorem 34
2.6. Conditional Probability and Conditional Expectation 37
2.7. Bayes' Theorem 41
2.8. Exercises 43
2.9. Bibliography 45

Chapter 3. Brownian Motion 47
3.1. Definition of Brownian Motion 47
3.2. Brownian Motion and the Heat Equation 49
3.3. Solution of the Heat Equation by Random Walks 50
3.4. The Wiener Measure 54
3.5. Heat Equation with Potential 56
3.6. Physicists' Notation for Wiener Measure 64
3.7. Another Connection Between Brownian Motion and the Heat Equation 66
3.8. First Discussion of the Langevin Equation 68
3.9. Solution of a Nonlinear Differential Equation by Branching Brownian Motion 73

3.10. A Brief Introduction to Stochastic ODEs 75
3.11. Exercises 77
3.12. Bibliography 80

Chapter 4. Stationary Stochastic Processes 83
4.1. Weak Definition of a Stochastic Process 83
4.2. Covariance and Spectrum 86
4.3. Scaling and the Inertial Spectrum of Turbulence 88
4.4. Random Measures and Random Fourier Transforms 91
4.5. Prediction for Stationary Stochastic Processes 96
4.6. Data Assimilation 101
4.7. Exercises 104
4.8. Bibliography 106

Chapter 5. Statistical Mechanics 109
5.1. Mechanics 109
5.2. Statistical Mechanics 112
5.3. Entropy and Equilibrium 115
5.4. The Ising Model 119
5.5. Markov Chain Monte Carlo 121
5.6. Renormalization 126
5.7. Exercises 131
5.8. Bibliography 134

Chapter 6. Time-Dependent Statistical Mechanics 135
6.1. More on the Langevin Equation 135
6.2. A Coupled System of Harmonic Oscillators 138
6.3. Mathematical Addenda 140
6.4. The Mori-Zwanzig Formalism 145
6.5. More on Fluctuation-Dissipation Theorems 150
6.6. Scale Separation and Weak Coupling 152
6.7. Long Memory and the t-Model 153
6.8. Exercises 156
6.9. Bibliography 157

Index 161
Contents

CHAPTER 1

Preliminaries

1.1. Least Squares Approximation

Let V be a vector space with vectors $u, v, w, \ldots$ and scalars $\alpha, \beta, \ldots$. The space V is an inner product space if one has defined a function $(\cdot, \cdot)$ from $V \times V$ to the reals (if the vector space is real) or to the complex (if V is complex) such that for all $u, v \in V$ and all scalars α, the following conditions hold:

$$\begin{aligned} (u, v) &= \overline{(v, u)}, \\ (u + v, w) &= (u, w) + (v, w), \\ (\alpha u, v) &= \alpha(u, v), \\ (v, v) &\geq 0, \\ (v, v) &= 0 \Leftrightarrow v = 0, \end{aligned} \tag{1.1}$$

where the overbar denotes the complex conjugate. Two elements u, v such that $(u, v) = 0$ are said to be orthogonal.

The most familiar inner product space is $\mathbb{R}^n$ with the Euclidean inner product. If $u = (u_1, u_2, \ldots, u_n)$ and $v = (v_1, v_2, \ldots, v_n)$, then

$$(u, v) = \sum_{i=1}^{n} u_i v_i.$$

Another inner product space is $C[0, 1]$, the space of continuous functions on $[0, 1]$, with $(f, g) = \int_0^1 f(x)g(x)\,dx$.

When you have an inner product, you can define a norm, the "L_2 norm", by

$$\|v\| = \sqrt{(v, v)}.$$

This has the following properties, which can be deduced from the properties of the inner product:

$$\begin{aligned} \|\alpha v\| &= |\alpha|\|v\|, \\ \|v\| &\geq 0, \\ \|v\| &= 0 \Leftrightarrow v = 0, \\ \|u+v\| &\leq \|u\| + \|v\|. \end{aligned}$$

The last, called the triangle inequality, follows from the Schwarz inequality

$$|(u,v)| \leq \|u\|\|v\|.$$

In addition to these three properties, common to all norms, the L_2 norm has the "parallelogram property" (so called because it is a property of parallelograms in plane geometry)

$$\|u+v\|^2 + \|u-v\|^2 = 2(\|u\|^2 + \|v\|^2),$$

which can be verified by expanding the inner products.

Let $\{u_n\}$ be a sequence in V.

DEFINITION. A sequence $\{u_n\}$ is said to converge to $\hat{u} \in V$ if $\|u_n - \hat{u}\| \to 0$ as $n \to \infty$ (i.e., for any $\epsilon > 0$, there exists some $N \in \mathbb{N}$ such that $n > N$ implies $\|u_n - \hat{u}\| < \epsilon$).

DEFINITION. A sequence $\{u_n\}$ is a Cauchy sequence if given $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $m, n > N$ $\|u_n - u_m\| < \epsilon$.

A sequence that converges is a Cauchy sequence, although the converse is not necessarily true. If the converse is true for all Cauchy sequences in a given inner product space, then the space is called complete. All of the spaces we work with from now on are complete. Examples are $\mathbb{R}^n$, $\mathbb{C}^n$, L_2.

A few more definitions from real analysis:

DEFINITION. An open ball centered at x with radius $r > 0$ is the set $B_r(x) = \{u : \|u - x\| < r\}$.

DEFINITION. A set S is open if for all $x \in S$, there exists an open ball $B_r(x)$ such that $B_r(x) \subset S$.

DEFINITION. A set S is closed if every convergent sequence $\{u_n\}$ such that $u_n \in S$ for all n converges to an element of S.

An example of a closed set is the closed interval $[0,1] \subset \mathbb{R}$. An example of an open set is the open interval $(0,1) \subset \mathbb{R}$. The complement of an open set is closed, and the complement of a closed set is open. The empty set is both open and closed and so is $\mathbb{R}^n$.

Given a set S and some point b outside of S, we want to determine under what conditions there is a point $\hat{b} \in S$ closest to b. Let $d(b,S) = \inf_{x\in S} \|x-b\|$ be the distance from b to S. The quantity on the right of this definition is the greatest lower bound of the set of numbers $\|x-b\|$, and its existence is guaranteed by the properties of the real number system. What is not guaranteed in advance, and must be proved here, is the existence of an element $\hat{b}$ that satisfies $\|\hat{b}-b\| = d(b,S)$. To see the issue, take $S = (0,1) \subset \mathbb{R}$ and $b = 2$; then $d(b,S) = 1$, yet there is no point $\hat{b} \in (0,1)$ such that $\|\hat{b}-2\| = 1$.

THEOREM 1.1. *If S is a closed linear subspace of V and b is an element of V, then there exists $\hat{b} \in S$ such that $\|\hat{b}-b\| = d(b,S)$.*

PROOF. There exists a sequence of elements $\{u_n\} \subset S$ such that $\|b-u_n\| \to d(b,S)$ by definition of the greatest lower bound. We now show that this sequence is a Cauchy sequence.

From the parallelogram law we have

$$\left\|\frac{1}{2}(b-u_m)\right\|^2 + \left\|\frac{1}{2}(b-u_n)\right\|^2 = \frac{1}{2}\left\|b-\frac{1}{2}(u_n+u_m)\right\|^2 + \frac{1}{8}\|u_n-u_m\|^2. \tag{1.2}$$

S is a vector space; therefore,

$$\frac{1}{2}(u_n+u_m) \in S \Rightarrow \left\|b-\frac{1}{2}(u_n+u_m)\right\|^2 \geq d^2(b,S).$$

Then since $\|b-u_n\| \to d(b,S)$, we have

$$\left\|\frac{1}{2}(b-u_n)\right\|^2 \to \frac{1}{4}d^2(b,S).$$

From (1.2),

$$\|u_n-u_m\| \to 0,$$

and thus $\{u_n\}$ is a Cauchy sequence by definition; our space is complete and therefore this sequence converges to an element $\hat{b}$ in this space. $\hat{b}$ is in V because V is closed. Consequently

$$\|\hat{b}-b\| = \lim \|u_n-b\| = d(b,S).$$

■

We now wish to describe further the relation between b and $\hat{b}$.

THEOREM 1.2. *Let S be a closed linear subspace of V, let x be any element of S, b any element of V, and $\hat{b}$ an element of S closest to b. Then*

$$(x - \hat{b}, b - \hat{b}) = 0.$$

PROOF. If $x = \hat{b}$ we are done. Else set

$$\theta(x - \hat{b}) - (b - \hat{b}) = \theta x + (1 - \theta)\hat{b} - b = y - b.$$

Since y is in S and $\|y - b\| \geq \|\hat{b} - b\|$, we have

$$\begin{aligned}\|\theta(x - \hat{b}) - (b - \hat{b})\|^2 &= \theta^2\|x - \hat{b}\|^2 - 2\theta(x - \hat{b},\, b - \hat{b}) + \|b - \hat{b}\|^2 \\ &\geq \|b - \hat{b}\|^2.\end{aligned}$$

Thus $\theta^2\|x - \hat{b}\|^2 - 2\theta(x - \hat{b},\, b - \hat{b}) \geq 0$ for all θ. The left hand side attains its minimum value when $\theta = (x - \hat{b},\, b - \hat{b})/\|x - \hat{b}\|^2$ in which case $-(x - \hat{b},\, b - \hat{b})^2/\|x - \hat{b}\|^2 \geq 0$. This implies that $(x - \hat{b},\, b - \hat{b}) = 0$. ■

THEOREM 1.3. *$(b - \hat{b})$ is orthogonal to x for all $x \in S$.*

PROOF. By Theorem 1.2, $(x - \hat{b}, b - \hat{b}) = 0$ for all $x \in S$. When $x = 0$ we have $(\hat{b}, b - \hat{b}) = 0$. Thus $(x, b - \hat{b}) = 0$ for all x in S. ■

COROLLARY 1.4. *If S is a closed linear subspace, then $\hat{b}$ is unique.*

PROOF. Let $b = \hat{b} + n = \hat{b}_1 + n_1$, where n is orthogonal to $\hat{b}$ and n_1 is orthogonal to $\hat{b}_1$. Therefore,

$$\begin{aligned}\hat{b} - \hat{b}_1 \in S &\Rightarrow (\hat{b} - \hat{b}_1, n_1 - n) = 0 \\ &\Rightarrow (\hat{b} - \hat{b}_1, \hat{b} - \hat{b}_1) = 0 \\ &\Rightarrow \hat{b} = \hat{b}_1.\end{aligned}$$

■

One can think of $\hat{b}$ as the orthogonal projection of b on S and write $\hat{b} = \mathbb{P}b$, where the projection $\mathbb{P}$ is defined by the foregoing discussion.

We will now give a few applications of the above results.

EXAMPLE. Consider a matrix equation $Ax = b$, where A is an $m \times n$ matrix and $m > n$. This kind of problem arises when one tries to fit a large set of data by a simple model. Assume that the columns of A are linearly independent. Under what conditions does the system have a solution? To clarify ideas, consider the 3×2 case:

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}.$$

Let A_1 denote the first column vector of A, A_2 the second column vector, etc. In this case,

$$A_1 = \begin{bmatrix} a_{11} \\ a_{21} \\ a_{31} \end{bmatrix}, \quad A_2 = \begin{bmatrix} a_{12} \\ a_{22} \\ a_{32} \end{bmatrix}.$$

If $Ax = b$ has a solution, then one can express b as a linear combination of $A_1, A_2, \dots A_m$; for example, in the 3×2 case, $x_1 A_1 + x_2 A_2 = b$. If b does not lie in the column space of A (the set of all linear combinations of the columns of A), then the problem has no solution. It is often reasonable to replace the unsolvable problem by the solvable problem $A\hat{x} = \hat{b}$, where $\hat{b}$ is as close as possible to b and yet does lie in the column space of A. We know from the foregoing that the "best $\hat{b}$" is such that $b - \hat{b}$ is orthogonal to the column space of A. This is enforced by the m equations

$$(A_1, \hat{b} - b) = 0, \quad (A_2, \hat{b} - b) = 0, \quad \dots \quad , (A_m, \hat{b} - b) = 0.$$

Since $\hat{b} = A\hat{x}$, we obtain the equation

$$A^T(A\hat{x} - b) = 0 \quad \Rightarrow \quad \hat{x} = (A^T A)^{-1} A^T b.$$

One application of the above is to "fit" a line to a set of points on the Euclidean plane. Given a set of points, $(x_1, y_1), (x_2, y_2), \dots , (x_n, y_n)$ that come from some experiment and that we believe would lie on a straight line if it were not for experimental error, what is the line that "best approximates" these points? We hope that if it were not for the errors, we would have $y_i = ax_i + b$ for all i and for some fixed a and b; so we seek to solve a system of equations

$$\begin{bmatrix} x_1 & 1 \\ \vdots & \vdots \\ x_n & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

Example. Consider the system of equations given by $Ax = b$, where A is an $n \times m$ matrix and $n < m$ (there are more unknowns than equations). The system has infinitely many solutions. Suppose you want the solution of smallest norm; this problem arises when one tries to find the most likely solution to an underdetermined problem.

Before solving this problem, we need some preliminaries.

Definition. $S \subset V$ is an affine subspace if $S = \{y : y = x + c, c \neq 0, x \in X\}$, where X is a closed linear subspace of V. Note that S is not a linear subspace.

LEMMA 1.5. *If S is an affine subspace and $b' \notin S$, then there exists $\hat{x} \in X$ such that $d(b', S) = \|\hat{x} + c - b'\|$. Furthermore, $\hat{x} - (b' - c)$ is orthogonal to x for all $x \in X$. (Note that here we use b' instead of b, to avoid confusion with the system's right-hand side.)*

PROOF. We have $S = \{y : y = x + c, c \neq 0, x \in X\}$, where X is a closed linear subspace of V. Now,

$$\begin{aligned} d(b', S) &= \inf_{y \in S} \|y - b'\| = \inf_{x \in X} \|x + c - b'\| \\ &= \inf_{x \in X} \|x - (b' - c)\| = d(b' - c, X) \\ &= \|\hat{x} - (b' - c)\| = \|\hat{x} + c - b'\|. \end{aligned}$$

The point $\hat{x} \in X$ exists since X is a closed linear subspace. It follows from Theorem 1.3 that $\hat{x} - (b' - c)$ is orthogonal to X. Note that the distance between S and b' is the same as that between X and $b' - c$. ■

From the proof above, we see that $\hat{x} + c$ is the element of S closest to b'. For the case $b' = 0$, we find that $\hat{x} + c$ is orthogonal to X.

Now we return to the problem of finding the "smallest" solution of an underdetermined problem. Assume A has "maximal rank"; that is, m of the column vectors of A are linearly independent. We can write the solutions of the system as $x = x_0 + z$, where x_0 is a particular solution and z is a solution of the homogeneous system $Az = 0$. So the solutions of the system $Ax = b$ form an affine subspace. As a result, if we want to find the solution with the smallest norm (i.e., closest to the origin) we need to find the element of this affine subspace closest to $b' = 0$. From the above, we see that such an element must satisfy two properties. First, it has to be an element of the affine subspace (i.e., a solution to the system $Ax = b$) and second, it has to be orthogonal to the linear subspace X, which is the null space of A (the set of solutions of $Az = 0$). Now consider $x' = A^T(AA^T)^{-1}b$; this vector lies in the affine subspace of the solutions of $Ax = b$, as one can check by multiplying it by A. Furthermore, it is orthogonal to every vector in the space of solutions of $Az = 0$ because $(A^T(AA^T)^{-1}b, z) = ((AA^T)^{-1}b, Az) = 0$. This is enough to make x' the unique solution of our problem.

1.2. Orthonormal Bases

The problem presented in the previous section, of finding an element in a closed linear space that is closest to a vector outside the space, lies in the framework of approximation theory, where we are given a function (or a vector) and try to find an approximation to it as a linear combination of given functions (or vectors). This is done by requiring that the norm of the error (difference between the given

function and the approximation) be minimized. In what follows, we shall find coefficients for this optimal linear combination.

DEFINITION. Let S be a linear vector space. A collection of m vectors $\{u_i\}_{i=1}^m$ belonging to S are linearly independent if and only if $\lambda_1 u_1 + \cdots + \lambda_m u_m = 0$ implies $\lambda_1 = \lambda_2 = \cdots = \lambda_m = 0$.

DEFINITION. Let S be a linear vector space. A collection $\{u_i\}_{i=1}^m$ of vectors belonging to S is called a basis of S if $\{u_i\}$ are linearly independent and any vector in S can be written as a linear combination of them.

Note that the number of elements of a basis can be finite or infinite depending on the space.

THEOREM 1.6. *Let S be an m-dimensional linear inner-product space with m finite. Then any collection of m linearly independent vectors of S is a basis.*

DEFINITION. A set of vectors $\{e_i\}_{i=1}^m$ is orthonormal if the vectors are mutually orthogonal and each has unit length (i.e., $(e_i, e_j) = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise).

The set of all the linear combinations of the vectors $\{u_i\}$ is called the span of $\{u_i\}$ and is written as $\text{Span}\{u_1, u_2, \ldots, u_m\}$.

Suppose we are given a set of vectors $\{e_i\}_{i=1}^m$ that are an orthonormal basis for a subspace S of a real vector space. If b is an element outside the space, we want to find the element $\hat{b} \in S$, where $\hat{b} = \sum_{i=1}^m c_i e_i$ such that $\|b - \sum_{i=1}^m c_i e_i\|$ is minimized. Specifically, we have

$$\begin{aligned}
\left\| b - \sum_{i=1}^m c_i e_i \right\|^2 &= \left(b - \sum_{i=1}^m c_i e_i \, , b - \sum_{j=1}^m c_j e_j \right) \\
&= (b,b) - 2\sum_{i=1}^m c_i (b, e_i) + \left(\sum_{i=1}^m c_i e_i \, , \sum_{j=1}^m c_j e_j \right) \\
&= (b,b) - 2\sum_{i=1}^m c_i (b, e_i) + \sum_{i,j=1}^m c_i c_j (e_i, e_j) \\
&= (b,b) - 2\sum_{i=1}^m c_i (b, e_i) + \sum_{i=1}^m c_i^2 \\
&= \|b\|^2 - \sum_{i=1}^m (b, e_i)^2 + \sum_{i=1}^m (c_i - (b, e_i))^2,
\end{aligned}$$

where we have used the orthonormality of the e_i to simplify the expression. As is readily seen, the norm of the error is a minimum when

$c_i = (b, e_i), i = 1 \dots m$, so that $\hat{b}$ is the projection of b onto S. It is easy to check that $b - \hat{b}$ is orthogonal to any element in S. Also, we see that the following inequality, called Bessel's inequality, holds:

$$\sum_{i=1}^{m} (b, e_i)^2 \leq \|b\|^2.$$

When the basis is not orthonormal, steps similar to the above yield

$$\begin{aligned}
\left\| b - \sum_{i=1}^{m} c_i g_i \right\|^2 &= \left(b - \sum_{i=1}^{m} c_i g_i, b - \sum_{j=1}^{m} c_j g_j \right) \\
&= (b, b) - 2\sum_{i=1}^{m} c_i (b, g_i) + \left(\sum_{i=1}^{m} c_i g_i, \sum_{j=1}^{m} c_j g_j \right) \\
&= (b, b) - 2\sum_{i=1}^{m} c_i (b, g_i) + \sum_{i,j=1}^{m} c_i c_j (g_i, g_j).
\end{aligned}$$

If we differentiate the last expression with respect to c_i and set the derivatives equal to zero, we get

$$Gc = r,$$

where G is the matrix with entries $g_{ij} = (g_i, g_j)$, $c = (c_1, \dots, c_m)^T$, and $r = ((g_1, b), \dots, (g_m, b))^T$. This system can be ill-conditioned so that its numerical solution presents a problem. The question that arises is how to find, given a set of vectors, a new set that is orthonormal. This is done through the Gram-Schmidt process, which we now describe.

Let $\{u_i\}_{i=1}^m$ be a basis of a linear subspace. The following algorithm will give an orthonormal set of vectors $e_1, e_2, \dots, e_m$ such that $\text{Span}\{e_1, e_2, \dots, e_m\} = \text{Span}\{u_1, u_2, \dots, u_m\}$.

1. Normalize u_1 (i.e., let $e_1 = u_1/\|u_1\|$).
2. We want a vector e_2 that is orthonormal to e_1. In other words we look for a vector e_2 satisfying $(e_2, e_1) = 0$ and $\|e_2\| = 1$. Take $e_2 = u_2 - (u_2, e_1)e_1$ and then normalize.
3. In general, e_j is found recursively by taking

$$e_j = u_j - \sum_{i=1}^{j-1} (u_j, e_i) e_i$$

and normalizing.

The Gram-Schmidt process can be implemented numerically very efficiently. The solution of the recursion above is equivalent to finding

$e_1, e_2, \ldots, e_m$, such that the following holds:

$$\begin{aligned} u_1 &= b_{11}e_1, \\ u_2 &= b_{12}e_1 + b_{22}e_2, \\ &\vdots \\ u_m &= b_{1m}e_1 + b_{2m}e_2 + \cdots + b_{mm}e_m; \end{aligned}$$

that is, what we want to do is decompose the matrix U with columns $u_1, u_2, \ldots, u_m$ into a product of two matrices Q and R, where Q has as columns the orthonormal vectors $e_1, e_2, \ldots, e_m$ and R is the matrix

$$R = \begin{bmatrix} b_{11} & b_{12} & \ldots & b_{1m} \\ 0 & b_{22} & \ldots & b_{2m} \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & b_{mm} \end{bmatrix}.$$

This is the well-known QR decomposition, for which there exist very efficient implementations.

1.3. Fourier Series

Let $L_2[0, 2\pi]$ be the space of square integrable functions in $[0, 2\pi]$ (i.e., such that $\int_0^{2\pi} f^2 dx < \infty$). Define the inner product of two functions f and g belonging to this space as $(f, g) = \int_0^{2\pi} fg\, dx$ and the corresponding norm $\|f\| = \sqrt{(f, f)}$. The Fourier series of a function $f(x)$ in this space is defined as

$$f(x) = a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + \sum_{n=1}^{\infty} b_n \sin(nx), \tag{1.3}$$

where

$$\begin{aligned} a_0 &= \frac{1}{2\pi} \int_0^{2\pi} f(x)\, dx, \\ a_n &= \frac{1}{\pi} \int_0^{2\pi} \cos(nx) f(x)\, dx, \\ b_n &= \frac{1}{\pi} \int_0^{2\pi} \sin(nx) f(x)\, dx. \end{aligned}$$

Alternatively, consider the set of functions

$$\left\{ \frac{1}{\sqrt{2\pi}}, \frac{1}{\sqrt{\pi}} \cos(nx), \frac{1}{\sqrt{\pi}} \sin(nx), \ldots \right\}, \quad n = 1, 2, \ldots .$$

This set is orthonormal in $[0, 2\pi]$ and the Fourier series (1.3) can be rewritten as

$$f(x) = \frac{\tilde{a}_0}{\sqrt{2\pi}} + \sum_{n=1}^{\infty} \frac{\tilde{a}_n}{\sqrt{\pi}} \cos(nx) + \sum_{n=1}^{\infty} \frac{\tilde{b}_n}{\sqrt{\pi}} \sin(nx). \tag{1.4}$$

with

$$\begin{aligned} \tilde{a}_0 &= \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} f(x)\, dx, \\ \tilde{a}_n &= \frac{1}{\sqrt{\pi}} \int_0^{2\pi} \cos(nx) f(x)\, dx, \\ \tilde{b}_n &= \frac{1}{\sqrt{\pi}} \int_0^{2\pi} \sin(nx) f(x)\, dx. \end{aligned}$$

For any function in $L_2[0, 2\pi]$ (the set of square integrable functions on $[0, 2\pi]$) the series (1.4) converges to f in the L_2 norm; i.e., let

$$S_0 = \frac{\tilde{a}_0}{\sqrt{2\pi}}, \quad S_n = \frac{\tilde{a}_0}{\sqrt{2\pi}} + \sum_{m=1}^{n} \frac{\tilde{a}_m}{\sqrt{\pi}} \cos mx + \sum_{m=1}^{n} \frac{\tilde{b}_m}{\sqrt{\pi}} \sin mx \quad (\text{for } n \geq 1)$$

Then we have $\|S_n - f\| \to 0$ as $n \to \infty$.

For any finite truncation of the series (1.4), we have

$$\tilde{a}_0^2 + \sum_{i=1}^{n} \left(\tilde{a}_i^2 + \tilde{b}_i^2 \right) \leq \|f\|^2. \tag{1.5}$$

This is the Bessel inequality, which becomes an equality (Parseval equality) as $n \to \infty$.

The above series (1.4) can be rewritten in complex notation. Recall that

$$\cos(kx) = \frac{e^{ikx} + e^{-ikx}}{2}, \quad \sin(kx) = \frac{e^{ikx} - e^{-ikx}}{2i}. \tag{1.6}$$

After substitution of (1.6) into (1.4) and collection of terms, the Fourier series becomes

$$f(x) = \sum_{k=-\infty}^{\infty} \frac{c_k}{\sqrt{2\pi}} e^{ikx},$$

where f is now complex. (Note that f will be real if for $k \geq 0$, we have $c_{-k} = \overline{c_k}$.) Consider a vector space with complex scalars and introduce an inner product that satisfy axioms (1.1) and define the

norm $\|u\| = \sqrt{(u,u)}$. For the special case where the inner product is given by

$$(u,v) = \int_0^{2\pi} u(x)\bar{v}(x)\,dx,$$

the functions $(2\pi)^{-1/2}\,e^{ikx}$ with $k = 0, \pm 1, \pm 2, \ldots$ form an orthonormal set with respect to this inner product. Then the complex Fourier series of a complex function $f(x)$ is written as

$$f(x) = \sum_{k=-\infty}^{\infty} c_k \frac{1}{\sqrt{2\pi}} e^{ikx}, \quad c_k = \left(f(x), \frac{e^{ikx}}{\sqrt{2\pi}} \right).$$

Let $f(x)$ and $g(x)$ be two functions with Fourier series given respectively by

$$f(x) = \sum_{k=-\infty}^{\infty} \frac{a_k}{\sqrt{2\pi}} e^{ikx},$$
$$g(x) = \sum_{k=-\infty}^{\infty} \frac{b_k}{\sqrt{2\pi}} e^{ikx}.$$

Then for their inner product, we have

$$(f,g) = \int_0^{2\pi} f(x)\bar{g}(x)dx = \int_0^{2\pi} \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} \frac{a_k \bar{b}_l}{2\pi} e^{i(k-l)x} = \sum_{k=-\infty}^{\infty} a_k \bar{b}_k$$

(this is known as Parseval's identity), and for their ordinary product, we have

$$f(x)g(x) = \sum_{k=-\infty}^{\infty} \frac{c_k}{\sqrt{2\pi}} e^{ikx},$$

where

$$\begin{aligned} c_k &= \int_0^{2\pi} \left(\sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \frac{a_n b_m}{2\pi} e^{i(n+m)x} \right) \frac{e^{-ikx}}{\sqrt{2\pi}}\,dx \\ &= \frac{1}{\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} a_n b_m \delta(n+m-k) \\ &= \frac{1}{\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} a_n b_{k-n} = \frac{1}{\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} a_{k-n} b_n. \end{aligned}$$

1.4. Fourier Transform

Consider the space of periodic functions defined on the interval $[-\tau/2, \tau/2]$. The functions $\tau^{-1/2}\exp(2\pi ikx/\tau)$ are an orthonormal basis for this space. For a function $f(x)$ in this space we have

$$f(x) = \sum_{k=-\infty}^{\infty} c_k e_k(x), \quad c_k = (f, e_k(x)),$$

where

$$e_k(x) = \frac{\exp(2\pi ikx/\tau)}{\sqrt{\tau}}$$

and

$$c_k = (f, e_k) = \int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} f(x)\overline{e_k}(x)\, dx.$$

Substituting the expression for the coefficient in the series, we find

$$\begin{aligned} f(x) &= \sum_{k=-\infty}^{\infty} \left(\int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} f(s) \frac{\exp(-2\pi iks/\tau)}{\sqrt{\tau}}\, ds \right) \frac{\exp(2\pi ikx/\tau)}{\sqrt{\tau}} \\ &= \sum_{k=-\infty}^{\infty} \frac{1}{\tau} \left(\int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} f(s)\exp(-2\pi iks/\tau)\, ds \right) \exp(2\pi ikx/\tau). \end{aligned}$$

Define

$$\hat{f}(l) = \int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} f(s)e^{-ils}\, ds.$$

Then the quantity in parentheses above becomes $\hat{f}(l = 2\pi k/\tau)$ and we have

$$f(x) = \sum_{k=-\infty}^{\infty} \frac{1}{\tau}\hat{f}(2\pi k/\tau)\exp(2\pi ikx/\tau). \tag{1.7}$$

Pick τ large and assume that the function f tends to zero at $\pm\infty$ fast enough so that $\hat{f}$ is well defined and that the limit $\tau \to \infty$ is well defined. Write $\Delta = 1/\tau$. From (1.7) we have

$$f(x) = \sum_{k=-\infty}^{\infty} \Delta\hat{f}(2\pi k\Delta)\exp(2\pi ik\Delta x).$$

As $\Delta \to 0$, this becomes

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(2\pi t)\exp(2\pi itx)\, dt,$$

where we have replaced $k\Delta$ by the continuous variable t. By the change of variables $2\pi t = l$, this becomes

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(l) e^{ilx}\, dl.$$

Collecting results, we have

$$\hat{f}(l) = \int_{-\infty}^{\infty} f(s) e^{-ils}\, ds,$$
$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(l) e^{ilx}\, dl.$$

The last two expressions are the Fourier transform and the inverse Fourier transform, respectively. There is no universal agreement on where the quantity 2π that accompanies the Fourier transform should be. It can be split between the Fourier transform and its inverse as long as the product remains 2π. In what follows, we use the splitting

$$\hat{f}(l) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(s) e^{-ils}\, ds,$$
$$f(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(l) e^{ilx}\, dl.$$

Instead of $L_2[0, 2\pi]$, now our space of functions is $L_2(\mathbb{R})$ (i.e., the space of square integrable functions on the real line).

Consider two functions $u(x)$ and $v(x)$ with Fourier series given respectively by $\sum a_k \exp(ikx)/\sqrt{2\pi}$ and $\sum b_k \exp(ikx)/\sqrt{2\pi}$. Then, as we saw above the Fourier coefficients for their product are

$$c_k = \frac{1}{\sqrt{2\pi}} \sum_{k'=-\infty}^{\infty} a_{k'} b_{k-k'}.$$

We now consider what this formula becomes as we go to the Fourier transform; for two functions f and g with Fourier transforms $\hat{f}$ and $\hat{g}$,

we have

$$\begin{aligned}
\widehat{fg}(k) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x)g(x)e^{-ikx}dx \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(k')e^{ik'x}dk'\, g(x)\, e^{-ikx}dx \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(k') \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(x)e^{-i(k-k')x}dx\, dk' \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(k')\hat{g}(k-k')dk' \\
&= \frac{1}{\sqrt{2\pi}} (\hat{f} * \hat{g})(k),
\end{aligned}$$

where $*$ stands for "convolution." This means that up to a constant, the Fourier transform of a product of two functions equals the convolution of the Fourier transforms of the two functions.

Another useful property of the Fourier transform concerns the transform of the convolution of two functions. Assuming f and g are bounded, continuous, and integrable, the following result holds for their convolution $h(x) = (f * g)(x)$:

$$\begin{aligned}
\widehat{(f * g)}(k) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f(\xi)g(x-\xi)d\xi \right) e^{-ikx}\, dx \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\xi)e^{-i\xi x}g(x-\xi)e^{-ik(x-\xi)}dx\, d\xi \\
&= \sqrt{2\pi} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(\xi)e^{-ik\xi} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(y)e^{-iky}dy\, d\xi \\
&= \sqrt{2\pi}\hat{f}(k)\hat{g}(k).
\end{aligned}$$

We have proved that, up to a constant, the Fourier transform of a convolution of two functions is the product of the Fourier transforms of the functions.

In addition, Parseval's equality carries over to the Fourier transform and we have $\|f\|^2 = \|\hat{f}\|^2$, where $\|\cdot\|$ is the L_2 norm on $\mathbb{R}$. This is a

special case ($f = g$) of the following identity

$$\begin{aligned}(f,g) &= \int_{-\infty}^{\infty} f(x)\overline{g(x)}\,dx \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(\xi) e^{i\xi x} d\xi\, \overline{g(x)}\,dx \\ &= \int_{-\infty}^{\infty} \hat{f}(\xi) \frac{1}{\sqrt{2\pi}} \overline{\int_{-\infty}^{\infty} g(x)\, e^{-i\xi x} dx}\, d\xi \\ &= \int_{-\infty}^{\infty} \hat{f}(\xi)\overline{\hat{g}(\xi)} d\xi = (\hat{f},\hat{g}).\end{aligned}$$

Futhermore, consider a function f and its Fourier transform $\hat{f}$. Then for the transform of the function $f(x/a)$, we have

$$\widehat{f\left(\frac{x}{a}\right)}(k) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f\left(\frac{x}{a}\right) e^{-ikx}\,dx.$$

By the change of variables $y = x/a$, we find

$$\begin{aligned}\widehat{f\left(\frac{x}{a}\right)}(k) &= \frac{a}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(y) e^{-iaky}\,dy \\ &= a\hat{f}(ak).\end{aligned}$$

Finally, consider the function $f(x) = \exp(-x^2/2t)$, where $t > 0$ is a parameter. For its Fourier transform we have

$$\begin{aligned}\hat{f}(k) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{x^2}{2t}\right) e^{-ikx}\,dx \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left[-\left(\frac{x^2}{2t} + ikx\right)\right] dx.\end{aligned}$$

By completing the square in the exponent we get

$$\begin{aligned}\hat{f}(k) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left[-\left(\frac{x}{\sqrt{2t}} + ik\sqrt{\frac{t}{2}}\right)^2 - \frac{tk^2}{2}\right] dx \\ &= \frac{1}{\sqrt{2\pi}} e^{-tk^2/2} \int_{-\infty}^{\infty} \exp\left[-\left(\frac{x}{\sqrt{2t}} + ik\sqrt{\frac{t}{2}}\right)^2\right] dx. \qquad (1.8)\end{aligned}$$

The integral in the last expression can be evaluated by a change of variables, but we have to justify that such a change of variables is legitimate. To do that, we quote a result from complex analysis.

LEMMA 1.7. *Let $\phi(z)$ be an analytic function in the strip $|y| < b$ and suppose that $\phi(z)$ satisfies the inequality $|\phi(x+iy)| \le \Phi(x)$ in the strip where $\Phi(x) \ge 0$ is a function such that $\lim_{|x|\to\infty} \Phi(x) = 0$ and $\int_{-\infty}^{\infty} \Phi(x)\,dx < \infty$. Then the value of the integral $\int_{-\infty}^{\infty} \phi(x+iy)\,dx$ is independent of the point $y \in (-b, b)$.*

The integrand in (1.8) satisfies the hypotheses of the lemma and so we are allowed to perform the change of variables

$$y = \frac{x}{\sqrt{2t}} + ik\sqrt{\frac{t}{2}}.$$

Thus, (1.8) becomes

$$\begin{aligned}\hat{f}(k) &= \frac{1}{\sqrt{2\pi}} e^{-tk^2/2} \int_{-\infty}^{\infty} \exp(-y^2)\sqrt{2t}\,dy \\ &= \frac{1}{\sqrt{2\pi}} e^{-tk^2/2} \sqrt{2t\pi} \\ &= \sqrt{t}\, e^{-tk^2/2}.\end{aligned}$$

By setting $t = 1$, we see in particular that the function $f(x) = \exp(-x^2/2)$ is invariant under the Fourier transform.

1.5. Exercises

1. Find the polynomial of degree less than or equal to 2 that best approximates the function $f(x) = e^{-x}$ in the interval $[0, 1]$ in the L_2 sense.
2. Find the Fourier coefficients $\hat{u}_k$ of the function $u(x)$ defined by

$$u(x) = \begin{cases} x, & 0 \le x < \pi \\ x - 2\pi, & \pi \le x \le 2\pi. \end{cases}$$

   Check that $|k\hat{u}(k)| \to$ a constant as $|k| \to \infty$.
3. Find the Fourier transform of the function $e^{-|x|}$.
4. Find the point in the plane $x + y + z = 1$ closest to $(0, 0, 0)$. Note that this plane is not a linear space, and explain how our standard theorem applies.

5. Let $x = (x_1, x_2, \dots)$ and $b = (b_1, b_2, \dots)$ be vectors with complex entries and define $\|x\|^2 = \sum x_i \overline{x_i}$, where $\overline{x_i}$ is the complex conjugate of x_i. Show that the minimum of $\|x - \lambda b\|$ can be found by differentiated with respect to $\overline{\lambda}$, and treating $\lambda, \bar{\lambda}$ as independent.

6. Denote the Fourier transform by F, so that the Fourier transform of a function g is Fg. A function g is an eigenvector of F with an eigenvalue λ if $Fg = \lambda g$ (we have seen that $e^{-x^2/2}$ is such an eigenfunction with eigenvalue 1). Show that F can have no eigenvalues other than $\pm 1, \pm i$. (Hint: what do you get when you calculate $F^4 g$?

1.6. Bibliography

[1] H. Dym and H. McKean, *Fourier Series and Integrals*, Academic Press, New York, 1972.
[2] G. Golub and F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1983.
[3] A. Kolmogorov and S. Fomin, *Elements of the Theory of Functions and Real Analysis*, Dover, New York, 2000.
[4] P. Lax, *Linear Algebra*, Wiley, New York, 1997.

CHAPTER 2

Probability

2.1. Definitions

In weather forecasts, one often hears a sentence such as "the probability of rain tomorrow is 50 percent." What does this mean? Something like: "If we look at all possible tomorrows, in half of them there will be rain" or "if we make the experiment of observing tomorrow, there is a quantifiable chance of having rain tomorrow, and somehow or other this chance was quantified as being 1/2." To make sense of this, we formalize the notions of experimental outcome, event, and probability.

Suppose that you make an experiment and imagine all possible outcomes.

DEFINITION. A sample space Ω is the space of all possible outcomes of an experiment.

For example, if the experiment is "waiting until tomorrow, and then observing the weather," Ω is the set of all possible weathers tomorrow. There can be many weathers, some differing only in details we cannot observe and with many features we cannot describe precisely.

Suppose you set up a thermometer in downtown Berkeley and decide you will measure the temperature tomorrow at noon. The set of possible weathers for which the temperature is between 65 and 70 degrees is an "event," an outcome which is specified precisely and about which we can think mathematically. An event is subset of Ω, a set of outcomes, a subset of all possible outcomes Ω, that corresponds to a well-defined property that can be measured.

DEFINITION. An event is a subset of Ω.

The set of events we are able to consider is denoted by $\mathcal{B}$; it is a set of subsets of Ω. We require that $\mathcal{B}$ (the collection of events) be a σ-algebra; that is, $\mathcal{B}$ must satisfy the following axioms:

1. $\emptyset \in \mathcal{B}$ and $\Omega \in \mathcal{B}$ ($\emptyset$ is the empty set).
2. If $B \in \mathcal{B}$, then $CB \in \mathcal{B}$ (CB is the complement of B in Ω).
3. If $\mathcal{A} = \{A_1, A_2, \dots, A_n, \dots\}$ is a finite or countable collection in $\mathcal{B}$, then any union of the elements of $\mathcal{A}$ is in $\mathcal{B}$.

It follows from these axioms that any intersection of a countable number of elements of $\mathcal{B}$ also belongs to $\mathcal{B}$.

Consider the tosses of a die. In this case, $\Omega = \{1, 2, 3, 4, 5, 6\}$.

1. If we are only interested in whether something happened or not, we may consider a set of events
$$\mathcal{B} = \{\{1, 2, 3, 4, 5, 6\}, \emptyset\}.$$
The event $\{1, 2, 3, 4, 5, 6\}$ means "something happened," while the event $\emptyset$ means "nothing happened."
2. If we are interested in whether the outcome is odd or even, then we may choose
$$\mathcal{B} = \{\{1, 3, 5\}, \{2, 4, 6\}, \{1, 2, 3, 4, 5, 6\}, \emptyset\}.$$
3. If we are interested in which particular number appears, then $\mathcal{B}$ is the set of all subsets of Ω; $\mathcal{B}$ is generated by $\{\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{6\}\}$.

Observe that $\mathcal{B}$ in case (1) is the smallest σ-algebra on the sample space (in the sense of having fewest elements), while $\mathcal{B}$ in case (3) is the largest.

Definition. A probability measure $P(A)$ is a function $P : \mathcal{B} \to \mathbb{R}$ defined on the sets $A \in \mathcal{B}$ such that:

1. $P(\Omega) = 1$.
2. $0 \leq P \leq 1$.
3. If $\{A_1, A_2, \dots, A_n, \dots\}$ is a finite or countable collection of events such that $A_i \in \mathcal{B}$ and $A_i \cap A_j = \emptyset$ for $i \neq j$, then $P(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} P(A_i)$ (the probability of the simultaneous occurrence of incompatible events is the sum of the probabilities of the individual events).

Definition. The triple $(\Omega, \mathcal{B}, P)$ is called a probability space.

In brief, the σ-algebra $\mathcal{B}$ defines the objects to which we assign probabilities and P assigns probabilities to the elements of $\mathcal{B}$.

Definition. A random variable $\eta : \Omega \to \mathbb{R}$ is a $\mathcal{B}$-measurable function defined on Ω, where "$\mathcal{B}$-measurable" means that the subset of elements ω in Ω for which $\eta(\omega) \leq x$ is an element of $\mathcal{B}$ for every x. In other words, it is possible to assign a probability to the occurrence of the inequality $\eta \leq x$ for every x.

Loosely speaking, a random variable is a real variable whose numerical values are determined by experiment, with the proviso that it is possible to assign probabilities to the occurrence of the various values.

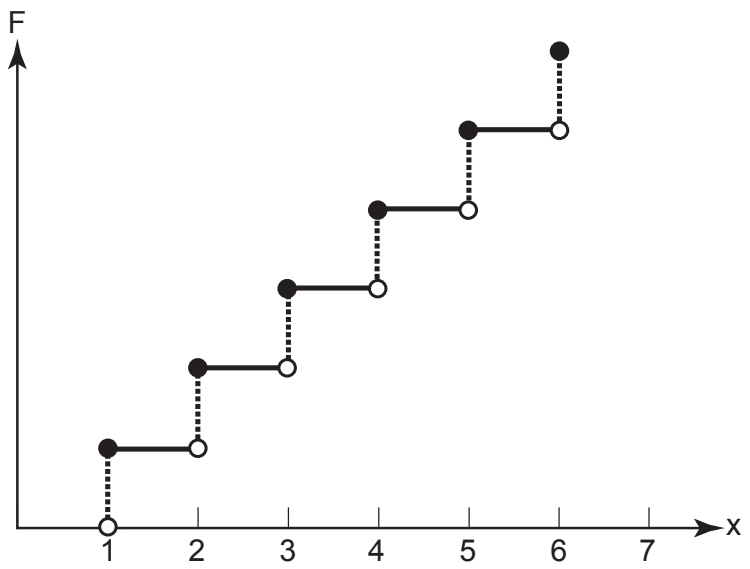

FIGURE 2.1. Probability distribution for a fair six-sided die.

Given a probability measure $P(A)$, the probability distribution function of a random variable η is defined by

$$F_\eta(x) = P\left(\{\omega \in \Omega \mid \eta(\omega) \leq x\}\right) = P(\eta \leq x).$$

The existence of such a function is guaranteed by the definition of a random variable.

Now consider several examples.

EXAMPLE. Let $\mathcal{B} = \{A_1, A_2, A_1 \cup A_2, \emptyset\}$. Let $P(A_1) = P(A_2) = 1/2$. Define a random variable

$$\eta(\omega) = \begin{cases} -1, & \omega \in A_1 \\ +1, & \omega \in A_2. \end{cases}$$

Then

$$F_\eta(x) = \begin{cases} 0, & x < -1 \\ 1/2, & -1 \leq x < 1 \\ 1, & x \geq 1. \end{cases}$$

EXAMPLE. Suppose that we are tossing a die. $\Omega = \{1, 2, 3, 4, 5, 6\}$ and $\eta(\omega) = \omega$. Take $\mathcal{B}$ to be the set of all subsets of Ω. The probability distribution function of η is the one shown in Figure 2.1.

Suppose that Ω is the real line and the range of a random variable η also is the real line (e.g., $\eta(\omega) = \omega$). In this case, one should be sure that the σ-algebra $\mathcal{B}$ is large enough to include all of the sets of the form

$\{\omega \in \Omega \mid \eta(\omega) \leq x\}$. The minimal σ-algebra satisfying this condition is the σ-algebra of the "Borel sets" formed by taking all the possible countable unions and complements of all of the half-open intervals in $\mathbb{R}$ of the form $(a, b]$.

Suppose that $F'_\eta(x)$ exists. Then $f_\eta(x) = F'_\eta(x)$ is the probability density of η. Since $F_\eta(x)$ is nondecreasing, $f_\eta(x) \geq 0$. Obviously,

$$\int_{-\infty}^{\infty} f_\eta(x)dx = F_\eta(\infty) - F_\eta(-\infty) = 1.$$

If $F'_\eta(x)$ exists and is continuous, then

$$P(x < \eta \leq x + dx) = F_\eta(x + dx) - F_\eta(x) = f_\eta(x)\, dx.$$

The following probability density functions (pdfs) are often encountered:

1. Equidistribution density
$$f(x) = \begin{cases} 1, & 0 \leq x \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$
2. Gaussian density
$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-m)^2}{2\sigma^2}\right), \tag{2.1}$$
where m and σ are constants.
3. Exponential density
$$f(x) = \begin{cases} e^{-x}, & x \geq 0 \\ 0, & x < 0. \end{cases}$$

2.2. Expected Values and Moments

Definition. Let $(\Omega, \mathcal{B}, P)$ be a probability space and η be a random variable. Then the expected value, or mean, of the random variable η is defined as the integral of η over Ω with respect to the measure P:

$$E[\eta] = \int_\Omega \eta(\omega)\, dP.$$

In this notation, the symbol dP is a reminder of the measure with respect to which the integral is taken; when there is a need for more specificity, we shall also sometimes write $P(d\omega)$ instead of dP. When Ω is a discrete set, this integral is just the sum of the products of the values of η with the probabilities that η assumes these values.

This definition can be rewritten in another way involving the Stieltjes integral. Let F be a nondecreasing and bounded function. Define the Stieltjes integral of a function $g(x)$ on an interval $[a, b]$ as follows. Let $a = x_0 < x_1 < \cdots < x_{n-1} < x_n = b$, $\Delta_i = x_{i+1} - x_i$, and $x_i^* \in [x_i, x_{i+1}]$. Then

$$\int_a^b g(x)\, dF(x) = \lim_{\Delta_i \to 0} \sum_{i=0}^{n-1} g(x_i^*)(F(x_{i+1}) - F(x_i))$$

(where we have written F instead of F_η for short). Let $x_i^* = x_i = -k + i/2^k$ for $i = 0, 1, \ldots, n = k \cdot 2^{k+1}$, when k is an integer, so that $-k \le x_i \le k$. Define the indicator function χ_B of a set B by $\chi_B(x) = 1$ if $x \in B$, $\chi_B(x) = 0$ if $x \notin B$. Set $\Delta_i = 1/2^k$. The expected value of η is

$$\begin{aligned}
\int_\Omega \eta(d\omega) P(d\omega) &= \int_\Omega \lim_{n\to\infty} \sum_{i=0}^{n-1} x_i \chi_{\{\omega | x_i < \eta \le x_{i+1}\}} P(d\omega) \\
&= \lim_{n\to\infty} \sum_{i=0}^{n-1} x_i\, P\left(\{\omega | x_i < \eta(\omega) \le x_{i+1}\}\right) \\
&= \lim_{n\to\infty} \sum_{i=0}^{n-1} x_i^*\, (F(x_{i+1}) - F(x_i)) \\
&= \lim_{k\to\infty} \int_{-k}^{k} x\, d\, F(x) + O\left(\frac{1}{2^k}\right) \\
&= \int_{-\infty}^{\infty} x\, d\, F(x).
\end{aligned}$$

If η is a random variable, then so is $a\eta$, where a is a constant. If η is a random variable and $g(x)$ is a continuous function defined on the range of η, then $g(\eta)$ is also a random variable, and

$$E[g(\eta)] = \int_{-\infty}^{\infty} g(x)\, dF(x).$$

The special cases

$$E[\eta^n] = \int_{-\infty}^{\infty} x^n\, dF(x)$$

and

$$E[(\eta - E[\eta])^n] = \int_{-\infty}^{\infty} (x - E[\eta])^n\, dF(x)$$

are called the n^{th} moment and the n^{th} centered moment of η, respectively. (Of course, these integrals may fail to converge for some random variables.) The second centered moment is the variance of η.

DEFINITION. The variance $\mathrm{Var}(\eta)$ of the random variable η is

$$\mathrm{Var}(\eta) = E[(\eta - E[\eta])^2]$$

and the standard deviation of η is

$$\sigma = \sqrt{\mathrm{Var}(\eta)}.$$

EXAMPLE. The Gaussian pdf (2.1) has $E[\eta] = m$ and $\mathrm{Var}(\eta) = \sigma^2$.

DEFINITION. Two events A and B are independent if $P(A \cap B) = P(A)P(B)$. Two random variables η_1 and η_2 are independent if the events $\{\omega \in \Omega \mid \eta_1(\omega) \leq x\}$ and $\{\omega \in \Omega \mid \eta_2(\omega) \leq y\}$ are independent for all x and y.

DEFINITION. If η_1 and η_2 are random variables, then the joint distribution function of η_1 and η_2 is defined by

$$F_{\eta_1\eta_2}(x, y) = P(\{\omega \in \Omega \mid \eta_1(\omega) \leq x, \eta_2(\omega) \leq y\}) = P(\eta_1 \leq x, \eta_2 \leq y).$$

If the second mixed derivative $\partial^2 F_{\eta_1\eta_2}(x, y)/\partial x\, \partial y$ exists, it is called the joint probability density of η_1 and η_2 and is denoted by $f_{\eta_1\eta_2}$. In this case,

$$F_{\eta_1\eta_2}(x, y) = \int_{-\infty}^{x} \int_{-\infty}^{y} f_{\eta_1\eta_2}(s, t)\, dt\, ds.$$

Clearly, if η_1 and η_2 are independent, then

$$F_{\eta_1\eta_2}(x, y) = F_{\eta_1}(x) F_{\eta_2}(y)$$

and

$$f_{\eta_1\eta_2}(x, y) = f_{\eta_1}(x) f_{\eta_2}(y).$$

We can view two random variables η_1 and η_2 as a single vector-valued random variable $\eta = (\eta_1, \eta_2) = \eta(\omega)$ for $\omega \in \Omega$. We say that η is measurable if the event $\eta \in S$ with $S \subset \mathbb{R}^2$ is measurable for a suitable family of S's (i.e., the event $Z = \{\omega \in \Omega : \eta(\omega) \in S\} \in \mathcal{B}$, where $\mathcal{B}$ is a σ-algebra on Ω). Suppose that the joint probability distribution function of the two random variables exists and is denoted by $F_{\eta_1\eta_2}(x, y) = P(\eta_1 \leq x, \eta_2 \leq y)$. Note that $F_{\eta_1\eta_2}(x, y) = F_{\eta_2\eta_1}(y, x)$ and $F_{\eta_1\eta_2}(\infty, y) = F_{\eta_2}(y)$. If the joint density exists, then $\int_{-\infty}^{\infty} f_{\eta_1\eta_2}(x, y)\, dx = f_{\eta_2}(y)$.

DEFINITION. The covariance of two random variables η_1 and η_2 is

$$\mathrm{Cov}(\eta_1, \eta_2) = E[(\eta_1 - E[\eta_1])(\eta_2 - E[\eta_2])].$$

If $\mathrm{Cov}(\eta_1, \eta_2) = 0$, then the random variables are uncorrelated. It is in general not true that uncorrelated variables are independent.

EXAMPLE. Let η_1 and η_2 be two random variables with joint probability distribution

$$(\eta_1, \eta_2) = \begin{cases} (\frac{1}{2}, \frac{1}{4}) & \text{with probability } \frac{1}{4} \\ (\frac{1}{2}, -\frac{1}{4}) & \text{with probability } \frac{1}{4} \\ (-\frac{1}{2}, 0) & \text{with probability } \frac{1}{2}. \end{cases}$$

Then we have $E[\eta_1] = 0$, $E[\eta_2] = 0$, and $E[\eta_1\eta_2] = 0$. However, the random variables are not independent because $P\left(\eta_1 = -\frac{1}{2}, \eta_2 = \frac{1}{4}\right) \neq P\left(\eta_1 = -\frac{1}{2}\right) P\left(\eta_2 = \frac{1}{4}\right)$.

Finally, a vector-valued random variable is Gaussian (or, equivalently, a sequence of random variables is jointly Gaussian) if

$$P(x_1 \leq \eta_1 \leq x_1 + dx_1, \ldots, x_n \leq \eta_n \leq x_n + dx_n) = \frac{1}{Z} e^{-\frac{1}{2}(x-m)^T A^{-1}(x-m)}\, dx,$$

where $x = (x_1, x_2, \ldots, x_n)$, $m = (m_1, m_2, \ldots, m_n)$, $dx = dx_1 \cdots dx_n$, and A is a symmetric, positive definite $n \times n$ matrix. The normalization constant Z can be shown to be $Z = (2\pi)^{n/2}|A|^{1/2}$, where $|A|$ is the determinant of A. In the case of jointly Gaussian random variables, the covariance matrix C with entries $C_{ij} = E[(\eta_i - E[\eta_i])(\eta_j - E[\eta_j])]$ equals the matrix A. If $C_{ij} = 0$, then η_i and η_j are uncorrelated. Furthermore, two Gaussian variables that are uncorrelated are also independent.

We now discuss several useful properties of the mathematical expectation E.

LEMMA 2.1. $E[\eta_1 + \eta_2] = E[\eta_1] + E[\eta_2]$.

PROOF. We assume for simplicity that the joint density $f_{\eta_1\eta_2}(x, y)$ exists. Then the density $f_{\eta_1}(x)$ of η_1 is given by

$$f_{\eta_1}(x) = \int_{-\infty}^{\infty} f_{\eta_1\eta_2}(x, y)\, dy$$

and the density $f_{\eta_2}(y)$ of η_2 is given by

$$f_{\eta_2}(y) = \int_{-\infty}^{\infty} f_{\eta_1\eta_2}(x, y)\, dx;$$

therefore

$$\begin{aligned} E[\eta_1 + \eta_2] &= \int (x+y) f_{\eta_1\eta_2}(x,y)\,dx\,dy \\ &= \int x f_{\eta_1\eta_2}(x,y)\,dx\,dy + \int y f_{\eta_1\eta_2}(x,y)\,dx\,dy \\ &= \int x\,dx \int f_{\eta_1\eta_2}(x,y)\,dy + \int y\,dy \int f_{\eta_1\eta_2}(x,y)\,dx \\ &= \int x f_{\eta_1}(x)\,dx + \int y f_{\eta_2}(y)\,dy = E[\eta_1] + E[\eta_2]. \end{aligned}$$

■

LEMMA 2.2. *If η_1 and η_2 are independent random variables, then*

$$\operatorname{Var}[\eta_1 + \eta_2] = \operatorname{Var}[\eta_1] + \operatorname{Var}[\eta_2].$$

PROOF. For simplicity, we assume that η_1 and η_2 have densities with mean zero. Then

$$\begin{aligned} \operatorname{Var}[\eta_1 + \eta_2] &= E[(\eta_1 + \eta_2 - E[\eta_1 + \eta_2])^2] = E[(\eta_1 + \eta_2)^2] \\ &= \int (x+y)^2 f_{\eta_1\eta_2}(x,y)\,dx\,dy \\ &= \int x^2 f_{\eta_1\eta_2}(x,y)\,dx\,dy + \int y^2 f_{\eta_1\eta_2}(x,y)\,dx\,dy \\ &\quad + 2\int xy f_{\eta_1\eta_2}(x,y)\,dx\,dy. \end{aligned}$$

The first two integrals are equal to $\operatorname{Var}(\eta_1)$ and $\operatorname{Var}(\eta_2)$, respectively. The third integral is zero. Indeed, because η_1 and η_2 are independent, $f_{\eta_1\eta_2}(x,y) = f_{\eta_1}(x) f_{\eta_2}(y)$ and

$$\int xy f_{\eta_1\eta_2}(x,y)\,dx\,dy = \int x f_{\eta_1}(x)\,dx \int y f_{\eta_2}(y)\,dy = E[\eta_1]E[\eta_2] = 0.$$

■

Another simple property of the variance is that $\mathrm{Var}(a\eta) = a^2\mathrm{Var}(\eta)$, where a is a constant. Indeed,

$$\begin{aligned}\mathrm{Var}(a\eta) &= \int (ax - E[a\eta])^2 f_\eta(x)\,dx \\ &= \int (ax - aE[\eta])^2 f_\eta(x)\,dx \\ &= a^2 \int (x - E[\eta])^2 f_\eta(x)\,dx \\ &= a^2\mathrm{Var}(\eta).\end{aligned}$$

We now prove a very useful estimate due to Chebyshev.

LEMMA 2.3. *Let η be a random variable. Suppose $g(x)$ is a non-negative, nondecreasing function (i.e., $g(x) \geq 0$ and $a < b \Rightarrow g(a) \leq g(b)$). Then, for any a,*

$$P(\eta \geq a) \leq \frac{E[g(\eta)]}{g(a)}.$$

PROOF.

$$\begin{aligned}E[g(\eta)] = \int_{-\infty}^{\infty} g(x)f(x)\,dx &\geq \int_a^{\infty} g(x)f(x)\,dx \\ &\geq g(a)\int_a^{\infty} f(x)\,dx = g(a)P(\eta \geq a).\end{aligned}$$

■

Suppose η is a non-negative random variable. We define $g(x)$ to be 0 when $x \leq 0$ and x^2 when $x \geq 0$. Let a be any positive number. Then

$$P(\eta \geq a) \leq \frac{E[g(\eta)]}{g(a)} = \frac{E[\eta^2]}{a^2}.$$

Consider now a special case. Let η be a random variable and define $\xi = |\eta - E[\eta]|$. Then we obtain the following inequality:

$$P(|\eta - E[\eta]| \geq a) \leq \frac{\mathrm{Var}(\eta)}{a^2}$$

for any $a > 0$. Now take $a = \sigma k$, where k is an integer. Then

$$P(|\eta - E[\eta]| \geq \sigma k) \leq \frac{\mathrm{Var}(\eta)}{(\sigma k)^2} = \frac{1}{k^2}.$$

In other words, it is very unlikely that η differs from its expected value by more than a few standard deviations.

Suppose $\eta_1, \eta_2, \ldots, \eta_n$ are independent, identically distributed random variables. Let

$$\eta = \frac{1}{n}\sum_{i=1}^{n} \eta_i.$$

Then

$$E[\eta] = E[\eta_1], \quad \mathrm{Var}(\eta) = \frac{1}{n}\mathrm{Var}(\eta_1), \quad \sigma(\eta) = \frac{\sigma(\eta_1)}{\sqrt{n}}.$$

Therefore,

$$P\left(|\eta - E[\eta]| \geq k n^{-1/2}\sigma(\eta_1)\right) \leq \frac{1}{k^2}.$$

This tells us that if we use the average of n independent samples of a given distribution to estimate the mean of the distribution, then the error in our estimates decreases as $1/\sqrt{n}$. This discussion brings the notion of expected value closer to the intuitive, every-day notion of "average."

2.3. Monte Carlo Methods

With Monte Carlo methods, one evaluates a nonrandom quantity as an expected value of a random variable.

A pseudo-random sequence is a computer-generated sequence that cannot be distinguished by simple tests from a random sequence with independent entries, yet is the same each time one runs the appropriate program. For the equidistribution density, number theory allows us to construct the appropriate pseudo-random sequence. Suppose that we want to generate a sequence of This can be done in the following way. Let $F(\eta) = \xi$, where η is the random variable we want to sample and ξ is equidistributed in $[0, 1]$. Take η such that $\eta = F^{-1}(\xi)$ holds (if there are multiple solutions, pick one arbitrarily). Then η will have the desired distribution. To see this, consider the following example. Let η be a random variable with

$$\eta = \begin{cases} \alpha_1 \text{ with probability } p_1 \\ \alpha_2 \text{ with probability } p_2 \\ \alpha_3 \text{ with probability } p_3, \end{cases}$$

where $\sum_{i=1}^{3} p_i = 1$ and $p_i \geq 0$ for $i = 1, 2, 3$. Then $F(\eta) = \xi$ implies

$$\eta = \begin{cases} \alpha_1 \text{ if } \xi \in [0, p_1] \\ \alpha_2 \text{ if } \xi \in (p_1, p_1 + p_2] \\ \alpha_3 \text{ if } \xi \in (p_1 + p_2, 1]. \end{cases}$$

This can be generalized to any countable number of discrete values in the range of η, and since any function can be approximated by a step function, the results hold for any probability distribution function F.

EXAMPLE. Let η be a random variable with the exponential pdf. Then $F(\eta) = \xi$ gives

$$\int_0^\eta e^{-s}\,ds = \xi \implies \eta = -\log(1-\xi).$$

EXAMPLE. If f exists, then by differentiating $\int_{-\infty}^{\eta} f(s)\,ds = \xi$, we get $f(\eta)d\eta = d\xi$. The following algorithm (the "Box-Muller" algorithm) allows us to sample pairs of independent variables with Gaussian densities with zero mean and variance σ^2. Let

$$\eta_1 = \sqrt{-2\sigma^2 \log \xi_1}\cos(2\pi\xi_2),$$
$$\eta_2 = \sqrt{-2\sigma^2 \log \xi_1}\sin(2\pi\xi_2),$$

where ξ_1 and ξ_2 are equidistributed in $[0,1]$; then η_1, η_2 are Gaussian variables with means zero and variances σ^2, as one can see from

$$\left|\begin{vmatrix} \dfrac{\partial\eta_1}{\partial\xi_1} & \dfrac{\partial\eta_1}{\partial\xi_2} \\ \dfrac{\partial\eta_2}{\partial\xi_1} & \dfrac{\partial\eta_2}{\partial\xi_2} \end{vmatrix}^{-1}\right| d\eta_1\,d\eta_2 = d\xi_1\,d\xi_2$$

(the outer vertical lines denote an absolute value) which becomes, with the equations above,

$$\frac{1}{2\pi\sigma^2}\exp\left(-\frac{\eta_1^2+\eta_2^2}{2\sigma^2}\right)d\eta_1\,d\eta_2 = d\xi_1\,d\xi_2.$$

Now we present the Monte Carlo method. Consider the problem of evaluating the integral $I = \int_a^b g(x)f(x)\,dx$, where $f(x) \geq 0$ and $\int_a^b f(x)\,dx = 1$. We have

$$I = \int_a^b g(x)f(x)\,dx = E[g(\eta)],$$

where η is a random variable with pdf $f(x)$. Suppose that we can sample η; that is, make n independent experiments with outcomes $\eta_1, \ldots, \eta_n$. Then, as can be seen from the Chebyshev inequality, we can approximate $E[g(\eta)]$ by

$$E[g(\eta)] \sim \frac{1}{n}\sum_{i=1}^{n} g(\eta_i).$$

The error in this approximation will be of the order of $\sigma(g(\eta))/\sqrt{n}$, where $\sigma(g(\eta))$ is the standard deviation of the variable $g(\eta)$. The integral I is the estimand, $g(\eta)$ is the estimator, and $n^{-1}\sum_{i=1}^{n} g(\eta_i)$ is the estimate. The estimator is unbiased if its expected value is the estimand.

EXAMPLE. Let

$$I = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(x) e^{-x^2/2}\, dx.$$

If η is a Gaussian random variable with mean 0 and variance 1, then

$$I = E[g(\eta)] \sim \frac{1}{n} \sum_{i=1}^{n} g(\eta_i).$$

There are two ways to reduce the error of a Monte Carlo method, as can be seen from the error estimate. One way is to take a larger number of samples. The other way is to reduce the variance of the function $g(\eta)$. One way to reduce the variance is "importance sampling."

We start with an extreme case. Suppose we want to evaluate the integral $I = \int_a^b g(x) f(x)\, dx$ as above. Suppose that the function g is non-negative; then the quantity $q(x)$ given by $q(x) = f(x)g(x)/I$ has the following properties:

$$q(x) \geq 0, \quad \int_a^b q(x)\, dx = 1.$$

Further, suppose we can generate a pseudo-random sequence with pdf $q(x)$. Then we have

$$\int_a^b g(x) f(x)\, dx = I \int_a^b \frac{g(x) f(x)}{I}\, dx = I \int_a^b q(x)\, dx = IE[1],$$

where 1 is the function that takes the value 1 for all samples. Then, the Monte Carlo method has zero error. However we need to know the value of I, which is exactly what we want to compute. If we know the value of the quantity that we want to compute, Monte Carlo can give us the exact result with no error.

However, it is possible to reduce the error of the Monte Carlo method along similar lines without knowing the result we want to compute. Suppose that we can find a function $h(x)$ with the following properties:

1. The integral $I_1 = \int_a^b f(x)h(x)\, dx$ is easily evaluated.
2. $h(x) \geq 0$.
3. We can sample a variable with pdf $f(x)h(x)/I_1$ easily.
4. $g(x)/h(x)$ varies little.

Then we have

$$I = \int_a^b g(x)f(x)\,dx = \int_a^b \frac{g(x)}{h(x)}f(x)h(x)\,dx = I_1 \int_a^b \frac{g(x)}{h(x)}\frac{f(x)h(x)}{I_1}\,dx$$
$$= I_1 E\left[\frac{g}{h}(\eta)\right] \sim \frac{I_1}{n}\sum_{i=1}^{n}\frac{g(\eta_i)}{h(\eta_i)}, \tag{2.2}$$

where η has pdf $f(x)h(x)/I_1$. Since $g(\eta)/h(\eta)$ varies little, its variation and the error will be smaller. The new random variable puts more points where g is large, hence the name of the method "importance sampling"; one puts more samples where g is large, or "important."

Example. Suppose that we want to compute via Monte Carlo the integral $I = \int_0^1 \cos(x/5)e^{-5x}\,dx$. We can do that by applying the basic Monte Carlo formula without any attempt at importance sampling. That would mean sampling n times an equipartitioned variable ξ and then approximating I by

$$I \approx \frac{1}{n}\sum_{i=1}^{n}\cos(\xi_i/5)e^{-5\xi_i},$$

where the ξ_i are the successive independent samples of ξ. However, due to the large variation of the function $\cos(x/5)e^{-5x}$, the corresponding error would be large (the large variation of the function is due to the presence of the factor e^{-5x}). Alternatively, we can perform the Monte Carlo integration using importance sampling. There are different ways of doing that and one of them is as follows. Let $I_1 = \int_0^1 e^{-5x}\,dx = (1-e^{-5})/5$. Then we have

$$I = \int_0^1 \cos(x/5)e^{-5x}\,dx = I_1\int_0^1 \cos(x/5)\frac{e^{-5x}}{I_1}\,dx.$$

Let η be a random variable with pdf

$$f(x) = \begin{cases} \dfrac{e^{-5x}}{I_1}, & 0 \le x \le 1 \\ 0, & \text{elsewhere,} \end{cases}$$

then I can be written as $I = I_1 E[\cos(\eta/5)]$. As can be readily seen, the function $\cos(x/5)$ has smaller variation in the range of integration $[0,1]$ than the previous integrand. In order to perform the Monte Carlo integration, we need to sample the variable η. As shown above, this can be done by solving the equation $\int_0^\eta e^{-5x}/I_1\,dx = \xi$, where ξ is equidistributed in $[0,1]$. An easy calculation gives $\eta = -\frac{1}{5}\log(1-5I_1\xi)$.

We can use this formula to sample η n times and, thus, the Monte Carlo approximation to I will read

$$I \approx \frac{I_1}{n} \sum_{i=1}^{n} \cos(\eta_i/5).$$

2.4. Parametric Estimation

Suppose η is a random variable that someone has sampled and given you the sample $(x_1, x_2, \dots, x_n)$. Now try to guess the pdf. Suppose you know the type of distribution you have, but not the parameters in the distribution. For example, suppose you know that the distribution is Gaussian, but you do not know the mean and the variance.

DEFINITION. Any function of a sample is called a "statistic."

Suppose you want to estimate a parameter θ of the pdf by a statistic $\hat{\theta}(x_1, x_2, \dots, x_n)$.

DEFINITION. The estimate is unbiased if

$$E[\hat{\theta}(\eta_1, \eta_2, \dots, \eta_n)] = \theta$$

(i.e., if, on the average, the estimate is exact).

For example, the sample mean defined by

$$\bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i$$

is an unbiased estimate of the mean, whereas the sample variance

$$\frac{1}{n} \sum_{i=1}^{n} (x_i - \bar{x})^2$$

is not an unbiased estimate of the variance (see the exercises). However, one can check that

$$\frac{1}{n-1} \sum_{i=1}^{n} (x_i - \bar{x})^2$$

is an unbiased estimate of the variance. It is of course desirable that one use unbiased estimators.

We now present a useful method for finding estimators. Suppose you know that the pdf of η that gave you the independent sample $\hat{x} = (x_1, x_2, \dots, x_n)$ is $f(x|\theta)$ (a function of x and of the parameter θ).

What is a good estimate of θ given the sample $\hat{x}$? Suppose you know θ. Then the probability of getting the given sample is proportional to

$$L = \prod_{i=1}^{n} f(x_i \mid \theta).$$

L is called a likelihood function. It is plausible that a good estimate of θ is the one that maximizes L (i.e., which makes the outcome you see as likely as possible). This is the "maximum likelihood estimate." In general, it is easier to maximize $\log L$, which has a maximum at the same value of the argument.

EXAMPLE. Suppose you think that $x_1, x_2, \ldots, x_n$ are independent samples of a Gaussian distribution with mean m and variance σ^2. Then

$$L = \prod_{i=1}^{n} \frac{e^{-(x_i-m)^2/2\sigma^2}}{\sqrt{2\pi\sigma^2}}.$$

Find the maximum of $\log L$:

$$\log L = \sum_{i=1}^{n} \left(-\frac{(x_i - m)^2}{2\sigma^2} - \frac{1}{2}\log 2\pi - \log \sigma \right),$$

$$\frac{\partial \log L}{\partial m} = \sum_{i=1}^{n} \frac{x_i - m}{\sigma^2} = 0.$$

Hence,

$$\sum_{i=1}^{n} x_i - nm = 0,$$

and we get the sample mean as the maximum likelihood estimate of $\hat{m}$:

$$m = \frac{1}{n}\sum_{i=1}^{n} x_i.$$

Similarly,

$$\frac{\partial \log L}{\partial \sigma} = -\frac{n}{\sigma} + \sum_{i=1}^{n} \frac{(x_i - m)^2}{\sigma^3} = 0;$$

hence, the maximum likelihood estimate of the variance of a Gaussian variable is the sample variance (which, as we know, is not unbiased):

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n} (x_i - \hat{m})^2.$$

2.5. The Central Limit Theorem

Suppose that $\eta_1, \eta_2, \dots, \eta_n$ are independent, identically distributed random variables. We can assume without loss of generality that they have mean 0 and variance 1. Suppose the η_i's have a pdf f. Define a new random variable

$$S_n = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \eta_i.$$

What can we say about the pdf of S_n? The answer for this question is given by the following:

THEOREM 2.4 (The Central Limit Theorem). *Let $\eta_1, \eta_2, \dots, \eta_n$ be independent and identically distributed random variables with finite variance and zero mean. Let us also assume for simplicity that* $\mathrm{Var}(\eta_i) = 1$. *Then*

$$S_n = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \eta_i$$

converges weakly to a Gaussian variable with mean 0 and variance 1.

PROOF. We will assume that the η_i have pdf f and that $f^{(n)}$ is the pdf of S_n. We want to show that

$$\lim_{n\to\infty} \int_a^b f^{(n)}(x)\,dx = \int_a^b \frac{e^{-x^2/2}}{\sqrt{2\pi}}\,dx$$

for any a, b. Note that $n^{-1}\sum \eta_i = n^{-1/2}(n^{-1/2}\sum \eta_i)$, where $n^{-1/2}\sum \eta_i$ tends to a Gaussian; thus, the central limit theorem contains information as to how $n^{-1}\sum \eta_i \to 0$ (i.e., for large n, $n^{-1}\sum \eta_i \approx$ Gaussian$/\sqrt{n}$). Suppose η_1 and η_2 are random variables with respective pdfs f_1 and f_2. What is the density of $\eta_1 + \eta_2$? We know that

$$P(\eta_1 + \eta_2 \le x) = F_{\eta_1+\eta_2}(x) = \int\int_{x_1+x_2\le x} f_1(x_1) f_2(x_2)\,dx_1\,dx_2.$$

With the change of variables $x_1 = t$ and $x_1 + x_2 = y$ (note that the Jacobian is 1), we obtain,

$$F_{\eta_1+\eta_2}(x) = \int_{-\infty}^{x} dy \int_{-\infty}^{\infty} f_1(t) f_2(y-t)\,dt.$$

Thus, the density of $\eta_1 + \eta_2 = f_{\eta_1+\eta_2}$ is just $\int f_1(t) f_2(y-t)\,dt = f_1 * f_2$ and, hence, $\hat{f}_{\eta_1+\eta_2} = \sqrt{2\pi}\,\hat{f}_1\hat{f}_2$.

Hence, if we assume that the random variables η_i have the same density function for all i, then $\sum_{i=1}^n \eta_i$ has density $f^{(n)} = f * f * \cdots * f$ (f appears n times), where $*$ is the convolution. Furthermore,

$$P(a < S_n \le b) = P\left(a < \frac{1}{\sqrt{n}} \sum \eta_i \le b\right) = P(\sqrt{n}a < \sum \eta_i \le \sqrt{n}b)$$
$$= \int_{\sqrt{n}a}^{\sqrt{n}b} f^{(n)}(x)\,dx = \int_a^b \sqrt{n} f^{(n)}(y\sqrt{n})\,dy. \quad (2.3)$$

The last step involves the change of variables $y = x/\sqrt{n}$.

What we want to show is that $\int_a^b \sqrt{n} f^{(n)}(y\sqrt{n})\,dy$ converges to

$$\int_a^b \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx.$$

Pick some nice function ϕ and consider

$$I = \int_{-\infty}^{\infty} \sqrt{n} f^{(n)}(x\sqrt{n})\phi(x)\,dx.$$

Let $\check{\phi}(k) = \hat{\phi}(-k)$ be the inverse Fourier transform of ϕ; that is,

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \check{\phi}(k) e^{-ikx}\,dk.$$

Then

$$\begin{aligned} I &= \int_{-\infty}^{\infty} \sqrt{n} f^{(n)}(x\sqrt{n})\phi(x)\,dx \\ &= \int_{-\infty}^{\infty} \sqrt{n} f^{(n)}(x\sqrt{n}) \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \check{\phi}(k)\, e^{-ikx}\,dk\,dx \\ &= \int_{-\infty}^{\infty} \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \sqrt{n} f^{(n)}(x\sqrt{n}) e^{-ikx} dx\right) \check{\phi}(k)\,dk \\ &= \int_{\infty}^{\infty} \widehat{f^{(n)}}\left(k/\sqrt{n}\right) \check{\phi}(k)\,dk \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \left[\sqrt{2\pi}\,\hat{f}\,\left(\frac{k}{\sqrt{n}}\right)\right]^n \check{\phi}(k)\,dk. \end{aligned}$$

Here

$$\hat{f}\left(\frac{k}{\sqrt{n}}\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x)\, e^{-ikx/\sqrt{n}}\,dx,$$

and we used that $\widehat{f * g} = \sqrt{2\pi}\hat{f}\cdot\hat{g}$. Expand $e^{-ikx/\sqrt{n}}$ in a Taylor series:

$$e^{-ikx/\sqrt{n}} = 1 - \frac{ixk}{\sqrt{n}} - \frac{x^2k^2}{2n} + O\left(\frac{1}{n^{3/2}}\right).$$

Recall that

$$\int f(x)\,dx = 1, \quad \int xf(x)\,dx = 0, \quad \int x^2 f(x)\,dx = 1.$$

Hence,

$$\begin{aligned}\hat{f}\left(\frac{k}{\sqrt{n}}\right) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\left(1 - \frac{k^2x^2}{2n} + \cdots\right) f(x)\,dx \\ &= \frac{1}{\sqrt{2\pi}}\left(1 - \frac{k^2}{2n}\right) + \text{small terms.}\end{aligned}$$

Remember that

$$\lim_{n\to\infty}\left(1 - \frac{a}{n}\right)^n = e^{-a}.$$

The contribution of the small terms in the expansion can be shown to be negligible, and we get

$$\lim_{n\to\infty}\left[\sqrt{2\pi}\,\hat{f}\left(\frac{k}{\sqrt{n}}\right)\right]^n = \lim_{n\to\infty}\left(1 - \frac{k^2}{2n} + \text{small}\right)^n = e^{-k^2/2}.$$

Returning to the integral I, we obtain

$$\begin{aligned}I &\to \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-k^2/2}\,\check{\phi}(k)\,dk \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-k^2/2}\left(\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\phi(x)e^{ikx}dx\right)\,dk \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\phi(x)\,dx\left(\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-k^2/2}e^{ikx}\,dk\right) \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\phi(x)e^{-x^2/2}\,dx.\end{aligned}$$

Now, taking ϕ to be a smooth function that approximates

$$\Phi(x) = \begin{cases}1, & a \le x \le b \\ 0, & \text{otherwise,}\end{cases}$$

we get the desired result. ■

It will be useful later to consider the central limit theorem in a slightly different form. Let the random variables η_i for $i = 1, 2, \ldots$ be independent and have each the pdf f, with mean 0 and variance 1 as above, and construct the following sequence of random variables:

$$T_{0,1} = \eta_1, \quad T_{0,2} = \eta_2, \quad T_{0,3} = \eta_2, \ldots \tag{2.4}$$

$$T_{1,1} = \frac{1}{\sqrt{2}}(\eta_1 + \eta_2), T_{1,2} = \frac{1}{\sqrt{2}}(\eta_3 + \eta_4), T_{1,3} = \frac{1}{\sqrt{2}}(\eta_5 + \eta_6), \ldots \tag{2.5}$$

and

$$T_{n+1,1} = \frac{1}{\sqrt{2}}(T_{n,1} + T_{n,2}),\ \ T_{n+1,2} = \frac{1}{\sqrt{2}}(T_{n,3} + T_{n,4}), \ldots \tag{2.6}$$

for $n \geq 1$, where $T_{n,1}, T_{n,2}$ are disjoint sums of 2^n variables in the set. It is easy to see that $T_n = S_{2^n}$, where S_{2^n} are the sums of 2^n of the random variables that appeared in the statement of the central limit theorem. Let the pdf of T_n be f_n with $f_0 = f$; if the pdf of S_n converge to a limit as n tends to infinity, so do the f_n. We have a formula for the pdf of a sum of two variables, and we know that if a variable ξ has the pdf $g(x)$ and a is a a positive constant, then ξ/a has the pdf $ag(ax)$; this yields:

$$f_{n+1}(x) = \sqrt{2}\int_{-\infty}^{+\infty} f_n(t) f_n(\sqrt{2}\,x - t)dt. \tag{2.7}$$

If the f_n converge to a limit f_∞, this equation becomes

$$f_\infty(x) = \sqrt{2}\int_{-\infty}^{+\infty} f_\infty(t) f_\infty(\sqrt{2}\,x - t)dt. \tag{2.8}$$

The central limit theorem says that if the variance of the η_i is finite, this last equation has a solution, which is Gaussian. The iteration (2.6) converges to that solution, and its limit is independent of the starting point f, just like a convergent iterative solution of algebraic equation converges to a limit independent of the starting point.

2.6. Conditional Probability and Conditional Expectation

Suppose we make an experiment and observe that event A has happened, with $P(A) \neq 0$. How does this knowledge affect the probability that another event B also happens? We define the probability of B given A to be

$$P(B|A) = \frac{P(A \cap B)}{P(A)}.$$

If A and B are independent, then $P(A \cap B) = P(A)P(B)$ and so

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{P(A)P(B)}{P(A)} = P(B).$$

If A is fixed and B is any member of $\mathcal{B}$ (i.e., any event), then $P(B|A)$ defines a perfectly good probability measure on $\mathcal{B}$; this is the probability conditional on A:

$$(\Omega, \mathcal{B}, P) \to (\Omega, \mathcal{B}, P(B|A)).$$

Suppose η is a random variable on Ω. Then the average of η given A is

$$E[\eta|A] = \int \eta(\omega) P(d\omega|A).$$

Thus if $\eta = \sum c_i \chi_{B_i}$ then

$$E[\eta|A] = \int \sum c_i \chi_{B_i}(\omega) P(d\omega|A) = \sum c_i \, P(B_i|A).$$

EXAMPLE. Suppose we throw a die. Let η be the value on top. Then

$$E[\eta] = \frac{1}{6} \sum_{i=1}^{6} i = 3.5.$$

Suppose we know that the outcome is odd. Then the probability that the outcome is 1, given this information, is

$$P(\{1\}|\text{outcome is odd}] = \frac{P(\{1\} \cap \{1,3,5\})}{P(\{1,3,5\})} = \frac{1/6}{1/2} = \frac{1}{3};$$

and the average of η given $A = \{1,3,5\}$ is

$$E[\eta|\text{outcome is odd}] = \frac{1}{3}(1+3+5) = 3.$$

The probability of a particular even outcome given A is

$$P(2|A) = P(4|A) = P(6|A) = 0,$$

whereas the total probability of an odd outcome given A is

$$P(1|A) + P(3|A) + P(5|A) = 1.$$

Suppose $Z = \{Z_i\}$ is an at most countable disjoint measurable partition of Ω. This means that the number of Z_i's is finite or countable, each Z_i is an element of $\mathcal{B}$, $\Omega = \bigcup_i Z_i$, and $Z_i \cap Z_j = \emptyset$ if $i \neq j$.

EXAMPLE. $Z = \{A, CA\}$, where A is a measurable subset of Ω and CA is the complement of A.

DEFINITION. Suppose B is an event. Then $\chi_B(\omega)$ is a random variable equal to 1 when $\omega \in B$ and 0 when $\omega \notin B$.

Observe that $E[\chi_B(\omega)] = P(B)$ and $E[\chi_B|A] = P(B|A)$.

DEFINITION. Let $Z = \{Z_i\}$ be a partition of Ω as above. Let η be a random variable and construct the random variable $E[\eta|Z]$ as follows:

$$E[\eta|Z] = \sum_i E[\eta|Z_i] \chi_{Z_i}.$$

This is a function of ω, whose definition depends on the choice of partition Z. In words, we average η over each element Z_i of the partition and then we assign this average to be the value of the variable $E[\eta|Z]$ for all ω in Z_i. If one could think of the elements of ω as people and the values of η as those people's heights, one could then partition the people by ethnic origin and assign an average height to each ethnic group. Given a person, the new variable would assign to that person not his or her real height but the average height of his or her ethnic group.

Note that Z generates a σ-algebra. It is a coarser σ-algebra than $\mathcal{B}$ (i.e., it is contained in $\mathcal{B}$). The variable $E[\eta|Z]$ is the best estimate of the original random variable when the instruments you use to measure the outcomes (which define the σ-algebra generated by Z) are too coarse.

EXAMPLE. Return to the example of the die. Let η be the number on top. Let A be the event that outcome is odd. Let $Z = \{A, CA\}$. Then

$$E[\eta|A] = \frac{1}{3}(1+3+5) = 3,$$

$$E[\eta|CA] = \frac{1}{3}(2+4+6) = 4,$$

and, finally,

$$E[\eta|Z] = 3\chi_A + 4\chi_{CA},$$

where χ_A, χ_{CA} are the indicator functions of the sets A, CA.

We now want to define the notion of conditional expectation of one random variable η given another random variable ξ. For simplicity, we assume at first that ξ takes only finitely many values $\xi_1, \xi_2, \ldots, \xi_n$. Let Z_i be the inverse image of ξ_i (the set of ω such that $\eta(\omega) = \xi_i$). Then $Z = \{Z_1, Z_2, \ldots, Z_n\}$ is a finite disjoint partition of Ω. Thus, we can construct $E[\eta|Z]$ as defined above.

DEFINITION. We define $E[\eta|\xi]$ to be the random variable $E[\eta|Z]$.

We observe that $E[\eta|\xi]$ is a random variable and, at the same time, a function of ξ. Indeed, when ξ has value ξ_i, then $E[\eta|\xi] = E[\eta|Z_i]$; thus, $E[\eta|\xi]$ is a function of ξ. We now show that $E[\eta|\xi]$ is actually the best least square approximation of η by a function of ξ. This property can serve as an alternative definition of conditional expectation.

THEOREM 2.5. *Let $g(\xi)$ be any function of ξ. Then*

$$E[(\eta - E[\eta|\xi])^2] \leq E[(\eta - g(\xi))^2].$$

PROOF. We remind the reader that

$$\int_0^1 (f(x) - c)^2 \, dx$$

is minimized when c is the average of $f(x)$ on $[0, 1]$ (i.e., when $c = \int_0^1 f(x) \, dx$). Similarly, we want to minimize

$$\begin{aligned} E[(\eta - g(\xi))^2] &= \int_\Omega (\eta - g(\xi(\omega))^2 P(d\omega) \\ &= \sum_i P(Z_i) \int_{Z_i} (\eta - g(\xi(\omega))^2 \frac{P(d\omega)}{P(Z_i)}. \end{aligned}$$

Since $g(\xi(\omega)) = g(\xi_i)$ for all ω in Z_i, each of the integrals

$$\int_{Z_i} (\eta - g(\xi(\omega)))^2 P(d\omega)/P(Z_i)$$

is minimized when $g(\xi_i) = E[\eta|Z_i]$ (i.e., when $g(\xi(\omega))$ is the average of η on Z_i). Thus, $E[\eta|\xi]$ is the best least squares approximation of η by a function of ξ. ■

Let $h(\xi)$ be a function of ξ. Then

$$E\left[(\eta - E[\eta|\xi])h(\xi)\right] = 0.$$

To see this, assume $\alpha = E\left[(\eta - E[\eta|\xi])h(\xi)\right] \neq 0$ for some function $h(\xi)$ and set $\epsilon = \alpha/E[(h(\xi))^2]$. Then

$$\begin{aligned} E\left[(\eta - E[\eta|\xi] - \epsilon h(\xi))^2\right] &= E\left[(\eta - E[\eta|\xi])^2\right] \\ &\quad + \epsilon^2 E[(h(\xi))^2] - 2\epsilon \, E\left[(\eta - E[\eta|\xi])h(\xi)\right] \\ &= E\left[(\eta - E[\eta|\xi])^2\right] - \alpha^2/E[(h(\xi))^2]. \end{aligned}$$

But this contradicts Theorem 2.5, so $\alpha = 0$ for all $h(\xi)$. We can give this result a geometric interpretation.

Consider the space of all square integrable random variables. It is a vector space and the functions of ξ form a linear subspace. Let η_1 and η_2 be random variables and define the inner product by

$$(\eta_1, \eta_2) = E[\eta_1 \eta_2].$$

Since $E\left[(\eta - E[\eta|\xi])h(\xi)\right]$ vanishes for all $h(\xi)$, we see that $\eta - E[\eta|\xi]$ is perpendicular to all functions $h(\xi)$. Set $P\eta = E[\eta|\xi]$. Then $\eta = P\eta + (\eta - P\eta)$ with $(\eta - P\eta, P\eta) = 0$, and we can interpret $P\eta$ as the orthogonal projection of η onto the subspace of random variables that are functions of ξ and have finite variance.

We now consider the special case where η and ξ are random variables whose joint density $f_{\eta\xi}$ is known:

$$P(s < \eta \le s + ds, t < \xi \le t + dt) = f_{\eta\xi}(s,t)\, ds\, dt.$$

We want to calculate $E[h(\eta,\xi)|\xi]$, where $g(\eta,\xi)$ is some function of η and ξ. $E[g(\eta,\xi)|\xi]$ is a random variable and a function of ξ. What is this function? Specifically, what is the value of this random variable when $\xi = a$?

To answer this question, we first define a discrete approximation $\hat{\xi}$ to ξ that takes the value $\hat{\xi} = (i + 1/2)h$ when $\xi \in (ih, (i+1)h]$. This happens with probability $\int_{ih}^{(i+1)h} f_\xi(t)\, dt$, where $f_\xi(t)$ is given by

$$f_\xi(t) = \int_{-\infty}^{\infty} f_{\eta\xi}(s,t)\, ds.$$

Now, we replace $E[g(\eta,\xi)|\xi]$ by $E[g(\eta,\xi)|\hat{\xi}]$. (We are committing many mathematical sins here, but sin should be enjoyed.) Suppose we fix an a and pick a value $a_i = (i + 1/2)h$ of $\hat{\xi}$ such that $a \in [ih, (i+1)h]$. Then, dropping the subscripts η, ξ of f, we find

$$\begin{aligned} E[g(\eta,\xi)|\xi]_{\xi=a} &\approx E\left[g(\eta,\xi)|\hat{\xi}\right]_{\hat{\xi}=a_i} \\ &\approx \frac{\int_{-\infty}^{\infty} ds \int_{ih}^{(i+1)h} g(s,t) f(s,t) dt}{\int_{-\infty}^{\infty} h f(s,(i+1/2)h)\, ds} \\ &\to \frac{\int_{-\infty}^{\infty} g(s,a) f(s,a)\, ds}{\int_{-\infty}^{\infty} f(s,a)\, ds} \end{aligned}$$

as $h \to 0$. Thus,

$$E[g(\eta,\xi)|\xi]_{\xi=a} = \frac{\int_{-\infty}^{\infty} g(s,a) f(s,a)\, ds}{\int_{-\infty}^{\infty} f(s,a)\, ds}, \tag{2.9}$$

and

$$E[g(\eta,\xi)|\xi] = \frac{\int_{-\infty}^{\infty} g(s,\xi) f(s,\xi)\, ds}{\int_{-\infty}^{\infty} f(s,\xi)\, ds}. \tag{2.10}$$

This is just what one would expect: $E[g(\eta,\xi)|\xi]$ is the mean of $g(\eta,\xi)$ when we keep the value of ξ fixed but allow η to take any value it wants.

2.7. Bayes' Theorem

Recall the definition of conditional probability:

DEFINITION. Let A and B be two events with $P(A) \neq 0$ and $P(B) \neq 0$. The conditional probability of B given A, $P(B|A)$, is

$$P(B|A) = \frac{P(A \cap B)}{P(A)}. \tag{2.11}$$

Similarly, the conditional probability of A given B is

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \tag{2.12}$$

Combining (2.11) and (2.12), we get Bayes' theorem:

THEOREM 2.6. *Let A and B be two events with $P(A) \neq 0$ and $P(B) \neq 0$. Then*

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{2.13}$$

Suppose $Z = \{Z_j\}, j = 1, 2, \ldots$, is a finite or countable partition of the sample space Ω as above; then for the probability $P(A)$ of an event A, we have

$$P(A) = \sum_j P(A \cap Z_j) = \sum_j \frac{P(A \cap Z_j)}{P(A)} P(A) = \sum_j P(Z_j|A)P(A).$$

Suppose that $P(Z_j) \neq 0$ for all j. Then we can also rewrite $P(A)$ as

$$P(A) = \sum_j P(A \cap Z_j) = \sum_j \frac{P(A \cap Z_j)}{P(Z_j)} P(Z_j) = \sum_j P(A|Z_j)P(Z_j). \tag{2.14}$$

Using Bayes' theorem (2.13) for the events A and Z_j and expressing $P(A)$ by (2.14), we get

$$P(Z_j|A) = \frac{P(A|Z_j)P(Z_j)}{\sum_i P(A|Z_i)P(Z_i)}. \tag{2.15}$$

This is the second form of Bayes' theorem. We can use the second form to address the following question: Suppose we have an experimental sample and we know that we have sampled some probability distribution that depends on a parameter θ. We do not know what value θ takes in the case at hand, but we have an idea a priori (i.e., a "prior" idea) that the set of possible values of θ can be viewed as a random variable with a density g_{old} (the "prior" distribution). Now that we have made an experiment and obtained data, we should be able to learn from these data how to improve the prior ideas and obtain a

new density g_{new}, the "posterior" density, which improves the "prior" density in light of the data. We show how to do it in an example.

EXAMPLE. Let η_1 and η_2 be two independent and identically distributed random variables with

$$\eta_1, \eta_2 = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1-p. \end{cases}$$

For the sum $\eta_1 + \eta_2$, we can deduce

$$\eta_1 + \eta_2 = \begin{cases} 2 & \text{with probability } p^2 \\ 1 & \text{with probability } 2p(1-p) \\ 0 & \text{with probability } (1-p)^2. \end{cases}$$

Suppose that before the experiment we thought that the parameter p had the value $p = 1/4$ with probability $1/4$ and the value $p = 1/2$ with probability $3/4$. This is the "prior distribution." Now, we make an experiment and find $\eta_1 + \eta_2 = 1$. We want to use the second form of Bayes' theorem (2.15) to see how the experiment affects our beliefs about the distribution of the parameter p. To do that, let A be the event that $\eta_1 + \eta_2 = 1$, let Z_1 be the event that $p = 1/4$, and let Z_2 be the event that $p = 1/2$ (note that $Z_1 \cup Z_2 = \Omega$). Then we have

$$\begin{aligned} P(Z_1|A) &= \frac{P(A|Z_1)P(Z_1)}{\sum_j P(A|Z_j)P(Z_j)} \\ &= \frac{\left(2 \times \frac{1}{4} \times \frac{3}{4}\right) \times \frac{1}{4}}{\left(2 \times \frac{1}{4} \times \frac{3}{4}\right) \times \frac{1}{4} + \left(2 \times \frac{1}{2} \times \frac{1}{2}\right) \times \frac{3}{4}} \\ &= \frac{1}{5}, \end{aligned}$$

as opposed to $1/4$ a priori. In words, the probability that $p = 1/4$ now that we know the outcome of the experiment equals the ratio of the product of the probability that the outcome is what it is when $p = 1/4$ and the prior probability that $p = 1/4$, normalized by the sum of the probabilities of the outcome we have for the various prior probabilities.

Of course, the taint of possible error in the prior ideas has not completely disappeared.

2.8. Exercises

1. Write a Monte Carlo program to evaluate
$$\int_0^1 \frac{e^{-\sqrt{1-x^2}}}{\sqrt{x}}\,dx$$
(you may have to do importance sampling to get anything at all). As you sum the samples, estimate the variance and the error; by making several runs with different sample sizes, check that your error estimates are realistic. Estimate the number of samples needed to get an error $\leq \epsilon$, where $\epsilon > 0$ is a tolerance. Find the value of the integral with an error $\leq 1\%$.

2. Let $H_0, H_1, H_2, \ldots$ be Hermite polynomials: H_n is a polynomial of degree n with
$$\int_{-\infty}^{+\infty} \frac{H_m H_n e^{-x^2}}{\sqrt{\pi}}\,dx = \delta_{nm}.$$
Suppose you want to evaluate $I = \pi^{-1/2}\int_{-\infty}^{+\infty} g(x)e^{-x^2}dx$, where g is a given function; let ξ be a Gaussian variable with mean 0 and variance 1/2. Show that for all a, b, $I = E[g(\xi) + aH_1(\xi) + bH_2(\xi)]$. However, the variance of the estimator is not independent of a, b. What values should a, b take to yield an estimator of least variance?

3. Let η be a random variable that takes the value 1/2 with probability 1/2 and the value $-1/2$ also with probability 1/2. Let $\Xi_n = (\sum_1^n \eta_i)/\sqrt{n}$, where the η_i are independent variables with the same distribution as η. Find the values that Ξ_n can take and their probabilities for $n = 3, 6, 9$, and plot their histograms together with the pdf of the limit of Ξ_n as $n \to \infty$.

4. Let η be again a random variable that takes the value 1/2 with probability 1/2 and the value $-1/2$ with probability 1/2, and form the variable $\Xi_n^\alpha = (\sum_1^n \eta_i)/n^\alpha$, where $\alpha \geq 0$. What can you say about the limit of the pdf of Ξ_n^α as $n \to \infty$? What equation does this limit satisfy?

5. Check the derivation of Box-Muller sampling scheme.

6. An exponential variable with parameter λ has the density $f = \lambda e^{-\lambda x}$, $\lambda > 0$. If you are given n independent samples of such a variable, how do you find the maximum likelihood estimate of λ?

7. Suppose you have n independent samples $x_1, ..., x_n$ of a random variable η; show that if $m = n^{-1}\sum_{i=1}^n x_i$, then $n^{-1}\sum(x_i - m)^2$ is not an unbiased estimate of the variance of η, whereas $(n-1)^{-1}\sum(x_i - m)^2$

is an unbiased estimate. Suggestion: To see what is going on, try first the case $n = 2$. Note: These calculations are independent of any assumed form for the density.

8. Consider a vector-valued Gaussian random variable ξ_1, ξ_2, with pdf

$$f(x_1, x_2) = f(x) = \frac{\alpha}{2\pi} \exp(-(x - m, A(x - m)/2)),$$

where A is a symmetric positive definite matrix. Show that $\alpha = \sqrt{\det A}$ and $A = C^{-1}$, where C is the covariance matrix.

9. Let $(\Omega, \mathcal{B}, P)$ be a probability space, A an event with $P(A) > 0$, and $P_A(B) = P(B|A)$ for every event B in $\mathcal{B}$. Show that $(\Omega, \mathcal{B}, P_A)$ satisfies all the axioms for a probability space.

10. let η_1, η_2 be two random variables with joint pdf

$$Z^{-1} \exp(-x_1^2 - x_2^2 - x_1^2 x_2^2),$$

where Z is a normalization constant. Evaluate $E[\eta_1 \eta_2^2 | \eta_1]$.

11. Let η be the number that comes up when you throw a die. Evaluate $E[\eta|(\eta - 3)^2]$ (you may want to present it as a table of its values for different values of $(\eta - 3)^2$).

12. Suppose η is a random variable such that $\eta = 0$ with probability p and $\eta = 1$ with probability $1 - p$. Suppose your prior distribution of p is $P(p = 1/2) = 0.5$ and $P(p = 3/4) = 0.5$. Now, you make an experiment and find $\eta = 1$. What is the posterior distribution of p? Suppose you make another, independent, experiment, and find, again, $\eta = 1$. What happens to the posterior distribution? Suppose you keep on making experiments and keep on finding $\eta = 1$. What happens to the posterior distributions? Why does this make sense?

2.9. Bibliography

[1] T. Amemiya, *Introduction to Statistics and Econometrics*, Harvard University Press, Cambridge, MA, 1994.
[2] P. Bickel and K. Doksum, *Mathematical Statistics: Basic Ideas and Selected Topics*, Prentice Hall, Upper Saddle River, NJ, 2001.
[3] A. Chorin, Hermite expansions in Monte-Carlo computation, J. Comput. Phys. 8 (1971), pp. 472-482.
[4] K.L. Chung, *A Course in Probability Theory*, Academic Press, New York, 1974.

[5] H. Dym and H. McKean, *Fourier Series and Integrals*, Academic Press, New York, 1972.
[6] A. N. Kolmogorov, *Foundations of the Theory of Probability*, Chelsea, New York, 1932.
[7] J. Lamperti, *Probability*, Benjamin, New York, 1966.

CHAPTER 3

Brownian Motion

3.1. Definition of Brownian Motion

In the chapter that follows we will provide a reasonably systematic introduction to stochastic processes; we start, however, here by considering a particular stochastic process that is of particular importance both in the theory and in the applications.

DEFINITION. A stochastic process (in the strict sense) is a function $v(\omega, t)$ of two arguments where $(\Omega, \mathcal{B}, P)$ is a probability space, $\omega \in \Omega$, and $t \in \mathbb{R}$, such that for each ω, $v(\omega, t)$ is a function of t and for each t, $v(\omega, t)$ is a random variable.

If t is a space variable, the stochastic process is also often called a random field.

DEFINITION. "Brownian motion" (in mathematical terminology) is a stochastic process $w(\omega, t)$, $\omega \in \Omega$, $0 < t < 1$, that satisfies the following four axioms:

1. $w(\omega, 0) = 0$ for all ω.
2. For each ω, $w(\omega, t)$ is a continuous function of t.
3. For each $0 \leq s \leq t, w(\omega, t) - w(\omega, s)$ is a Gaussian variable with mean zero and variance $t - s$.
4. $w(\omega, t)$ has independent increments; i.e., $0 \leq t_1 < t_2 < \cdots < t_n$ then $w(\omega, t_i) - w(\omega, t_{i-1})$ for $i = 1, 2, \ldots, n$ are independent.

Note that what is called in mathematics Brownian motion (BM) is called in physics the Wiener process. Also what is called in physics BM is a different process which is called in mathematics the Ornstein-Uhlenbeck process, which we shall discuss later.

First, one must show that a process that satisfies all of these conditions exists. This is not a trivial issue; we shall see shortly that if the second condition is replaced by the requirement that w be differentiable, then there is no way to satisfy the conditions. The original proof of Wiener consisted of showing that the Fourier series

$$\frac{\pi}{2\sqrt{2}} \sum_{k=1}^{\infty} \frac{a_k}{k} \sin(\pi k t/2),$$

where the a_k are independent Gaussian variables with mean 0 and variance 1, converges, and its sum satisfies the above conditions for $0 \le t \le 1$. Each coefficient is a random function defined on some probability space $(\Omega, \mathcal{B}, P)$ and the resulting BM is also a function on the very same Ω. For longer times, one can construct the process by stringing the processes constructed by this series end to end. We refer the reader to the literature.

Next, we derive some consequences of the definition of BM.

1. The correlation function of BM is $E[w(t_1)w(t_2)] = \min\{t_1, t_2\}$. Indeed, assuming $t_1 < t_2$, we get
$$\begin{aligned} E[w(t_1)w(t_2)] &= E[w(t_1)(w(t_1) + (w(t_2) - w(t_1))] \\ &= E[w(t_1)w(t_1)] + E[w(t_1)(w(t_2) - w(t_1))] = t_1. \end{aligned}$$
In this equation, the variables $w(t_1)$ and $w(t_2) - w(t_1)$ are independent and each has mean 0.
2. Consider the variable
$$\frac{w(\omega, t + \Delta t) - w(\omega, t)}{\Delta t}.$$
It is Gaussian with mean 0 and variance $(\Delta t)^{-1}$, which tends to infinity as Δt tends to zero. Therefore, one can guess that the derivative of $w(\omega, t)$ for any fixed ω exists nowhere with probability 1.

Nondifferentiable functions may have derivatives in the sense of distributions. The derivative in the sense of distributions $v(\omega, s)$ of a BM $w(\omega, t)$ is called "white noise" and is defined by the property:
$$\int_{t_1}^{t_2} v(\omega, s)\, ds = w(\omega, t_2) - w(\omega, t_1).$$
The origin of the name will be clarified in the next chapter.

Two-dimensional BM is $(w_1(\omega, t), w_2(\omega, t))$, where w_1, w_2 are independent BMs, and similarly for n-dimensional BM.

We also consider random walks $W_n(t)$ that approximate BM, constructed as follows: Consider the time interval $[0, 1]$ and divide it into n pieces of equal lengths; define $W_n(0) = 0$ and $W_n(i/n) = W_n((i-1)/n) + W_i$, where the W_i without an argument are independent Gaussian variables with means 0 and variances $1/n$ and i is a positive integer. Then define $W_n(t)$ for intermediate values of t by linear interpolation. Clearly, $W_n(t)$ for large n resembles BM: for all t, $W_n(t)$ is a Gaussian random variable with mean 0; For large n, its variance is at least approximately equal to t. Furthermore, if t_1, t_2, t_3, and t_4 in $[0, t]$ are such that $t_4 > t_3 > t_2 > t_1$ and, furthermore, $t_3 \ge (t_2 + 1/n)$,

then the variables $W_n(t_4) - W_n(t_3)$ and $W_n(t_2) - W_n(t_1)$ are independent. The discussion of the precise relation between $W_n(t)$ and BM is outside the scope of this volume, but we shall take for granted that the convergence of $W_n(t)$ to BM is good enough for the limiting arguments presented below to be valid.

3.2. Brownian Motion and the Heat Equation

We first solve the heat equation

$$v_t = \frac{1}{2} v_{xx}, \quad v(x,0) = \phi(x) \tag{3.1}$$

on $-\infty < x < \infty$, $t > 0$, by Fourier transforms. Let

$$v(x,t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{ikx} \hat{v}(k,t) dk.$$

Then

$$v_x(x,t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} ike^{ikx} \hat{v}(k,t) dk$$

$$v_{xx}(x,t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} (ik)^2 e^{ikx} \hat{v}(k,t) dk$$

$$v_t(x,t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{ikx} \partial_t \hat{v}(k,t) dk.$$

Inserting in (3.1) we obtain

$$\partial_t \hat{v}(k,t) = -\frac{1}{2} k^2 \hat{v}(k,t)$$
$$\hat{v}(k,0) = \hat{\phi}(k).$$

The solution of this ordinary differential equation is

$$\hat{v}(k,t) = e^{-\frac{1}{2}k^2 t} \hat{\phi}(k).$$

Using the expression for $\hat{v}$ in the formula for v and completing the square we get

$$\begin{aligned} v(x,t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{ikx} e^{-\frac{1}{2}k^2 t} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-ikx'} \phi(x') dx' dk \\ &= \int_{-\infty}^{\infty} \frac{e^{-\frac{(x-x')^2}{2t}}}{\sqrt{2\pi t}} \phi(x') \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}\left(k\sqrt{t} - i\left(\frac{x-x'}{\sqrt{t}}\right)\right)^2}}{\sqrt{2\pi}} dk\sqrt{t}\ dx' \\ &= \int_{-\infty}^{\infty} \frac{e^{-\frac{(x-x')^2}{2t}}}{\sqrt{2\pi t}} \phi(x') dx' \\ &= \int_{-\infty}^{\infty} \frac{e^{-\frac{(x')^2}{2t}}}{\sqrt{2\pi t}} \phi(x+x') dx' \end{aligned} \tag{3.2}$$

The function

$$G(x) = \frac{e^{-x^2/2t}}{\sqrt{2\pi t}}$$

is the Green function of the heat equation and we have shown that the solution of the heat equation is the convolution of the initial data with the Green function.

Since the Green function G is also the probability density function for a random variable η with mean zero and variance t we can rewrite (3.2) as

$$v(x,t) = E[\phi(x + \eta(\omega))].$$

Remember that if $w(\omega, t)$ is BM, then for a fixed t, $w(\omega, t)$ is a Gaussian variable with mean 0 and variance t, hence,

$$v(x,t) = E[\phi(x + w(\omega,t))]. \tag{3.3}$$

This result has a geometrical interpretation: consider the point (x,t) at which we want to evaluate w. Start BMs going backwards in time from (x,t); they intersect the x-axis at time t at the points $x + w(\omega,t)$. Find the initial values of v at the points of intersection, and average them over all BMs. This average is $v(x,t)$.

3.3. Solution of the Heat Equation by Random Walks

We now rederive the result above in a more instructive way that will be useful in the analysis of a more general situation. We construct a grid on which to approximate the heat equation (3.1), solve the resulting discrete equations by a random walk, and take a limit that will reproduce the result of the previous section. To construct the grid,

draw horizontal and vertical lines in the (x,t) plane. The distance between the horizontal lines is k (*not* the Fourier variable!), and between the vertical lines is h. The points at which these lines intersect will carry values of an approximation V of the solution $v(x,t)$ of the heat equation. That is, each gridpoint (ih, nk) carries a value of the grid function $V_i^n \sim v(ih,nk) = v_i^n$. Construct a difference approximation of the derivatives in (3.1):

$$v_t \sim \frac{v_i^{n+1} - v_i^n}{k} \sim \frac{V_i^{n+1} - V_i^n}{k} \tag{3.4}$$

$$v_{xx} \sim \frac{v_{i+1}^n + v_{i-1}^n - 2v_i^n}{h^2} \sim \frac{V_{i+1}^n + V_{i-1}^n - 2V_i^n}{h^2}. \tag{3.5}$$

Substituting (3.4) and (3.5) into (3.1), we obtain an equation for the V_i^n:

$$\frac{V_i^{n+1} - V_i^n}{k} = \frac{1}{2}\,\frac{V_{i+1}^n + V_{i-1}^n - 2V_i^n}{h^2}. \tag{3.6}$$

Starting from the initial data $V_i^0 = \phi(ih)$, we can find a solution of (3.6) at time $t = nk$ for any n by the recurrence formula

$$V_i^{n+1} = V_i^n + \lambda(V_{i+1}^n + V_{i-1}^n - 2V_i^n) = (1-2\lambda)V_i^n + \lambda V_{i+1}^n + \lambda V_{i-1}^n, \tag{3.7}$$

where

$$\lambda = \frac{1}{2}\frac{k}{h^2}.$$

Define the local "truncation error"

$$\tau_i^n = \frac{v_i^{n+1} - v_i^n}{k} - \frac{1}{2}\,\frac{v_{i+1}^n + v_{i-1}^n - 2v_i^n}{h^2},$$

where v is a smooth solution of the differential equation (3.1). Using Taylor's formula one finds that

$$\tau_i^n = O(k) + O(h^2).$$

In numerical analysis the fact that τ_i^n tends to zero as $h \to 0, k \to 0$ is called "consistency." Thus the exact solution of the differential equation satisfies the difference equations, up to a small error.

Now we show that for $\lambda \le 1/2$, the approximate solution V converges to the exact solution v as h and k tend to zero. It is easy to check that the error $e_i^n = v_i^n - V_i^n$ satisfies the equation

$$e_i^{n+1} = (1-2\lambda)e_i^n + \lambda e_{i+1}^n + \lambda e_{i-1}^n + k\tau_i^n.$$

Taking the absolute value of both sides we get

$$|e_i^{n+1}| \le (1-2\lambda)|e_i^n| + \lambda|e_{i+1}^n| + \lambda|e_{i-1}^n| + k|\tau_i^n|,$$

where we have assumed that $1 - 2\lambda \geq 0$ (or $\lambda \leq 1/2$). Define

$$E^n = \max_i |e_i^n| \tag{3.8}$$

and let

$$\tau^n = \max_i |\tau_i^n|, \quad \tau = \max_{nk \leq t} |\tau^n|. \tag{3.9}$$

Then,

$$E^{n+1} \leq E^n + k\tau^n \leq E^n + k\tau;$$

thus,

$$E^{n+1} \leq E^n + k\tau \leq E^{n-1} + 2k\tau \leq \cdots \leq E^0 + (n+1)k\tau.$$

If we start from the exact solution, then $E^0 = 0$ and, hence,

$$E^n \leq nk\tau = t\tau.$$

Recall that the local truncation error tends to zero as $h, k \to 0$ and consider the solution of the heat equation on a finite t interval $0 \leq t \leq T$ for some finite T; then E^n tends to zero as h and k tend to zero provided $\lambda = k/(2h^2)$ is less than or equal to $1/2$. That means that the approximate solution converges to the exact solution for $\lambda \leq 1/2$.

Choose $\lambda = 1/2$. Then (3.7) becomes

$$V_i^{n+1} = \frac{1}{2}(V_{i+1}^n + V_{i-1}^n). \tag{3.10}$$

Using (3.10) and iterating backward in time, we can write V_i^n in terms $V_i^0 = \phi(ih)$:

$$\begin{aligned} V_i^n &= \frac{1}{2}V_{i+1}^{n-1} + \frac{1}{2}V_{i-1}^{n-1} \\ &= \frac{1}{4}V_{i-2}^{n-2} + \frac{2}{4}V_i^{n-2} + \frac{1}{4}V_{i+2}^{n-2} \\ &\vdots \\ &= \sum_{j=0}^{n} C_{j,n}\phi((-n + 2j + i)h). \end{aligned}$$

It is easy to see that the numbers $C_{j,n}$ are the binomial coefficients divided by 2^n:

$$C_{j,n} = \frac{1}{2^n}\binom{n}{j}. \tag{3.11}$$

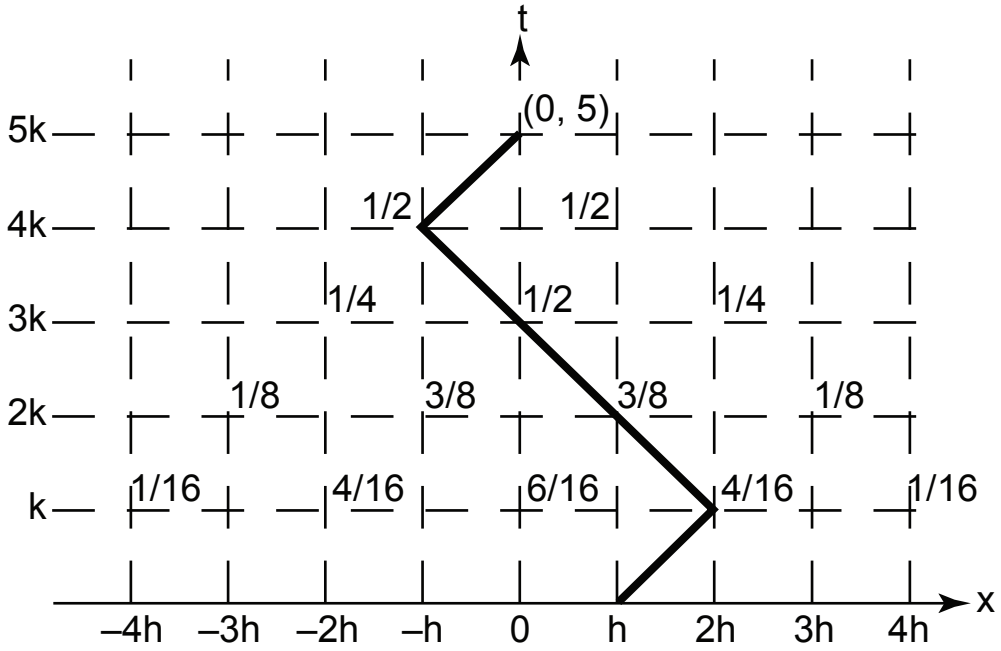

FIGURE 3.1. Backward walk for the heat equation.

Thus,

$$V_i^n = \sum_{j=0}^{n} \frac{1}{2^n} \binom{n}{j} \phi((-n+2j+i)h). \tag{3.12}$$

We can interpret the numbers $C_{j,n}$ as follows: Imagine that a drunkard makes a step h to the left with probability $1/2$ or a step h to the right with probability $1/2$ starting from the point (ih, nk) (see Figure 3.1). Assume that her successive steps are independent. The probability that she will reach the point $((-n+2j+i)h, 0)$ after n steps is exactly $C_{j,n}$. We can represent this drunken walk as a sum of n random variables

$$\eta_k = \begin{cases} h & \text{probability } \frac{1}{2} \\ -h & \text{probability } \frac{1}{2}, \end{cases}$$

with $k = 1, 2, \ldots, n$. This representation gives us another expression for $C_{j,n}$:

$$C_{j,n} = P\left(\sum_{k=1}^{n} \eta_k = (-n+2j)\, h \right). \tag{3.13}$$

According to the central limit theorem, the sum $\sum_{k=1}^{n} \eta_k$ tends to a Gaussian variable with mean 0 and variance nh^2 as $n \to \infty$. Recall that

$\lambda = k/(2h^2) = 1/2$; consequently, $h^2 = k$ and, hence, $nh^2 = nk = t$. So, $\sum_{k=1}^{n} \eta_k$ tends to a Gaussian variable with mean 0 and variance t as $n \to \infty$, $h \to 0$ and $nh^2 = t$. Hence,

$$P\left(\sum_{k=1}^{n} \eta_k = (-n + 2j)h\right) \sim \frac{e^{-(x')^2/2t}}{\sqrt{2\pi t}} \cdot 2h,$$

where $x' = (-n + 2j)h$. Therefore,

$$V_i^n = \sum_{j=0}^{n} C_{j,n}\phi((-n + 2j + i)h) \to \int_{-\infty}^{\infty} \frac{e^{-(x-x')^2/2t}}{\sqrt{2\pi t}}\phi(x')\, dx' \quad (3.14)$$

as $n \to \infty$. We have used the central limit theorem to derive the solution formula for the heat equation.

3.4. The Wiener Measure

Having described the properties of Brownian motion, the goal in the present section is to construct a probability measure on the space of continuous functions in such a way that the set of Brownian motions (the samples of BM, the functions $w(\omega, t)$ for various values of ω) have probability 1. In other words, we are trying to define a measure on a space of functions in such a way that the only functions that count are the BMs. To this end, consider the space of continuous functions $u(t)$ such that $u(0) = 0$. This collection is now our sample space Ω. Ω is the space of experimental outcomes and our experiment consists of creating an instance of a continuous function with $u(0) = 0$.

Next we need to define a σ-algebra. Pick an instant in time, say t_1, and associate with this instant a window of values $(a_1, b_1]$, where $-\infty \le a_1, b_1 \le \infty$. Consider the subset of all the continuous functions that pass through this window and denote it as C_1. This subset is called a cylinder set. For every instant and every window, we can define a corresponding cylinder set; i.e., C_i is the subset of all continuous functions that pass through the window $(a_i, b_i]$ at the instant t_i. Next, consider two cylinder sets C_1 and C_2. Then $C_1 \cap C_2$ is the set of functions that pass through both windows. Similarly, $C_1 \cup C_2$ is the set of functions that pass through either C_1 or C_2. It can be shown that the class of finite disjoint unions of intersections of cylinders is closed under finite unions, intersections and complements, i.e., they form an algebra on the space of continuous functions u in $[0, 1]$ with $u(0) = 0$.

The next step in our construction is to define a measure (i.e., a rule by which to attach probabilities to the cylinder sets). We want to define the measure in such a way that is appropriate for BMs. Take

the cylinder set C_1. If the functions that belong to this cylinder set are Brownian motions, the probability of the cylinder set is

$$P(C_1) = \int_{a_1}^{b_1} \frac{e^{-s_1^2/2t_1}}{\sqrt{2\pi t_1}}\, ds_1.$$

Assign this P to this set, and similarly for other cylinder sets constructed in the same way at different values of t.

Next, consider the intersection $C_1 \cap C_2$ of two cylinder sets C_1 and C_2 with $t_2 > t_1$. By the property of Brownian motion that nonoverlapping increments are independent random variables with Gaussian distributions, we conclude that the probability we should assign to $C_1 \cap C_2$ is

$$P(C_1 \cap C_2) = \int_{a_1}^{b_1} \frac{e^{-s_1^2/2t_1}}{\sqrt{2\pi t_1}}\, ds_1 \int_{a_2}^{b_2} \frac{e^{-(s_2-s_1)^2/2(t_2-t_1)}}{\sqrt{2\pi(t_2-t_1)}}\, ds_2.$$

Similarly, we can define a probability for the intersection any finite number of cylinder sets. The cylinder sets can be embedded in a variety of σ-algebras. These are not equivalent, but we choose a σ-algebra that contains the set of all continuous functions with $u(0) = 0$.

It can be shown that the measure defined in this way can be extended to a probability measure on the σ-algebra. We shall not give the details but refer the reader to the literature. The identity $P(\Omega) = 1$ can be seen from the evaluation of the Gaussian integrals in the interval $(-\infty, +\infty)$. The measure we defined was introduced by Wiener and carries his name.

Suppose that F is a number attached to a continuous function. For example, if $u(s)$ is a continuous function with $u(0) = 0$ and $0 \le s \le 1$, then we could define $F = \int_0^1 u^2(s)\, ds$. Any mapping that attaches a number to a function is, for historical reasons, called a functional. Also for historical reasons, a functional acting on a function $u(\cdot)$ is written as $F[u(\cdot)]$. F is a function on Ω, the space of continuous functions that start from the origin.

If one has a measure, one has an integral. Denote the integral with respect to the Wiener measure by $\int dW$. In particular, if χ_C is the indicator function of the set C ($\chi_C = 1$ if ω is in C, $\chi_C = 0$ otherwise), then $\int \chi_C\, dW = P(C)$. If we attach to each BM w a number $F[w(\cdot)]$ (the number is attached to the BM viewed as a whole, not to particular point values), then the integral $\int F[w(\cdot)]\, dW$ is, by definition, the expected value of F as w runs over all the possible BMs.

EXAMPLE. Suppose $F[w(\cdot)] = w^2(1)$; that is, we take a BM w, look at the value of w when $t = 1$, and square that number. This

is a number attached to w. $w(1)$ is by definition a Gaussian random variable with mean 0 and variance 1. Then

$$\int F\,dW = \int_{-\infty}^{+\infty} u^2 \frac{e^{-u^2/2}}{\sqrt{2\pi}}\,du = 1.$$

EXAMPLE. Fubini's theorem can be extended to integrals more abstract than the elementary finite-dimensional integral and, in particular, we can show that it is legitimate, under appropriate conditions, to interchange the order of integration with respect to the Wiener measure and ordinary integration. For instance, if $F[w(\cdot)] = \int_0^1 w^2(s)\,ds$ (a perfectly reasonable way to attach a number to the function $w(t)$), then

$$\int dW \int_0^1 w^2(s)\,ds = \int_0^1 ds \int dW\, w^2(s) = \int_0^1 s\,ds = \frac{1}{2}$$

because $w(s)$ is a Gaussian variable with variance s and mean 0.

3.5. Heat Equation with Potential

Now consider the initial value problem

$$v_t = \frac{1}{2}v_{xx} + U(x)v, \quad v(x,0) = \phi(x). \tag{3.15}$$

(Note that with the addition of the imaginary i in front of the time derivative, this would be a Schroedinger equation and U would be a potential.) Generalizing what has been done before, approximate this equation by

$$\frac{V_i^{n+1} - V_i^n}{k} = \frac{1}{2}\frac{V_{i-1}^n + V_{i+1}^n - 2V_i^n}{h^2} + \frac{1}{2}\left(U_{i-1}V_{i-1}^n + U_{i+1}V_{i+1}^n\right), \tag{3.16}$$

where $U_i = U(ih)$ and V_i^n is, as before, a function defined on the nodes (ih, nk) of a grid. Note the clever split of the term Uv into two halves; we now show that the addition of these terms does not destroy the convergence of the approximation to the solution of the differential equation. To check consistency we let v_i^n be the exact solution evaluated at the gridpoints (ih, nk). As before,

$$\frac{v_i^{n+1} - v_i^n}{k} = v_t + O(k), \quad \frac{v_{i+1}^n + v_{i-1}^n - 2v_i^n}{h^2} = v_{xx} + O(h^2).$$

For the potential term we find

$$\begin{aligned}\frac{1}{2}\left(U_{i+1}v_{i+1}^n + U_{i-1}v_{i-1}^n\right) &= \frac{1}{2}\left(2U_iv_i^n + h^2(Uv)_{xx} + h^2O(h^2)\right)\\ &= U_iv_i^n + O(h^2).\end{aligned}$$

We can therefore define the truncation error by

$$\begin{aligned}\tau_i^n &= \frac{v_i^{n+1}-v_i^n}{k} - \frac{1}{2}\frac{v_{i+1}^n + v_{i-1}^n - 2v_i^n}{h^2} - \frac{1}{2}\left(U_{i+1}v_{i+1}^n + U_{i-1}v_{i-1}^n\right)\\ &= v_t - \frac{1}{2}v_{xx} - U(x)v + O(k) + O(h^2)\\ &= O(k) + O(h^2).\end{aligned}$$

Thus the truncation error is small.

Now we show that the approximate solution converges to the exact solution as k and h tend to zero. Let $\lambda = k/(2h^2)$, as before. The exact solution of (3.15) satisfies

$$v_i^{n+1} = (1-2\lambda)v_i^n + \lambda v_{i+1}^n + \lambda v_{i-1}^n + \frac{k}{2}\left(U_{i+1}v_{i+1}^n + U_{i-1}v_{i-1}^n\right) + k\tau_i^n,$$

while the approximate solution satisfies

$$V_i^{n+1} = (1-2\lambda)V_i^n + \lambda V_{i+1}^n + \lambda V_{i-1}^n + \frac{k}{2}\left(U_{i+1}V_{i+1}^n + U_{i-1}V_{i-1}^n\right).$$

Thus the error $e_i^n = v_i^n - V_i^n$ satisfies

$$e_i^{n+1} = (1-2\lambda)e_i^n + \lambda e_{i+1}^n + \lambda e_{i-1}^n + \frac{k}{2}(U_{i+1}e_{i+1}^n + U_{i-1}e_{i-1}^n) + k\tau_i^n.$$

Taking the absolute value of both sides and choosing $\lambda \le 1/2$, we get

$$\begin{aligned}|e_i^{n+1}| \le (1-2\lambda)|e_i^n| + \lambda|e_{i+1}^n| + \lambda|e_{i-1}^n| +\\ + \frac{k}{2}(|U_{i+1}||e_{i+1}^n| + |U_{i-1}||e_{i-1}^n|) + k|\tau_i^n|.\end{aligned}$$

Assume that the potential is bounded,

$$|U(x)| \le M,$$

and recall the definitions of E^n (3.8) and τ^n (3.9). It follows that

$$E^{n+1} \le E^n + MkE^n + k\tau^n \le E^n(1+Mk) + k\tau$$

and hence

$$E^{n+1} \le e^{kM}E^n + k\tau.$$

Then,

$$\begin{aligned} E^{n+1} &\le e^{kM}E^n + k\tau \\ &\le e^{kM}(e^{kM}E^{n-1} + k\tau) + k\tau \\ &= e^{2kM}E^{n-1} + k\tau(1 + e^{kM}) \\ &\le \cdots \\ &\le e^{(n+1)kM}E^0 + k\tau\left(1 + e^{kM} + e^{2kM} + \cdots + e^{nkM}\right) \\ &= e^{(n+1)kM}E^0 + k\tau\frac{e^{(n+1)kM} - 1}{e^{kM} - 1}. \end{aligned}$$

Since we start to compute the approximate solution from the given initial condition $v(x, 0) = \phi(x)$, we may assume that $E^0 = 0$. Therefore, at time $t = nk$, E^n is bounded by

$$E^n \le k\tau\frac{e^{tM} - 1}{e^{kM} - 1} \le \frac{\tau}{M}(e^{tM} - 1).$$

We see that E^n tends to zero as k and h tend to zero with $\lambda \le 1/2$. Thus, the approximation is convergent.

Now set $\lambda = 1/2$. Then for the approximate solution, we have

$$\begin{aligned} V_i^{n+1} &= \frac{1}{2}(V_{i-1}^n + V_{i+1}^n) + \frac{k}{2}(U_{i+1}V_{i+1}^n + U_{i-1}V_{i-1}^n) \\ &= \frac{1}{2}(1 + kU_{i+1})V_{i+1}^n + \frac{1}{2}(1 + kU_{i-1})V_{i-1}^n. \end{aligned}$$

By induction the approximate solution V can be expressed as

$$V_i^n = \sum_{\ell_1=\pm1,\dots,\ell_n=\pm1} \frac{1}{2^n}\left(1 + kU_{i+\ell_1}\right)\cdots\left(1 + kU_{i+\ell_1+\cdots+\ell_n}\right)V_{i+\ell_1+\cdots+\ell_n}^0.$$

Here, unlike in the case $U = 0$, each movement to the right or to the left brings in not just a factor $\frac{1}{2}$ but a factor of $\frac{1}{2}(1 + kU(x))$. Each choice of $\ell_1, \dots, \ell_n$ corresponds to a path. We simply connect the points (ih, nk), $(ih + \ell_1 h, (n-1)k), \dots, (ih + \ell_1 h + \cdots + \ell_n h, 0)$, see Figure 3.2

We will interpret the approximate solution probabilistically. Let $\{\eta_m\}$ be a collection of independent random variables with $P(\eta_m = h) = P(\eta_m = -h) = \frac{1}{2}$. Since $P(\eta_1 = \ell_1 h, \dots, \eta_n = \ell_n h) = \frac{1}{2^n}$ we see that

$$V_i^n = E_{\text{all paths}}\left\{\prod_{m=1}^{n}(1 + kU(ih + \eta_1 + \cdots + \eta_m))\phi(ih + \eta_1 + \cdots + \eta_n)\right\}.$$

To describe the random paths we use linear interpolation. Let $t_n = nk$ and $s_m = mk$. If $s_{m-1} \le s \le s_m$ set

$$\tilde{w}(s) = \eta_1 + \cdots + \eta_{m-1} + \frac{s - s_{m-1}}{k}\eta_m.$$

Each path starting at $(x,t) = (ih, nk)$ is then of the form $(x+\tilde{w}(s), t-s)$ for $0 \le s \le t$ and

$$V_i^n = E_{\substack{\text{all broken}\\ \text{line paths}}}\left\{\prod_{m=1}^{n}\left(1 + kU(x+\tilde{w}(s_m))\right)\phi(x+\tilde{w}(t))\right\}.$$

If $k|U| < 1/2$ then $(1+kU) = \exp(kU+\epsilon)$ where $|\epsilon| \le Ck^2$, so we can write the product as

$$\prod_{m=1}^{n}\left(1 + kU(x+\tilde{w}(s_m))\right) = \exp\left(k\sum_{m=1}^{n}U(x+\tilde{w}(s_m)) + \epsilon'\right)$$

where $|\epsilon'| \le nCk^2 = Ctk$. Since $k\sum_{m=1}^{n}U(x+\tilde{w}(s_m))$ is a Riemann sum for the integral $\int_0^t U(x+\tilde{w}(s))\,ds$ it follows that

$$V_i^n = E_{\substack{\text{all broken}\\ \text{line paths}}}\left\{e^{\int_0^t U(x+\tilde{w}(s))ds}\phi(x+\tilde{w}(t))\right\} + \text{ small terms.}$$

As h, k tend to zero the broken line paths $x+\tilde{w}(s)$ look more and more like Brownian motion paths $x + w(s)$ so in the limit

$$\begin{aligned} v(x,t) &= E_{\substack{\text{all Brownian}\\ \text{motion paths}}}\left\{e^{\int_0^t U(x+w(s))ds}\phi(x+w(t))\right\} \\ &= \int dW e^{\int_0^t U(x+w(s))ds}\phi(x+w(t)). \end{aligned} \tag{3.17}$$

This is the Feynman-Kac formula. It reduces to the solution formula for the heat equation when $U = 0$. This result is useful in quantum mechanics and in other fields.

We now use the path integral representation to derive a time-dependent perturbation scheme, together with extremely useful graphical representations of the various terms, known as "Feynman diagrams". To introduce the ideas, consider first the ordinary differential equation

$$\frac{du}{dt} = au + f(t), \tag{3.18}$$

where $u = u(t)$ is the unknown function, $u(0) = u_0$ is an initial value, a is a constant, and $f = f(t)$ is a given function. As is easily shown, the solution is

$$u(t) = u_0 e^{at} + \int_0^t f(s)e^{a(t-s)}ds. \tag{3.19}$$

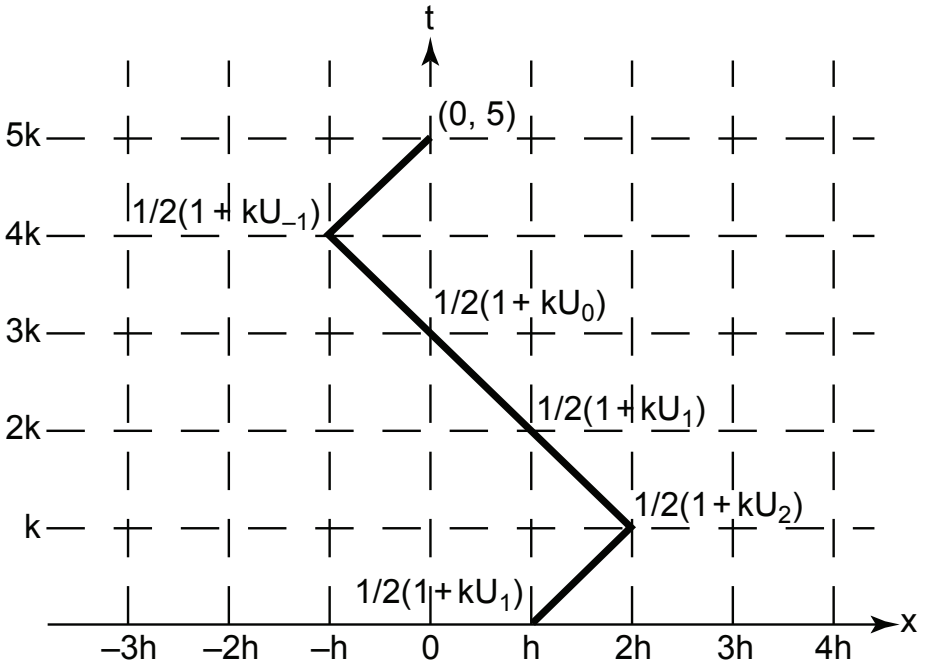

FIGURE 3.2. Backward walk for the heat equation with potential.

The function e^{at}, which solves the equation when $u_0 = 1$ and $f(t) = 0$, is a "propagator", or "vacuum propagator", in the language that

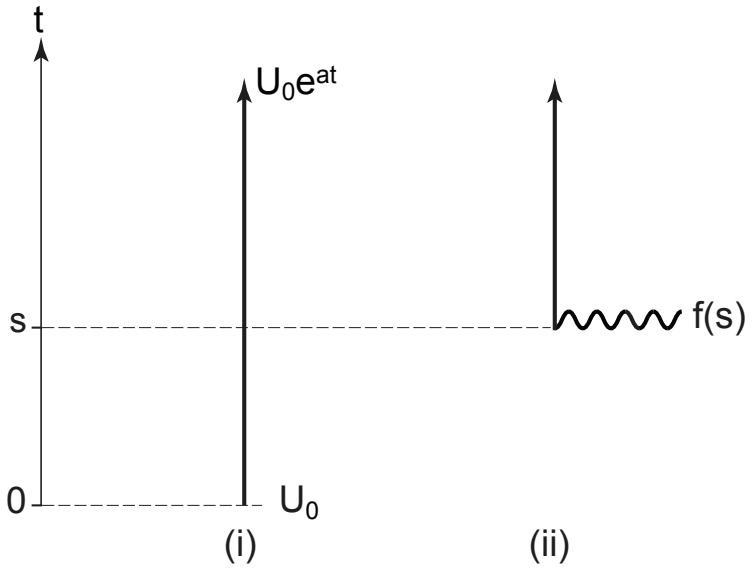

FIGURE 3.3. Diagrams for $\frac{du}{dt} = au + f(t)$.

will soon be introduced. The first term in the solution, $u_0 e^{at}$, can be picturesquely described as a propagation in time, via the vacuum propagator, of the initial value u_0, and can be graphically represented as in part (i) of Figure 3.3.

The second term in the expression for $u(t)$ can be understood by considering what the solution would be like if the function $f(t)$ were to vanish except in a small interval $[s, s+ds]$ with $0 < s < t$. Before s the solution would have only the first term. The presence of f would change the solution at s by an amount $f(s)ds$, that variation would be added to the first term, and then would also propagate by the vacuum propagator. The effect of the "outside force" f would be a sum of such perturbations, all propagating with the same vacuum propagator but starting at different times. This can be described by the picture (ii) in Figure 3.3; the wavy line represents an interaction with the "external force " f, which then propagates as just described, it being understood that the interaction and its aftereffect must be summed over all times s between 0 and t. The fact that the graph starts at s and not at the initial line represents the fact that the second term is independent of the initial value u_0.

We now generalize this representation to the heat equation perturbed by a potential. One should expect some kind of averaging to be necessary not only in time but also in space. First introduce an ϵ in front of the potential U in the equation, so that it reads $v_t = \frac{1}{2}v_{xx} + \epsilon U(x)u$, with the Feynman-Kac formula acquiring an ϵ in the obvious place; the presence of this ϵ suggests that our calculations are more likely to be useful when ϵ is small, but more important, it will allow us to label the various terms by the power of ϵ that precedes them. Next, expand $\exp\left(\int_0^t \epsilon U(x+w(s))ds\right)$ in Taylor series:

$$\exp\left(\int_0^t \epsilon U(x+w(s))ds\right) = 1 + \epsilon\int_0^t U\left(x+w(s)\right)ds) + \frac{1}{2}\epsilon^2(\int_0^t U(x+w(s))ds))^2 + ... \quad (3.20)$$

and substitute the series into the Wiener integral representation of $u(x,y)$. Write

$$K(z,s) = \frac{1}{\sqrt{2\pi s}}\, e^{-z^2/2s}, \quad (3.21)$$

so that the first term in the series, which would be the whole solution in the absence of U, becomes:

$$T_0 = \int_{-\infty}^{\infty} \frac{e^{-y^2/2t}}{\sqrt{2\pi t}} \phi(x+y)dy \tag{3.22}$$

$$= \int_{-\infty}^{+\infty} K(x-z,t)\phi(z)dz. \tag{3.23}$$

K is the "vacuum propagator"; indeed, one can think of the BM's that define the solution as propagating in space, with a motion modified by the potential U along their paths; if $U = 0$ as in this first term, one can think of them propagating in a "vacuum". T_0 can be represented graphically as in the "Feynman diagram" (i) of Figure 3.4: the straight line represents vacuum propagation, which starts from (x,t) and goes to $(z,0)$ in the plane, it being understood that an integration over z is to be performed.

The second term T_1 in the expansion has a coefficient ϵ multiplying the integral

$$\int dW \int_0^t U(x+w(s))\phi(x+w(t))ds = \int_0^t ds \int dW\, U(x+w(s))\phi(x+w(t)). \tag{3.24}$$

The variables $x+w(s), x+w(t)$ are both Gaussian, but they are not independent, so that in order to average one has to find their joint pdf. It is easier to express the integrand as a function of two independent variables; clearly $s \leq t$ so that $w(t) = w(s) + (w(t) - w(s))$, and $w(s), w(t)-w(s)$ are independent, by the definition of BM. Now $x+w(s)$ is a Gaussian variable with mean x and variance s, $w(t) - w(s)$ is a Gaussian variable with mean 0 and variance $t - s$, so T_1 becomes

$$T_1 = \epsilon \int_0^t ds \int_{-\infty}^{+\infty} dz_1 \int_{-\infty}^{+\infty} dz_2 K(z_1 - x, s) \cdot \cdot U(z_1)K(z_2, t-s)\phi(z_1+z_2). \tag{3.25}$$

This term can be represented graphically as in Figure 3.4 (ii): vacuum propagation from (x,t) to $(z_1, t-s)$, interaction with the potential U at z_1 (represented by a wavy line), followed by a vacuum propagation from $(z_1, t-s)$ to $(z_1+z_2, 0)$, it being understood that one integrates over all intermediate quantities, s, z_1, z_2.

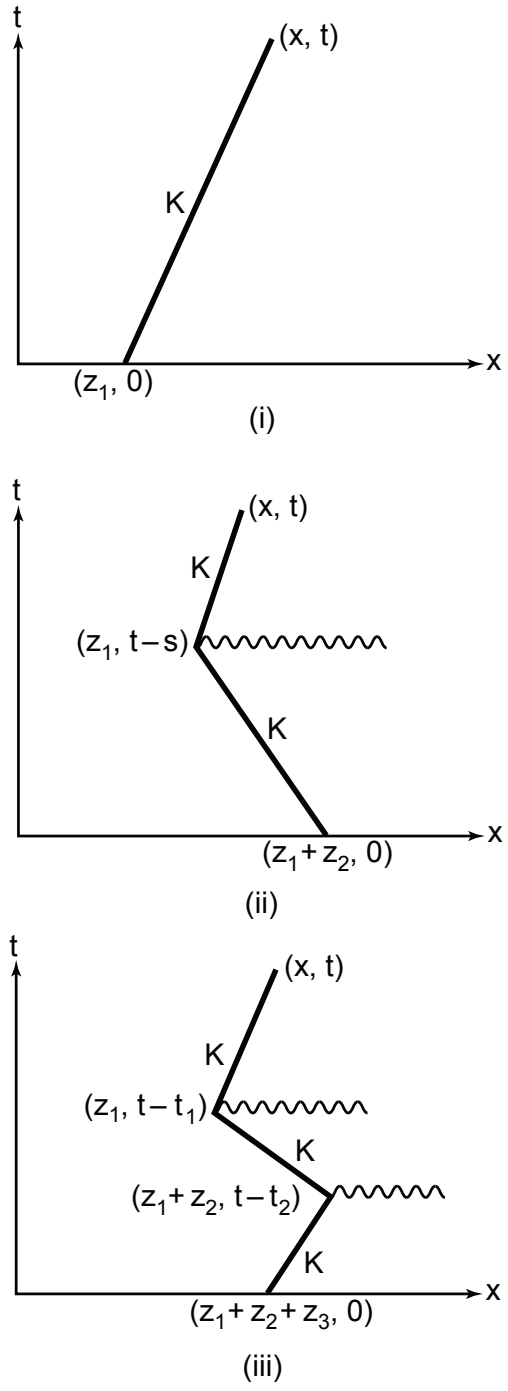

FIGURE 3.4. Feynman diagrams.

To evaluate the second term, we need the identity

$$\left(\int_0^t f(s)ds\right)^2 = 2\int_0^t d\tau_2 \int_0^{\tau_2} d\tau_1 f(\tau_1)f(\tau_2), \tag{3.26}$$

which is easily proved by differentiating both sides; note that in this formula $t \geq \tau_2 \geq \tau_1$. The second term T_2 then becomes ϵ^2 multiplying

$$\int dW \int_0^t d\tau_2 \int_0^{\tau_2} d\tau_1 U(x + w(\tau_1))U(x + w(\tau_2))\phi(x + w(t)). \tag{3.27}$$

As before, write $x + w(\tau_2) = x + w(\tau_1) + w(\tau_2) - w(\tau_1)$, and $x + w(t) = x + w(\tau_1) + w(\tau_2) - w(\tau_1) + w(t) - w(\tau_2)$ to create independent variables, and note that $x+w(\tau_1)$ is Gaussian with mean x and variance τ_1, $w(\tau_2) - w(\tau_1)$ is Gaussian with mean 0 and variance $\tau_2 - \tau_1$, and $w(t) - w(\tau_2)$ is Gaussian with mean 0 and variance $t - \tau_2$. T_2 then becomes:

$$\begin{aligned} T_2 =& \epsilon^2 \int_0^t d\tau_2 \int_0^{\tau_2} d\tau_1 \int_{-\infty}^{\infty} dz_1 \int_{-\infty}^{\infty} dz_2 \int_{-\infty}^{\infty} dz_3 \cdot \\ & \cdot K(z_1 - x, \tau_1)U(z_1)K(z_2, \tau_2 - \tau_1)U(z_1 + z_2)\cdot \\ & \cdot K(z_3, t - \tau_2)\phi(z_1 + z_2 + z_3). \end{aligned} \tag{3.28}$$

This can be represented by the diagram (iii) of Figure 3.4. The higher order terms follow the same pattern. The point of the diagrams is that they are much easier to generate and visualize than the corresponding integral expressions.

3.6. Physicists' Notation for Wiener Measure

Physicists use an interesting notation for the Wiener measure that can be useful when one uses Wiener integrals in problems of mechanics and quantum mechanics. There are no new ideas here, just new notation. Before proceeding, we recall a number of results already established. In the construction of cylinder sets for the Wiener measure, pick an event $C = \bigcap C_i$, where C_i is associated with the interval $(a_i, b_i]$ and $t_i = ih$. Additionally, assume that the windows are of small width (i.e., $b_i - a_i = du_i$). The probability attached to such a set is

$$P = \int_{a_1}^{b_1} \frac{e^{-u_1^2/2h}}{\sqrt{2\pi h}}\, du_1 \int_{a_2}^{b_2} \frac{e^{-(u_2-u_1)^2/2h}}{\sqrt{2\pi h}}\, du_2 \cdots \int_{a_n}^{b_n} \frac{e^{-(u_n-u_{n-1})^2/2h}}{\sqrt{2\pi h}}\, du_n. \tag{3.29}$$

For sufficiently narrow windows, each integral in (3.29) can be approximated by

$$\int_{a_i}^{b_i} \frac{e^{-(u_i-u_{i-1})^2/2h}}{\sqrt{2\pi h}}\,du_i \approx \frac{e^{-(u_i^*-u_{i-1}^*)^2/2h}}{\sqrt{2\pi h}}\,du_i,$$

where $u_i^* \in [a_i, b_i]$. Therefore, P can be approximated by

$$P \approx \frac{1}{Z}\exp\left(-\sum_{i=1}^{n}\frac{(u_i^*-u_{i-1}^*)^2 h}{2h^2}\right)[du],$$

where $[du] = du_1 du_2 \cdots du_n$ and Z is an appropriate normalization constant. Thus, formally (this means "not rigorously" or "not really"),

$$dW = \frac{1}{Z}e^{-\frac{1}{2}\int_0^t \left(\frac{du}{ds}\right)^2 ds}[du], \quad Z = (2\pi h)^{n/2}.$$

This expression is formal in the sense that neither the integral in the exponent, nor the limiting Z, nor the product of du's hidden in $[du]$ exists. As mentioned previously, it can be shown that Brownian motion is not differentiable. However, this expression turns out to be useful.

Recall that, given the equation $v_t = \frac{1}{2}v_{xx} + U(x)v$ with the initial data $v(x,0) = \phi(x)$, we have

$$v(x,t) = \int e^{\int_0^t U(x+w(s))\,ds}\phi(x+w(t))\,dW.$$

In the new notation, this last integral can be written as

$$v(x,t) = \frac{1}{Z}\int e^{-\int_0^t \left[\frac{1}{2}\left(\frac{dw}{ds}\right)^2 - U(x+w(s))\right]ds}\phi(x+w(t))[dw]. \tag{3.30}$$

By definition, (3.30) is a "sum over paths." In principle, one can evaluate it by taking many Brownian motion paths, evaluating the integrals for each path, and then averaging the results. The formal (i.e., meaningless if one looks too closely) expression $[dw]$ is often written as "$d$path" (or something similar). Note, and this is an important point, that the exponent is an integral of what we will see is a Lagrangian. Similar integrals appear in quantum mechanics (with an additional imaginary factor i in the exponent).

If one is given an expression for a measure in the form (3.30), one can interpret it properly by retracing the steps that led to that form. The integral of the squared derivative denotes the Wiener measure, the other part of the integral can be discretized, and the terms in the resulting sums become the probabilities of a "path" belonging to a cylinder set.

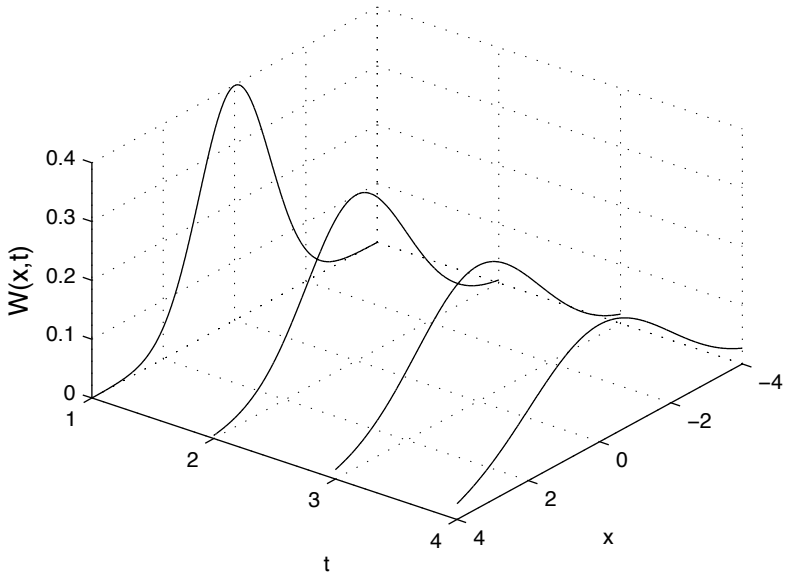

FIGURE 3.5. $W(x,t)$ for Brownian motion.

3.7. Another Connection Between Brownian Motion and the Heat Equation

We now examine a relationship between BM going forward in time and related stochastic processes on one hand, and parabolic equations, including the heat equation, on the other hand. This relationship will be generalized in the following section and will make it possible to use stochastic processes related to BM to solve yet broader classes of equations. Consider the random variables $w(\omega, t)$ as functions of ω (i.e., as random variables) for several values of t. Define the function $W = W(x,t)$ by

$$W(x,t)\,dx = P(x < w(t) \le x + dx), \tag{3.31}$$

where $w(t)$ is a Brownian motion. $W(x,t)$ is the probability density function of the Brownian motion $w(t)$ at the fixed moment t. As we know,

$$W(x,t) = \frac{e^{-x^2/2t}}{\sqrt{2\pi t}}.$$

The graphs of $W(x,t)$ for several values of t are shown in Figure 3.5.

We see that the graphs become lower and wider as t increases. A direct calculation shows that $\partial_t W = \frac{1}{2}\partial_x^2 W$. The increments of Brownian motion are independent. This means that if we know that at time t $w(t)$ is at x, then where it is at $t + \Delta t$ does not depend on

where it was prior to the moment t. The relation between $W(x,t)$ and $W(x,t+\Delta t)$ is given by the Chapman-Kolmogorov equation

$$W(x,t+\Delta t)=\int_{-\infty}^{\infty}W(x+y,t)\Psi(x,y,\Delta t)\,dy, \tag{3.32}$$

where

$$\Psi(x,y,\Delta t)=\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}} \tag{3.33}$$

is the "transition kernel." This equation states that the probability of reaching x at time $t+\Delta t$ is the sum of the probabilities of reaching $x+y$ at time t multiplied by the probability of going from $x+y$ to x during the time interval Δt.

Expand $W(x+y,t)$ in a Taylor series in y:

$$W(x+y)=W(x)+yW_x(x)+\frac{y^2}{2}W_{xx}(x)+\frac{y^3}{6}W_{xxx}(x)+O(y^4)$$

and substitute it into (3.32):

$$\begin{aligned}\int_{-\infty}^{\infty}W(x+y,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy=&\int_{-\infty}^{\infty}W(x,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy\\&+\int_{-\infty}^{\infty}yW_x(x,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy\\&+\frac{1}{2}\int_{-\infty}^{\infty}y^2W_{xx}(x,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy\\&+\frac{1}{6}\int_{-\infty}^{\infty}y^3W_{xxx}(x,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy\\&+\int_{-\infty}^{\infty}O(y^4)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy;\end{aligned}$$

thus,

$$\int_{-\infty}^{\infty}W(x+y,t)\frac{e^{-y^2/2\Delta t}}{\sqrt{2\pi\Delta t}}\,dy=W(x)+0+\frac{W_{xx}(x)\Delta t}{2}+0+O(\Delta t^2).$$

Hence, we have

$$W(x,t+\Delta t)=W(x,t)+\frac{\Delta t}{2}W_{xx}(x,t)+O(\Delta t^2).$$

Dividing by Δt, we obtain

$$\frac{W(x,t+\Delta t)-W(x,t)}{\Delta t}=\frac{1}{2}W_{xx}(x,t)+O(\Delta t).$$

Letting $\Delta t \to 0$, we find

$$\frac{\partial W}{\partial t} = \frac{1}{2}\frac{\partial^2 W}{\partial x^2}.$$

This is a "Fokker-Planck equation"—an equation that describes the time evolution of a one-time probability density for a stochastic process. We see that the Fokker-Planck equation for BM is the heat equation. This observation provides another relation between BM and the heat equation.

A stochastic process is called a Markov process if what happens after time t is independent of what happened before time t; that is, if $t' > t$, then

$$E[u(\omega, t')|u(\omega, t)] = E[u(\omega, t')|u(\omega, s), s \le t].$$

In other words, if we know $u(\omega, t)$, then knowing in addition $u(\omega, s)$ for $s < t$ does not help us to predict $u(\omega, t')$ for $t' > t$.

As discussed above, if $P(x < u(t) \le x + dx) = W(x, t)\,dx$, then $W(x, t)$ satisfies the Chapman-Kolmogorov equation

$$W(x, t + \Delta t) = \int W(x + y, t)\Psi(x, y, \Delta t)\,dy,$$

where Ψ is the "transition probability" from a state $x + y$ at time t to the state x at time $t + \Delta t$. For a Markov process, the transition probability does not depend on $W(x, s)$ for $s < t$. Brownian motion is by construction a Markov process because it has independent increments.

3.8. First Discussion of the Langevin Equation

Let $u(t, \omega)$ be a stochastic process defined by the following (formal) equation:

$$\frac{du}{dt} = -au + \frac{dw}{dt}$$

where a is a positive constant and dw/dt is white noise, the derivative of a Brownian motion w. We know that this derivative does not exist in the classical sense; thus, the equation makes sense only formally (or else in the sense of distributions). A more sensible way to write the Langevin equation is

$$du = -au\,dt + dw, \tag{3.34}$$

where dw is the increment of Brownian motion. The meaning of (3.34) is defined by integrating from 0 to t:

$$u(t) - u(0) = -a \int_0^t u\, dt + \int_0^t dw = -a \int_0^t u\, dt + w(t).$$

This is the Langevin equation (also known in some mathematical circles as the Ornstein-Uhlenbeck equation). It is an instance of a stochastic differential equation. The equation contains a term that is a random function of t, and the solution $u = u(\omega, t)$ should also be a random function of t that satisfies the equation for every ω in the probability space on which the equation is defined. The solution of this equation is known to mathematicians as the Ornstein-Uhlenbeck process.

If we omit the noise term in this equation and retain only the "damping" term $-au$, the solution is the initial datum times e^{-at}, a pure decay. If, on the other hand, we keep the noise term but set $a = 0$, the solution of the equation is Brownian motion. In physics, this equation is used to model the motion of a heavy particle under bombardment by lighter particles; the collisions with the lighter particles provide random instantaneous bursts of added momentum while the mean effect of the collisions is to slow the heavy particle down. We will see in Section 6.1 that when this equation is used as a physical model, the coefficient a, as well as the coefficient of the noise term that we have, rather arbitrarily, set equal to 1, acquire a direct physical meaning. The solution of this equation, with the coefficients interpreted correctly, is what physicists call Brownian motion.

Similarly to what we did in the case of Brownian motion, we want to find the equation satisfied by the probability density function of u (i.e., the Fokker-Planck equation for this problem, also know to mathematicians as the Kolmogorov equation). We choose an approximation for (3.34): Integrating from nk to $(n+1)k$, where k is the time step, we find

$$u^{n+1} - u^n = -aku^n + w^{n+1} - w^n. \tag{3.35}$$

We choose k small enough so that $ak < 1$. The choice to evaluate the term $-aku$ at time nk is not just an arbitrary choice of approximation scheme but is dictated by the desire to a obtain a solution that constitutes a Markov process; as we are constructing the solution step-by-step in time, what we have to work with when we go from time $t = nk$ to time $t = (n+1)k$ is the value of u we have previously calculated and the sample of BM in that time interval, and this is all we can use. The quantity $w^{n+1} - w^n$ in (3.35) is an increment of Brownian motion, therefore, it is a Gaussian variable with mean 0 and variance k.

Equation (3.35) says that $u^{n+1} - u^n + aku^n$ is a Gaussian variable with mean 0 and variance k. If u^n is known, then $P(x < u^{n+1} \leq x + dx)$ is

$$P(x < u^{n+1} \leq x + dx) = \frac{\exp\left(-\frac{(x-u^n+aku^n)^2}{2k}\right)}{\sqrt{2\pi k}}\, dx. \tag{3.36}$$

Since u^n is known, this is exactly the transition probability from the point u^n at time nk to the point x at time $(n+1)k$. If we write $u^n = x + y$, then the Chapman-Kolmogorov equation is

$$W(x,(n+1)k) = \int_{-\infty}^{+\infty} W(x+y,nk)\Psi(x,y,k)\, dy.$$

Replacing Ψ by the expression we have just derived gives

$$W(x,(n+1)k) = \int_{-\infty}^{+\infty} W(x+y,nk)\frac{\exp\left(-\frac{(-y+ak(x+y))^2}{2k}\right)}{\sqrt{2\pi k}}\, dy.$$

After rearranging the exponent in the above, we have

$$W(x,t+k) = \int_{-\infty}^{+\infty} W(x+y,t)\frac{\exp\left(-\frac{((1-ak)y-akx)^2}{2k}\right)}{\sqrt{2\pi k}} dy \tag{3.37}$$

where $t = nk$. The next step is to expand $W(x+y,t)$ around x. Up to fourth order we have

$$W(x+y,t) = W(x,t) + yW_x(x,t) + \frac{y^2}{2}W_{xx}(x,t) + \frac{y^3}{6}W_{xxx}(x,t) + O(y^4). \tag{3.38}$$

The expansion of $W(x+y,t)$ is substituted in (3.37) and we evaluate the different integrals that appear. Consider

$$I_1 = \int_{-\infty}^{+\infty} W(x,t)\frac{\exp\left(-\frac{((1-ak)y-akx)^2}{2k}\right)}{\sqrt{2\pi k}}\, dy.$$

To evaluate I_1 we make the change of variables $z = (1-ak)y$ and find

$$\begin{aligned} I_1 &= W(x,t)\int_{-\infty}^{+\infty} \frac{\exp\Big(-\frac{(z-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\frac{dz}{1-ak} \\ &= \frac{W(x,t)}{1-ak}\int_{-\infty}^{+\infty} \frac{\exp\Big(-\frac{(z-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\,dz \\ &= \frac{W(x,t)}{1-ak} \\ &= W(x,t)(1+ak+O(k^2)) \\ &= W(x,t)(1+ak)+O(k^2). \end{aligned} \tag{3.39}$$

The second integral is

$$I_2 = \int_{-\infty}^{+\infty} yW_x(x,t)\frac{\exp\Big(-\frac{((1-ak)y-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\,dy.$$

With the same change of variables we get

$$\begin{aligned} I_2 &= W_x(x,t)\int_{-\infty}^{+\infty} \frac{z}{1-ak}\,\frac{\exp\Big(-\frac{(z-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\,\frac{dz}{1-ak} \\ &= \frac{W_x(x,t)}{(1-ak)^2}akx \\ &= W_x(x,t)(1+2ak+O(k^2))akx \\ &= W_x(x,t)akx+O(k^2). \end{aligned} \tag{3.40}$$

The third integral is

$$I_3 = \int_{-\infty}^{+\infty} \frac{y^2}{2}\,W_{xx}(x,t)\,\frac{\exp\Big(-\frac{((1-ak)y-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\,dy. \tag{3.41}$$

The same change of variables gives

$$\begin{aligned} I_3 &= W_{xx}(x,t)\int_{-\infty}^{+\infty} \frac{z^2}{2(1-ak)^2}\,\frac{\exp\Big(-\frac{(z-akx)^2}{2k}\Big)}{\sqrt{2\pi k}}\,\frac{dz}{1-ak} \\ &= W_{xx}(x,t)\frac{1}{2(1-ak)^3}(k+(akx)^2) \\ &= W_{xx}(x,t)\frac{k}{2}+O(k^2). \end{aligned} \tag{3.42}$$

The fourth integral is

$$I_4 = \int_{-\infty}^{+\infty} \frac{y^3}{6} W_{xxx}(x,t) \frac{\exp\left(-\frac{((1-ak)y-akx)^2}{2k}\right)}{\sqrt{2\pi k}}\, dy, \tag{3.43}$$

which becomes

$$\begin{aligned} I_4 &= W_{xxx}(x,t) \int_{-\infty}^{+\infty} \frac{z^3}{6(1-ak)^3} \frac{\exp\left(-\frac{(z-akx)^2}{2k}\right)}{\sqrt{2\pi k}} \frac{dz}{1-ak} \\ &= W_{xxx}(x,t)\, \frac{1}{6(1-ak)^4}\, (3axk^2 + (akx)^3) \\ &= W_{xxx}(x,t) O(k^2). \end{aligned} \tag{3.44}$$

The fourth integral contributes only terms of order $O(k^2)$ and higher; the same is true of the next terms in the expansion, which have been omitted. Collecting (3.37), (3.38), (3.39), (3.40), (3.42), and (3.44), we find

$$W(x,t+k) = W(x,t) + W(x,t)ak + W_x(x,t)akx + \frac{k}{2} W_{xx}(x,t) + O(k^2),$$

and

$$\frac{W(x,t+k) - W(x,t)}{k} = W(x,t)a + W_x(x,t)ax + \frac{1}{2} W_{xx}(x,t) + O(k),$$

and, finally, as we make $k \to 0$,

$$W_t(x,t) = (axW(x,t))_x + \frac{1}{2} W_{xx}(x,t).$$

This is the Fokker-Planck equation corresponding to the solution of the Langevin equation (3.34). One can readily see that the first term on the right tends to concentrate W on ever smaller sets near the origin, corresponding to the effect of the damping in the Langevin equation; the second term spreads the support of W, as the diffusion by BM should do. A balance between concentration and spreading is reached when $\partial W/\partial t = 0$; the corresponding stationary solution for W is a Gaussian function, a fact that will be significant in chapter 6.

Given a Markovian stochastic process, we can construct its Fokker-Planck equation and vice versa. An interesting pair of a stochastic ordinary differential equation and the corresponding Fokker-Planck equation arises in two-dimensional incompressible fluid mechanics. If we consider a fluid having velocity $\mathbf{u} = (u, v)$ and vorticity $\xi = v_x - u_y$,

where (x, y) represents a point in physical space, then the equation for the evolution of the vorticity is

$$\frac{\partial \xi}{\partial t} + (\mathbf{u} \cdot \nabla)\xi = \frac{1}{Re}\Delta\xi, \tag{3.45}$$

where Re is the Reynolds number of the flow ($1/Re$ is a dimensionless measure of the viscosity, i.e., of the friction). If we assume that $\xi \geq 0$ and $\int \xi \, dx \, dy = 1$, then (3.45) is the Fokker Planck equation of the following system of stochastic ordinary differential equations:

$$d\mathbf{x} = \mathbf{u} \, dt + \sqrt{\frac{2}{Re}} \, d\mathbf{W}.$$

Here, $\mathbf{x}$ is the position of the point where the vorticity is ξ, and $\mathbf{W}$ is a two-dimensional Brownian motion. Each of these particles carries a fixed amount of vorticity the corresponding evolution of the density solves the vorticity partial differential equation. There is one equation per point in the support of ξ (i.e., for every point (x, y) such that $\xi(x, y) \neq 0$). The velocity $\mathbf{u}$ depends on the whole vorticity field at each instant t, so this equation is nonlinear and couples the BMs that correspond to different points in physical space, as one should expect given that the original equation of motion is nonlinear.

3.9. Solution of a Nonlinear Differential Equation by Branching Brownian Motion

So far, with the exception of the short comments at the end of the previous section, all the equations we have been solving have been linear. Now we give an example of how a variant of BM can be used to solve a nonlinear partial differential equation. The equation we work with is the Kolmogorov-Petrovskii-Piskunov (KPP) equation,

$$v_t - \tfrac{1}{2} v_{xx} = v^2 - v,$$

for which we prescribe initial data $v(x, t = 0) = \phi(x)$. This equation is an important model in combustion theory and in biology. We are looking for a representation of the solution v at a point (x, t) that relies on BM, as in earlier sections.

Start a BM w going backward in time from (x, t) and let it run until a time $t - t_1$, with t_1 drawn at random from the exponential density, $P(y < t_1 \leq y + dy) = \exp(-y) \, dy$. Start two independent BMs running backward from $(x + w(t_1), t - t_1)$, until new times $t - t_1 - t_{11}$ and $t - t_1 - t_{12}$ with t_{11} and t_{12} drawn independently from the exponential density. At each stopping time, split the branch of the BM into two independent BMs. If the time becomes negative for any branch, stop. The result is a backward tree with roots that cross the x-axis. Let the intersections of

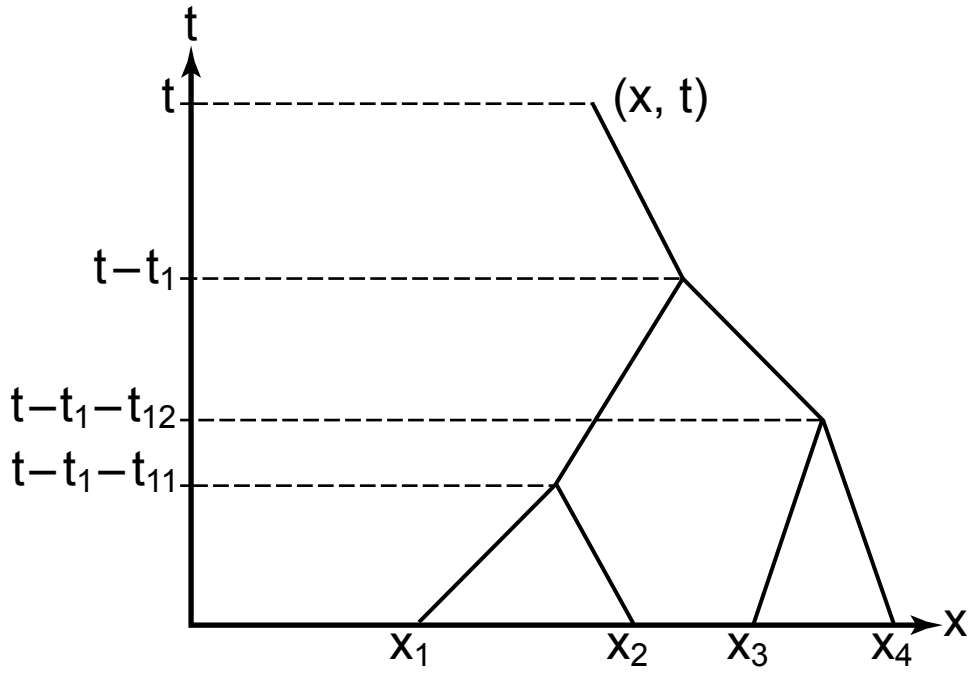

FIGURE 3.6. Branching Brownian motion.

the tree with the x-axis be $x_1, x_2, \ldots, x_n$, $n \geq 1$, and associate with the tree the product of initial values $\Xi = \phi(x_1)\phi(x_2)\cdots\phi(x_n)$; the claim is that the expected value of this product is the solution we want:

$$v(x,t) = E[\Xi] = E[\phi(x_1)\cdots\phi(x_n)]$$

(see Figure 3.6).

We take this opportunity to introduce a notation that will be widely used in chapter 6. Let Δ be the second derivative operator in the space variable x: $\Delta f = f_{xx}$ for a smooth function f. Just as the solution of the equation $v' - av = 0$, $v(0) = v_0$, $a =$ constant, is $e^{at}v_0$, we will symbolically write the solution of the heat equation $v_t - \frac{1}{2}\Delta v = 0$, $v(x,0) = \phi$, as $v(t) = e^{\frac{1}{2}t\Delta}\phi$ (this is the "semigroup" notation). For $v(x,t)$, which is the function $v(t)$ evaluated at x, we write $v(x,t) = (e^{\frac{1}{2}t\Delta}\phi)(x)$. We know that $(e^{\frac{1}{2}t\Delta}\phi)(x) = E[\phi(x + w(t))]$, where, as before, w is BM. One can readily understand the identity $e^{\frac{1}{2}(t+s)\Delta} = e^{\frac{1}{2}t\Delta}e^{\frac{1}{2}s\Delta}$ and check its validity (this is the "semigroup property").

We first check that the function $E[\Xi] = E[\phi(x_1)\cdots\phi(x_n)]$ satisfies the KPP equation. Write $V(x,t) = E[\Xi]$ with the backward branching walk starting at x, t. The probability that the first branching occurs at a time t_1 larger than t (so there is only one branch) is $\int_t^\infty e^{-s}\,ds = e^{-t}$

by definition; if this happens, the number Ξ attached to the tree is $\phi(x + w(t))$, whose expected value is $(e^{\frac{1}{2}t\Delta}\phi)(x)$.

Suppose to the contrary that t_1 occurs in a time interval $(s, s + ds)$ earlier than t (this happens with probability $e^{-s}\,ds$ by construction). Two branches of the tree start then at the point $(x + w(t_1), t - t_1)$. The two branches are independent by construction, and if we treat the point $(x + w(t_1), t - t_1)$ as fixed, the mean value of the product $E[\Xi]$ attached to each branch is $V(x + w(t_1), t - t_1)$, so that the mean value of $E[\Xi]$ at $(x + w(t_1), t - t_1)$ is $V^2((x + w(t_1), t - t_1)$. Now average $V^2((x + w(t_1), t - t_1)$ over $w(t_1)$, remembering the solution of the heat equation. This yields $e^{\frac{1}{2}s\Delta}V^2(t - s)$. Multiply this expression by the probability that the branching occurs at the time assumed, and sum over all first branching times between 0 and t.

Collecting all terms, we find

$$\begin{aligned} V = E[\Xi] &= e^{-t}e^{\frac{1}{2}t\Delta}\phi + \int_0^t e^{-s}e^{\frac{1}{2}s\Delta}V^2(t - s)\,ds \\ &= e^{-t}e^{\frac{1}{2}t\Delta}\phi + \int_0^t e^{s-t}e^{\frac{1}{2}(t-s)\Delta}V^2(s)\,ds, \end{aligned}$$

where the last identity is obtained by making the change of variables $s' = t - s$ and then dropping the prime on s. All that remains to be done is differentiate this expression for $V = E[\Xi]$ with respect to t, noting that $\Delta e^{-t} = e^{-t}\Delta$ (differentiation with respect to x and multiplication by a function of t commute), then calculate $\Delta E[\Xi]$ using the fact that $e^{\frac{1}{2}t\Delta}\phi$ and $e^{\frac{1}{2}(t-s)\Delta}V^2(s)$ are solutions of the heat equation; the equation we wish to solve appears. It is obvious that at $t = 0$ $E[\Xi] = \phi(x)$, and therefore $v(x, t) = V = E[\Xi]$ provided the solution of the KPP equation with given initial data is unique (and it is).

Figure 3.6 can be interpreted as a Feynman diagram (see Section 3.5): in picturesque language, one can say that an interaction with the nonlinear potential $u^2 - u$ has the effect of destroying an old particle and creating two new ones in its stead. Such interpretations are commonly encountered in physics.

3.10. A Brief Introduction to Stochastic ODEs

In Section 3.8 we solved a particular stochastic differential equation–the Langevin equation; we now make some comments about more general stochastic ordinary differential equations (SODEs) of the form

$$du = a(t, u(t))\,dt + b(t, u(t))\,dw, \tag{3.46}$$

where w is Brownian motion. The meaning of this equation is defined by

$$u(t) - u(0) = \int_0^t a(s, u(s))\, ds + \int_0^t b(s, u(s))\, dw.$$

The first integral is well defined, whereas, as we shall now see, the second is not. Integrals of the second form are called stochastic integrals. Let us figure out in what sense we can understand them.

Let $f(t)$ be a function defined on an interval $[a, b]$. A partition of $[a, b]$ is a set of points $\{t_i\}_{i=0}^n$ such that

$$a = t_0 < t_1 < t_2 < \cdots < t_n = b.$$

DEFINITION. The variation of $f(t)$ on $[a, b]$ is defined by

$$\text{Variation}(f(t)) = \sup_{\text{all partitions}} \sum_{i=0}^{n-1} |f(t_{i+1}) - f(t_i)|. \tag{3.47}$$

If the sup is finite f is said to have bounded variation; Brownian motion does not have bounded variation. Stieltjes integrals of the form $\int g(t)\, df(t)$ make sense only when the increment function f has bounded variation and, therefore,

$$\int_0^t b(s, u(s))\, dw$$

is not well defined as a Stieltjes integral.

The way to make sense of the stochastic integrals is to approximate $b(t, u(s))$ by a piecewise constant function; i.e.,

$$\int_0^t b(s, u(s))\, dw \approx \sum_{i=0}^{n-1} b_i\, dw_i = \sum_{i=0}^{n-1} b_i (w(t_{i+1}) - w(t_i)),$$

where $\{t_i\}_{i=0}^n$ is a partition of $[0, t]$, and then consider the limits of the sum as one makes the largest interval $t_i - t_{i-1}$ in the partition go to zero. Now one has to decide how to pick the b_i's. There are two common choices:

1. The b_i's are evaluated at the left ends of the intervals; i.e.,

$$b_i = b(t_i, u(t_i)).$$

2. The b_i's are the average of the endpoints

$$b_i = \frac{1}{2} \left[b(t_i, u(t_i)) + b(t_{i+1}, u(t_{i+1})) \right].$$

Choice 1 defines the Ito stochastic integral, whereas choice 2 defines the Stratonovich stochastic integral.

EXAMPLE. Suppose $b(t, u(t)) = w(t)$. Then, in the Ito case,

$$I_1 = \int_0^t w\, dw \approx \sum_{i=0}^{n-1} w(t_i)(w(t_{i+1}) - w(t_i)).$$

This is, of course, a random variable; the expected value of this random variable is zero, as one can see from the properties of BM:

$$E[I_1] = 0.$$

In the Stratonovich case, we find for the stochastic integral,

$$\begin{aligned} I_2 &= \int_0^t w\, dw \approx \sum_{i=0}^{n-1} \frac{1}{2}(w(t_{i+1}) + w(t_i))(w(t_{i+1}) - w(t_i)) \\ &= \sum_{i=0}^{n-1} \frac{1}{2}(w^2(t_{i+1}) - w^2(t_i)) \\ &= \frac{1}{2}\left[w^2(t_1) - w^2(t_0) + w^2(t_2) - w^2(t_1) + \cdots + w^2(t_n) - w^2(t_{n-1})\right] \\ &= \frac{1}{2}\left[w^2(t_n) - w^2(t_0)\right] = \frac{1}{2}w^2(t), \end{aligned}$$

and the expected value of this integral is

$$E[I_2] = \frac{t}{2}.$$

The fact that the expected values of the two integrals are so different is, of course, enough to show that the integrals themselves are different. This is very different from the situation in ordinary calculus, where the value of an integral is independent of the choice of points in the Riemann sums. How the stochastic integral is defined makes a big difference to the meaning of a stochastic differential equation. For the sake of definiteness, we will assume henceforth, when this makes a difference, that we are dealing with stochastic differential equations in the sense of Ito. When b in (3.46) is a constant (as has been the case so far in this volume), there is no ambiguity.

3.11. Exercises

1. Consider the partial differential equation $u_t = u_x$, with initial data $u(x, 0) = \phi(x)$. Solve it approximately as follows: Put a grid on the (x, t) plane, with mesh length h in the x-direction and k in the t-direction. Set $u_i^0 = \phi(ih)$. To calculate $u_{(i+1/2)h}^{(n+1/2)k}$ (halfway between

mesh points and halfway up the time interval k), proceed as follows: Pick a random number θ from the equidistribution density, one such choice for the whole half-step. Set

$$u_{(i+1/2)h}^{(n+1/2)k} = \begin{cases} u_i^n, & \theta \le \frac{1}{2} - \frac{k}{2h} \\ u_{i+1}^n, & \text{otherwise.} \end{cases}$$

The half-step from time $(n+1/2)k$ to $(n+1)k$ is similar. Show that if $k/h \le 1$, the solution of this scheme converges to the solution of the differential equation as $h \to 0$. (This is a special case of the Glimm or random choice scheme). Hint: The solution of the differential equation is $\phi(x+t)$ (i.e., initial values propagate along the lines $t = -x + \text{constant}$). Examine how the scheme propagates initial values: Show that an initial value u_i^0 moves in a time t by an amount η, where η is a random variable whose mean tends to $-t$ and whose variance tends to zero.

2. Consider the heat equation $v_t = (1/2)v_{xx}$ with initial data $v(x,0) = \phi(x)$ for $0 \le x \le 1$ and boundary conditions $v(0,t) = a$ and $v(1,t) = b$. Show that the solution at the point (x,t) can be obtained by starting BMs from (x,t) backward in time, attaching to each BM a number F as follows: If the BM first hits the portion of the x-axis between the boundaries $x = 0$ and $x = 1$, then $F = \phi(x + w(\omega,t))$; if it first hits the boundary $x = 0$, then $F = a$, and similarly at $x = 1$; finally, $v(x,t) = E[F]$. Hint: One way is to go through a finite-difference argument and then assume that the random walks converge to a BM.

3. Evaluate exactly $\int F\,dW$ for the following functionals F: (i) $F[w(\cdot)] = \int_0^1 w^4(s)\,ds$; (ii) $F = \sin(w^3(1))$, (iii) $F = \sin(w^2(1/2))\cos(w(1))$, (iv) $F = \int_0^{1/2} w^2(s)w^2(0.5+s)ds$.

4. Show that

$$\int dW(\int_0^t w^n(s)ds) = \int_{-\infty}^{+\infty} du \int_0^t ds(\sqrt{s}u)^n \exp(-u^2/2)/\sqrt{(2\pi)}$$

for all non-negative integers n.

5. Write the solution of the partial differential equation

$$v_t = (1/2)v_{xx} - xv,$$

with data $v(x,0) = \sin x$, as a Wiener integral.

6. Evaluate $\int F\,dW$, where

$$F[w(\cdot)] = e^{-\int_0^1 w^2(s)\,ds}\cos(w(1))$$

by Monte Carlo, along the following lines: Divide the time interval $[0,1]$ into n pieces and construct random walks w_n as follows: For t a multiple of $1/n$, set $w_n((i+1)h) = w_n(ih) + q$, where q is a Gaussian variable with mean 0 and variance $1/n$ (and, of course, $w_n(0) = 0$). For t between the multiples of $1/n$, construct w_n by linear interpolation. For each such w_n evaluate F and average over many walks, until the error (as measured by the difference between runs) is less than 1%. Do this for $n = 5$ and $n = 10$. Note that this integral is the solution at $(0,1)$ of some initial value problem for a differential equation. What is this problem?

7. In the previous problem, we discretized a Wiener integral by approximating the BMs by walks with n Gaussian increments. Write the solution of this discretized problem as an n-fold ordinary integral. (We shall see in chapter 5 how to evaluate such n-fold integrals, even for n large, by efficient Monte Carlo algorithms.)

8. Find the Fokker-Planck equation for the process that satisfies the equation $du = -dt + dw$, where w is Brownian motion. Does the pdf ever settle to a steady state?

9. Find a stochastic equation whose Fokker-Planck equation is $W_t = 5W + 5xW_x + 16W_{xx}$.

10. Consider particles moving in the plane, with coordinates that satisfy the pair of stochastic equations

$$dx_1 = a_1\,dt + dw_1, \quad dx_2 = a_2\,dt + dw_2,$$

where a_1, a_2 are constants and dw_1, dw_2 independent BMs. The density function $W = W(x,y,t)$ is the joint density of x_1 and x_2; find the partial differential equation (Fokker-Planck equation) that it satisfies.

11. Consider the tree in Figure 3.6. Let n be the (random) number of its intersections with the $x-$axis. Consider the function $u(t) = E[a^n]$, where $a > 0$ is a given constant. Show that u satisfies the equation $du/dt = u^2 - u$ with initial datum $u(0) = a$.

12. Consider again the tree in Figure 3.6. The set of branching points plus the set of intersections with the x–axis is the set of nodes of the tree. Associate with each intersection with the x-axis the given number $a > 0$. Each branching point X is attached to two nodes

below it, say Y and Z. If the number associated with Y is A and the number associated with Z is B, associate with X the number $AB+A$ (it is immaterial which point is Y and which is Z). Let D be the number associated with the first (from the top) branching point. Define $u(t) = E[D]$. Show that u satisfies the equation $du/dt = u^2$ with $u(0) = a$.

13. Prove that $e^{s\Delta}e^{t\Delta} = e^{(s+t)\Delta}$, where $\Delta = \partial^2/\partial x^2$. (You first have to figure out what this means and then check by means of formulas.)

14. Evaluate $e^{t\,\partial/\partial x}f$ for $f = \sin x$, at the point $x = 1, t = 1$.

15. Show that the solution of (3.34) is given by $u(\omega, t) = e^{-at}u(0) + w(\omega, t) - a\int_0^t e^{-a(t-\tau)}w(\omega, \tau)\,d\tau$.

16. Show that the solution of (3.35) is given by $u^n = (1-ak)^n u_0 + w^n - ak\,(w^{n+1} + (1-ak)w^{n-2} + \cdots + (1-ak)^{n-2}w^1)$.

3.12. Bibliography

[1] L. Arnold, *Stochastic Differential Equations*, Wiley, New York, 1973.

[2] R. Bhattacharya, L. Chen, S. Dobson, R. Guenther, C. Orum, M. Ossiander, E. Thomann, and E. Waymire, Majorizing kernels and stochastic cascades with applications to incompressible Navier-Stokes equations, Trans. Am. Math. Soc. 355 (2003), pp. 5003-5040.

[3] S. Chandrasekhar, Stochastic problems in physics and astronomy, Rev. Mod. Phys. 15 (1943), pp. 1-88; reprinted in N. Wax, *Selected Papers on Noise and Stochastic Processes*, Dover, New York, 1954.

[4] A.J. Chorin, Numerical study of slightly viscous flow, J. Fluid Mech. 57 (1973), pp. 785-796.

[5] A.J. Chorin, Vortex methods, in Les Houches Summer School of Theoretical Physics, 59, (1995), pp. 67-109.

[6] A.J. Chorin, Accurate evaluation of Wiener integrals, Math. Comp. 27 (1973), pp. 1-15.

[7] A.J. Chorin, Random choice solution of hyperbolic systems, J. Comput. Phys. 22, (1976), pp. 517-533.

[8] R. Feynman and A. Hibbs, *Quantum Mechanics and Path Integrals*, McGraw Hill, New York, 1965.

[9] C.W. Gardiner, *Handbook of Stochastic Methods*, Springer, New York, 1985.

[10] J. Glimm, Solution in the large of hyperbolic conservation laws, Comm. Pure Appl. Math. 18, (1965), pp. 69-82.
[11] O.H. Hald, Approximation of Wiener integrals, J. Comput. Phys. 69 (1987), pp. 460-470.
[12] K. Ito and H. McKean, *Diffusion Processes and Their Sample Paths*, Springer-Verlag, Berlin, 1974.
[13] M. Kac, *Probability and Related Topics in the Physical Sciences*, Interscience Publishers, London, 1959.
[14] P. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, Berlin, 1992.
[15] J. Lamperti, *Probability*, Benjamin, New York, 1966.
[16] H. McKean, Application of Brownian motion to the equation of Kolmogorov-Petrovskii-Piskunov, Commun. Pure Appl. Math. 27, (1975), pp. 323-331.
[17] S. Neftci, *An Introduction to the Mathematics of Financial Derivatives*, Academic, New York, 2000.
[18] R. Paley and N. Wiener, *Fourier Transforms in the Complex Domain*, Colloquium Publications Vol. 19, American Mathematical Society, Providence, RI, 1934.
[19] L. Schulman, *Techniques and Applications of Path Integration*, Wiley, New York, 1981.
[20] N. Wiener, *Nonlinear Problems in Random Theory*, MIT Press, Cambridge, MA, 1958.

CHAPTER 4

Stationary Stochastic Processes

4.1. Weak Definition of a Stochastic Process

This chapter is devoted to further topics in the theory of stochastic processes and of their applications. We start with a different, weaker, definition of a stochastic process, useful in the study of stationary processes.

Consider a collection of random variables $u(t,\omega) \in \mathbb{C}$ parametrized by t.

DEFINITION. We say that $u(t,\omega)$ is a weakly defined real valued stochastic process if for every finite set of points $t_1, \dots, t_n$, the joint distribution of $u(t_1,\omega)$, $\dots$, $u(t_n,\omega)$ is known:

$$F_{t_1,\dots,t_n}(y_1,\dots,y_n) = P(u(t_1) \le y_1, \dots, u(t_n) \le y_n).$$

The family of functions $F_{t_1,\dots,t_n}(y_1,\dots,y_n)$ must satisfy some natural requirements:

1. $F \ge 0$.
2. $F(\infty,\dots,\infty) = 1$ and $F(-\infty,\dots,-\infty) = 0$.
3. $F_{t_1,\dots,t_n}(y_1,\dots,y_m,\infty,\dots,\infty) = F_{t_1,\dots,t_m}(y_1,\dots,y_m)$.
4. If $(i_1,\dots,i_n)$ is a permutation of $(1,\dots,n)$, then

$$F_{t_{i_1},\dots,t_{i_n}}(y_{i_1},\dots,y_{i_n}) = F_{t_1,\dots,t_n}(y_1,\dots,y_n).$$

This definition has a natural extension to complex-valued processes.

A moment of $u(t,\omega)$ of order q is an object of the form

$$M_{i_1,\dots,i_n} = E[u^{i_1}(t_1)\cdots u^{i_n}(t_n)], \quad \sum_{j=1}^{n} i_j = q.$$

If a stochastic process has finite moments of order q, it is a process of order q. The moment

$$E[u(t,\omega)] = m(t)$$

is the mean of u at t. The function

$$E\left[(u(t_1,\omega) - m(t_1))\overline{(u(t_2,\omega) - m(t_2))}\right] = R(t_1,t_2)$$

is the covariance of u. Let us list the properties of the covariance of u:

1. $R(t_1, t_2) = \overline{R(t_2, t_1)}$.
2. $R(t_1, t_1) \geq 0$.
3. $|R(t_1, t_2)| \leq \sqrt{R(t_1, t_1)R(t_2, t_2)}$.
4. For all $t_1, \ldots, t_n$ and all $z_1, \ldots, z_n \in \mathbb{C}$,

$$\sum_{i=1}^{n} \sum_{j=1}^{n} R(t_i, t_j) z_i \overline{z_j} \geq 0.$$

The first three properties are easy to establish; the fourth is proved as follows: For any choice of complex numbers z_j, the sum

$$\sum_{i=1}^{n} \sum_{j=1}^{n} R(t_i, t_j) z_i \overline{z_j}$$

is by definition equal to

$$E\left[\left|\sum_{j=1}^{n} \left(u(t_j) - m(t_j)\right) z_j\right|^2\right] \geq 0$$

(i.e., to the expected value of a non-negative quantity).

Definition. A weakly defined stochastic process is stationary in the strict sense if for every $t_1, \ldots, t_n$ and for any $T \in \mathbb{R}$,

$$F_{t_1, \ldots, t_n}(y_1, \ldots, y_n) = F_{t_1+T, \ldots, t_n+T}(y_1, \ldots, y_n).$$

For a stochastic process that is stationary in this sense, all moments are constant in time and, in particular, $m(t) = m$ and $R(t_1, t_2) = R(t_1+T, t_2+T)$ for all T. Choose $T = -t_2$; then $R(t_1, t_2) = R(t_1-t_2, 0)$, and it becomes reasonable to define

$$R(t_2 - t_1) = R(t_2, t_1),$$

where the function R on the left side, which has only one argument, is also called R with the hope that there is no ambiguity. Note that $R(T) = R(t + T, t)$.

The above properties become, for the new function R,

1. $R(t) = \overline{R(-t)}$.
2. $R(0) \geq 0$.
3. $|R(t)| \leq R(0)$.
4. For all $t_1, \ldots, t_n$ and all $z_1, \ldots, z_n \in \mathbb{C}$,

$$\sum_{i}^{n} \sum_{j}^{n} R(t_j - t_i) \overline{z_i} z_j \geq 0. \tag{4.1}$$

DEFINITION. A stochastic process is stationary in the wide sense if it has a constant mean and its covariance depends only on the difference between the arguments; i.e.,

1. $m(t) = m$.
2. $R(t_2, t_1) = R(t_2 - t_1)$.

If a stochastic process is stationary in the weak sense and Gaussian, then it is stationary in the strict sense (because a Gaussian process is fully determined by its mean and covariances). Brownian motion is *not* stationary. White noise is stationary (but ill-defined without appeal to distributions).

We now consider some instances of processes that are stationary in the weak sense. Pick $\xi \in \mathbb{C}$ to be a random variable and $h(t)$ a nonrandom function of time, and consider the process $u(t, \omega) = \xi h(t)$. Assume for simplicity that $h(t)$ is differentiable, and determine when a process of this type is stationary in the wide sense. Its mean is

$$m(t) = E[\xi h(t)] = h(t)E[\xi],$$

which is constant if and only if $h(t)$ is constant or $E[\xi] = 0$. Suppose $E[\xi] = 0$. The covariance

$$R(t_2, t_1) = E[\xi h(t_2)\overline{\xi}\,\overline{h(t_1)}] = E[\xi\overline{\xi}]h(t_2)\overline{h(t_1)}$$

must depend only on the difference $t_2 - t_1$. Consider the special case $t_1 = t_2 = t$. In this case, the covariance $E[\xi\overline{\xi}]h(t)\overline{h(t)}$ must be $R(0)$; hence, $h(t)\overline{h(t)}$ must be constant. Therefore, $h(t)$ is of the form

$$h(t) = Ae^{i\phi(t)}.$$

Now we narrow the possibilities some more. Suppose h has the form $Ae^{i\phi(t)}$ with $A \neq 0$. Then

$$R(t_2 - t_1) = |A|^2 E[\xi\overline{\xi}]e^{i\phi(t_2) - i\phi(t_1)}.$$

Set $t_2 - t_1 = T$ and $t_1 = t$. Then

$$R(T) = |A|^2 E[\xi\overline{\xi}]e^{i[\phi(t+T) - \phi(t)]}$$

for all t, T. Since $R(T) = \overline{R(-T)}$ we see that

$$\frac{\phi(t+T) - 2\phi(t) + \phi(t-T)}{T^2} = 0.$$

Letting $T \to 0$ gives $\phi''(t) = 0$ for all t so $\phi(t) = \alpha t + \beta$. We have therefore shown that the process $u(t, \omega) = \xi h(t)$ is stationary in the wide sense if $h(t) = Ce^{i\alpha t}$ and $E[\xi] = 0$ or if $h(t)$ is constant.

4.2. Covariance and Spectrum

In the last section, we presented an example of a stationary stochastic process in the wide sense, given by $u(t,\omega) = \xi e^{i\lambda t}$, where ξ is a random variable with mean 0. This stochastic process has a covariance of the form

$$R(T) = R(t_2, t_1) = R(t_2 - t_1) = E[|\xi|^2]e^{i\lambda T},$$

where $T = t_2 - t_1$. Now we want to generalize this example. First, we try to construct a process of the form

$$u(t,\omega) = \xi_1 e^{i\lambda_1 t} + \xi_2 e^{i\lambda_2 t},$$

with $\lambda_1 \neq \lambda_2$. Then $E[u] = E[\xi_1]e^{i\lambda_1 t} + E[\xi_2]e^{i\lambda_2 t}$, which is independent of t if $E[\xi_1] = E[\xi_2] = 0$. The covariance is

$$\begin{aligned} &E\left[(\xi_1 e^{i\lambda_1 t_2} + \xi_2 e^{i\lambda_2 t_2})(\overline{\xi_1} e^{-i\lambda_1 t_1} + \overline{\xi}_2 e^{-i\lambda_2 t_1})\right] \\ &\quad = E\left[|\xi_1|^2 e^{i\lambda_1 T} + |\xi_2|^2 e^{i\lambda_2 T} + \xi_1\overline{\xi}_2 e^{i\lambda_1 t_2 - i\lambda_2 t_1} + \overline{\xi}_1\xi_2 e^{i\lambda_2 t_2 - i\lambda_1 t_1}\right], \end{aligned}$$

which can be stationary only if $E[\xi_1\overline{\xi}_2] = 0$. Then $u(t,\omega)$ is stationary and

$$R(T) = E[|\xi_1|^2]e^{i\lambda_1 T} + E[|\xi_2|^2]e^{i\lambda_2 T}.$$

More generally, a process $u = \sum_j \xi_j e^{i\lambda_j t}$ is wide sense stationary if $E[\xi_j\overline{\xi_k}] = 0$ when $j \neq k$ and $E[\xi_i] = 0$. In this case,

$$R(T) = \sum E\left[|\xi_j|^2\right] e^{i\lambda_j T}.$$

This expression can be rewritten in a more useful form as a Stieltjes integral. Recall that when g is a nondecreasing function of x, the Stieltjes integral of a function f with respect to g is defined to be

$$\int f\, dg = \lim_{\max\{x_{i+1} - x_i\} \to 0} \sum f(x_i^*)[g(x_{i+1}) - g(x_i)],$$

where $x_i \leq x_i^* \leq x_{i+1}$. If g is differentiable, then

$$\int_a^b f\, dg = \int_a^b f g'\, dx.$$

Suppose $g(x)$ is the step function

$$g(x) = \begin{cases} 0, & x < c \\ q, & x \geq c. \end{cases}$$

with $a \le c \le b$. Then $\int_a^b f\,dg = f(c)q$ if f is continuous at c. We define the function $F = F(k)$ by

$$F(k) = \sum_{\{j|\lambda_j \le k\}} E[|\xi_j|^2];$$

i.e., $F(k)$ is the sum of the expected values of the squares of the amplitudes of the complex exponentials with frequencies less than or equal to k. $R(T)$ becomes

$$R(T) = \int_{-\infty}^{+\infty} e^{ikT}\,dF(k).$$

We shall now see that under some technical assumptions, this relation holds for all wide sense stationary stochastic processes. Indeed, we have the following:

THEOREM 4.1 (Khinchin).

1. *If $R(T)$ is the covariance of a weakly defined wide sense stationary stochastic process such that*
$$\lim_{h\to 0} E\left[|u(t+h) - u(t)|^2\right] = 0,$$
then $R(T) = \int e^{ikT}\,dF(k)$ for some nondecreasing function $F(k)$.
2. *If a function $R(T)$ can be written as $\int e^{ikT}\,dF(k)$ for some nondecreasing function F, then there exists a weakly defined wide sense stationary stochastic process, satisfying the condition in part 1 of the theorem, that has $R(T)$ as its covariance.*

Khinchin's theorem follows from the inequalities we have proved for R; indeed, one can show (and we will not do so here) that a function that satisfies these inequalities is the Fourier transform of a non-negative function. If it so happens that $dF(k) = \phi(k)\,dk$, then $R(T) = \int e^{ikT}\phi(k)\,dk$ and $\phi(k)$ is called the spectral density of the process. Thus, Khinchin's theorem states that the covariance function is a Fourier transform of the spectral density. Hence, if we know $R(T)$, we can compute the spectral density by

$$\phi(k) = \frac{1}{2\pi}\int_{-\infty}^{+\infty} e^{-ikT}R(T)\,dT.$$

EXAMPLE. In the case of white noise we have $R(T) = \delta(T)$. Its spectral density (interpreted carefully) is $\phi(k) = 1/2\pi$; thus, all frequencies have the same amplitude. The adjective "white" comes from the fact that in white light, all frequencies are present with the same amplitude. Any random signal that is not white noise is called colored noise.

4.3. Scaling and the Inertial Spectrum of Turbulence

To illustrate these constructions, we now derive the "inertial range" spectrum of fully developed turbulence. The equations of motion will not be written down because they will not be used directly.

Consider turbulence far from walls, with the Reynolds number $Re = U\ell_0/\nu$ very large, where U is a typical velocity difference in the flow, ℓ_0 is a length scale for the flow, and ν is the viscosity; the dimensionless number Re measures the amount by which the "inertial" (i.e., nonlinear) terms in the equations of motion dominate the viscous dissipation and is large in fully developed turbulence. The movement of energy from scale to scale (i.e., from one wave number k to another), is described by the nonlinear terms in the equation of motion. The flow is driven by large-scale forcing (e.g., in the case of meteorology, by the rotation of the Earth around its axis and around the Sun); one assumes that as the energy moves to large wave numbers k (i.e., small wavelengths), the geometry of the forcing is forgotten and the flow can be viewed as approximately homogeneous (translation invariant) and isotropic (rotation invariant) on the small scales, and its spectral properties at large k are universal (i.e., independent of specific geometry and forcing).

The velocity field in three space dimensions is a vector quantity: $\mathbf{u} = (u_1, u_2, u_3)$. One can define a covariance tensor

$$R_{ij}(\mathbf{r}) = E[u_i(\mathbf{x})u_j(\mathbf{x}+\mathbf{r})],$$

where $\mathbf{r}$ is a three-component vector; then

$$R_{ii}(\mathbf{r}) = \int_{-\infty}^{\infty} e^{i\mathbf{k}\cdot\mathbf{r}}\, dF_{ii}(\mathbf{k}),$$

where $\mathbf{k} = (k_1, k_2, k_3)$ and $\mathbf{k}\cdot\mathbf{r}$ is the ordinary Euclidean inner product. Without loss of generality in what follows, one can write $dF_{ii}(\mathbf{k}) = \Psi_{ii}(\mathbf{k})\, dk_1\, dk_2\, dk_3$ (this is so because all we will care about is the dimensions of the various quantities, which is not affected by their possible lack of smoothness). Finally, one can define

$$E(k) = \int_{k_1^2+k_2^2+k_3^2=k^2} (\Psi_{11} + \Psi_{22} + \Psi_{33})\, dS(\mathbf{k}).$$

$E(k)$ is the average of the $\Psi_{ii}(\mathbf{k})$ over the sphere of radius k in wave number space and is a function of only $k = \sqrt{k_1^2 + k_2^2 + k_3^2}$. One can see from the identities in the previous section that $E[u^2] = \int_0^\infty E(k)\, dk$, where $u^2 = u_1^2 + u_2^2 + u_3^2$; u^2 is the kinetic energy of the flow and $E(k)$ is, by definition, the spectrum of the flow.

The kinetic energy of the flow is proportional to the square of the velocity, whereas energy dissipation, modeled in the equations of motion by the Laplace operator acting on the velocity, is proportional to the square of the derivatives of the velocity; in spectral variables (i.e., after Fourier transformation), the kinetic energy is proportional to the square of $\hat{u}(k)$, whereas the dissipation is proportional to the square of $k\hat{u}(k)$, where $\hat{u}$ is the (random) Fourier transform of the velocity $u = u(x)$. It is plausible that when Re is large, the kinetic energy resides in a range of k's disjoint from the range of k's where the dissipation is taking place, and, indeed, experimental data show it to be so; specifically, there exist wave numbers k_1 and k_2 such that

$$\int_0^{k_1} E(k)\,dk \sim \int_0^{\infty} E(k)\,dk, \quad \int_{k_2}^{\infty} k^2 E(k)\,dk \sim \int_0^{\infty} k^2 E(k)\,dk,$$

with $k_1 \ll k_2$. The range of k's such that $k_1 < k < k_2$ is the "inertial range" of wave numbers; the name is a bit of a misnomer because it implies that in that range, the mechanics is purely "inertial," free of viscosity effects; but we shall see that this is not so. This is the range of wave numbers k we now focus on.

We will be relying on dimensional analysis. Suppose a variable a is a function of variables $a_1, a_2, \dots, a_m, b_1, b_2, \dots, b_k$, where $a_1, \dots, a_m$ have independent units (for example a_1 could be a length and a_2 could be a time), while the units of $b_1, \dots, b_k$, can be formed from the units of $a_1, a_2, \dots, a_m$; in the example just used, b_1 could be a velocity. Then there exist dimensionless variables

$$\Pi = \frac{a}{a_1^{\alpha_1} \cdots a_m^{\alpha_m}}, \quad \Pi_i = \frac{b_i}{a_1^{\alpha_{i1}} \cdots a_m^{\alpha_{im}}}, \quad i = 1, \dots, k,$$

where the α_i, α_{ij} are simple fractions, such that Π is a function of the Π_i:

$$\Pi = \Phi(\Pi_1, \dots, \Pi_k).$$

This is just a consequence of the requirement that a physical relationship be independent of the magnitude of the units of measurement. At this stage, nothing can be said about the function Φ. Suppose the variables Π_i are small or large (the two cases are indistinguishable, because an unknown function of x is also an unknown function of $1/x$) and assume that the function Φ has a nonzero finite limit as its arguments tend to zero or to infinity; then $\Pi \sim$ constant, and one finds a power monomial relation between a and the a_i: $a = a_1^{\alpha_1} \cdots a_m^{\alpha_m}$. This is a "complete similarity" relation. If the function Φ does not have the assumed limit, it may happen that for Π_1 small or large, $\Phi(\Pi_1) = \Pi_1^{\alpha} \Phi_1(\Pi_1) + \cdots$, where the dots denote lower-order terms,

α is a constant, the other arguments of Φ have been omitted and, Φ_1 has a finite nonzero limit. One can then obtain a monomial expression for a in terms of the a_i and b_i, with undetermined powers that must be found by means other than dimensional analysis. The resulting power relation is an 'incomplete" similarity relation. The exponent α is known in the physics literature as an anomalous scaling exponent; in physics, incomplete similarity is usually discussed in the context of the renormalization group; see chapter 5. Of course, one may well have functions Φ with neither kind of similarity.

We now apply these scaling ideas to the spectrum. The spectrum in the inertial range $E(k)$ is a function of k, of the viscosity ν, of the length scale ℓ_0, of the amplitude U of the typical velocity in the flow, and of the rate of energy dissipation ϵ. That last variable belongs here because energy is transferred from the low k domain through the inertial range into the large k domain where it is dissipated; the fact that ϵ belongs in the list was the brilliant insight of Kolmogorov.

Our basic units are the units of length L and of time T. The units of the viscosity are L^2/T, those of ϵ are L^2/T^3, those of k are $1/L$, and the equation $E[u^2] = \int E(k)\,dk$ shows that the units of E are L^3/T^2. Dimensional analysis yields $E(k)(\epsilon^{-2/3}k^{5/3}) = \Phi(Re, \ell_0 k)$ for some unknown function Φ of the two large arguments Re and $\ell_0 k$; Re is large because this is the condition for fully developed turbulence to appear and $\ell_0 k$ is large in the inertial range of scales. If the function Φ has a finite nonzero limit C as its arguments grow (an assumption of complete similarity), one can deduce $E(k) = C\epsilon^{2/3}k^{-5/3}$—the famous Kolmogorov-Obukhov scaling law for the inertial range of fully developed turbulence, the cornerstone of turbulence theory.

This law is not fully satisfactory for various reasons, and a number of correction schemes have been proposed over the years. In recent years, it has been argued that the unknown function Φ behaves, as its arguments tend to infinity, as $C(Re)(\ell_0 k)^{-d/\log(Re)}\Phi_0(\ell_0 k, Re)$, where it is Φ_0 that tends to a non-zero constant as its arguments grow, $C(Re)$ is a function of Re, and d is a positive constant; the exponent $-d/\log(Re)$ is an anomalous exponent. This is an assumption of incomplete similarity, which leads, for large Re and $\ell_0 k$, to the relation

$$E(k) = C(Re)\epsilon^{2/3}k^{-5/3}(\ell_0 k)^{-d/\log(Re)}.$$

4.4. Random Measures and Random Fourier Transforms

We now show that arbitrary wide sense stationary processes can be represented as convolutions of nonrandom functions with certain simple processes (often Brownian motion). An important special case of this representation is the stochastic Fourier transform, which exists whenever the covariance function exists and does not require that the process itself have samples to which the standard Fourier transform can be applied; this is a key building block in the study of turbulence, signal processing, and quantum theory.

Given a probability space $(\Omega, \mathcal{B}, P)$, consider the set of random variables $f(\omega)$, where ω is in Ω, such that $E[f\bar{f}] < \infty$. We refer to this set as $L_2(\Omega, \mathcal{B}, P)$. We now construct a one-to-one mapping $L_2(\Omega, \mathcal{B}, P) \rightarrow L_2(A, \mu)$, where A is a subset of the t-axis and μ is a measure on A. Consider $\mathcal{A}$, an algebra of subsets of A, i.e, a collection of sets with the property that if the sets A_i are in $\mathcal{A}$, then so are their complements, and their finite unions and intersections; an algebra is much like a σ-algebra, with the exception that we do not require that the union of a countably infinite family of subsets belong to the algebra, a detail which is important in a rigorous analysis, but which we will disregard here.

Consider the triple $(A, \mathcal{A}, \mu)$, where μ is a rule which to each subset $A_i \in \mathcal{A}$ assigns a number such that

1. $\mu(A_i) \geq 0$.
2. $\mu(A_i)$ is finite.
3. $\mu(\emptyset) = 0$.
4. $A_i \cap A_j = \emptyset \Rightarrow \mu(A_i \cup A_j) = \mu(A_i) + \mu(A_j)$.

(Again, note that we are concerned only with finitely many A_i.) Next, construct a random variable $\rho = \rho(A_i, \omega)$, where $A_i \in \mathcal{A}$ and $\omega \in \Omega$ (remember that a random variable is a function defined on Ω) that has the following properties:

1. $A_i \cap A_j = \emptyset \Rightarrow \rho(A_i \cup A_j, \omega) = \rho(A_i, \omega) + \rho(A_j, \omega)$.
2. $\rho(A_i, \omega)$ is square integrable; i.e., $E[\rho(A_i, \omega)\bar{\rho}(A_i, \omega)] < \infty$.
3. $\rho(\emptyset, \omega) = 0$.
4. $A_i, A_j \subset A \Rightarrow E[\rho(A_i, \omega)\bar{\rho}(A_j, \omega)] = \mu(A_i \cap A_j)$.

The properties listed above imply that $\mu(A_i) \geq 0$ for all $A_i \epsilon \mathcal{A}$, since

$$\mu(A_i) = \mu(A_i \cap A_i) = E[\rho(A_i, \omega)\bar{\rho}(A_i, \omega)] \geq 0.$$

μ is called the structure function of ρ. Just as a stochastic process is a function of both ω and t, so is a random measure a function of both ω and the subsets A_i of A.

Now define $\chi_{A_i} = \chi_{A_i}(t)$, the characteristic function of the subset A_i of the t-axis, to be

$$\chi_{A_i} = \begin{cases} 1, & t \in A_i \\ 0, & \text{otherwise,} \end{cases}$$

and consider a function $q(t)$ of the form

$$q(t) = \sum c_i \chi_{A_i}(t).$$

We consider the case where $\{A_i\}$ is a finite partition of A; i.e., there are only finitely many A_i, $A_i \cap A_j = \emptyset$ for $i \neq j$, and $\bigcup A_i = A$. Thus, $q(t)$ takes on only a finite number of values. To this function $q(t)$ assign the random variable

$$f(\omega) = \sum c_i \rho(A_i, \omega).$$

Hence, each characteristic function of a subset is replaced by the random variable that the random measure assigns to the same subset; thus, this substitution transforms a function of t into a function of ω (i.e., into a random variable).

Now consider the product $q_1(t)\overline{q_2}(t)$ of two functions of the form

$$q_1 = \sum_{j=1}^{n} c_j \chi_{A_j}(t), \quad q_2 = \sum_{k=1}^{m} d_k \chi_{B_k}(t),$$

where the $\{B_i\}$ is another finite partition of A. It is not necessary for n and m to be equal. There is a finite number of intersections of the A_j's and B_k's, and on each of these subsets, the product

$$\begin{aligned} q_1\overline{q_2} &= \left(\sum_{j=1}^{n} c_j \chi_{A_j} \right) \left(\sum_{k=1}^{m} \overline{d_k} \chi_{B_k} \right) \\ &= \sum_{j=1}^{n} \sum_{k=1}^{m} c_j \overline{d_k} \chi_{A_j \cap B_k}, \end{aligned}$$

takes on a constant value $c_j\overline{d_k}$. Thus, the same construction allows us to assign a random variable $f_1\overline{f_2}$ to the product $q_1\overline{q_2}$. Since

$$f_1(\omega) = \sum c_j \rho(A_j, \omega), \quad f_2(\omega) = \sum d_k \rho(B_k, \omega),$$

we conclude that

$$\begin{aligned} E[f_1\overline{f_2}] &= E\left[\sum_{j=1}^{n}\sum_{k=1}^{m} c_j\overline{d_k}\rho(A_j,\omega)\bar{\rho}(B_k,\omega)\right] \\ &= \sum_{j=1}^{n}\sum_{k=1}^{m} c_j\overline{d_k}E\left[\rho(A_j,\omega)\bar{\rho}(B_k,\omega)\right] \\ &= \sum_{j=1}^{n}\sum_{k=1}^{m} c_j\overline{d_k}\mu(A_j\cap B_k) \\ &= \int q_1\overline{q_2}\mu(dt). \end{aligned} \tag{4.2}$$

Thus we have established a mapping between random variables with finite mean squares and functions of time with finite square integrals (i.e., between the random variables $f(\omega)$ and functions $q(t)$ such that $\int q_1(t)\overline{q_2}(t)\mu(dt)$ is finite). Although we have defined the mapping only for functions $q(t) = \sum c_i\chi_{A_i}(t)$, an argument that we omit enables us to extend the mapping to all random variables and functions of t with the square integrability properties listed above.

Example. We now show in detail how this construction works for a very special case. Say we are given a probability space (Ω, B, P) and three subsets of the t-axis: $A_1 = [0,1), A_2 = [1,3)$, and $A_3 = [3, 3\frac{1}{2}]$. Each A_i is assigned a real valued random variable $\rho_i(\omega) = \rho(A_i,\omega)$ that has mean 0 and variance equal to the length of A_i. For example, $\rho_1(\omega)$ has mean 0 and variance 1, and so forth. The variables ρ_1, ρ_2, and ρ_3 are independent, and $E[\rho_i\rho_j] = 0$ for $i \neq j$, where $E[\rho_i^2]$ is the length of the i^{th} interval. Moreover,

$$\chi_1 = \begin{cases} 1, & 0 \le t < 1 \\ 0, & \text{elsewhere,} \end{cases}$$

$$\chi_2 = \begin{cases} 1, & 1 \le t < 3 \\ 0, & \text{elsewhere,} \end{cases}$$

$$\chi_3 = \begin{cases} 1, & 3 \le t \le 3\frac{1}{2} \\ 0, & \text{elsewhere,} \end{cases}$$

where $\int \chi_i\chi_j\,dt = 0$ for $i \neq j$ and $\int \chi_i^2\,dt$ is the length of the i^{th} interval.

Now take a function of the form $q_1(t) = \sum_i c_i \chi_i(t)$, where the c_i's are constants. Then

$$q_1(t) \to f_1(\omega) = \sum_{i=1}^{3} c_i \rho_i(\omega).$$

Suppose we have another function $q_2(t)$ on the same partition:

$$q_2(t) = \sum_{j=1}^{3} d_j \chi_j(t) \to f_2(\omega) = \sum_{j=1}^{3} d_j \rho_j(\omega).$$

Then

$$\begin{aligned} E[f_1 \overline{f_2}] &= E\left[\sum_{i=1}^{3}\sum_{j=1}^{3} c_i \overline{d_j} \rho_i \rho_j\right] \\ &= \sum_{j=1}^{3} c_j \overline{d_j} E\left[\rho_j^2\right] \\ &= \sum_{j=1}^{3} c_j \overline{d_j} \mu(A_j), \end{aligned}$$

where $\mu(A_j)$ is the length of A_j. Notice also that

$$\begin{aligned} \int_0^{3\frac{1}{2}} q_1(t)\overline{q_2}(t)\, dt &= \int_0^{3\frac{1}{2}} \sum_{i=1}^{3}\sum_{j=1}^{3} c_i \overline{d_j} \chi_i(t)\chi_j(t)\, dt \\ &= \sum_j c_j \overline{d_j} \mu(A_j), \end{aligned}$$

which verifies that $q(t) \to f(\omega)$, so $E[f_1 \overline{f_2}] = \int q_1(t)\overline{q_2}(t)\mu(dt)$ as in (4.2).

Now approximate every square integrable function q on A (i.e., such that $\int_A q\bar{q}\, d\mu$ is finite) by a step function, construct the corresponding random variable, and take the limit, as the approximation improves, of the sequence of random variables obtained in this way. This makes for a mapping of square integrable functions on A onto random variables with finite mean squares. This mapping can be written as

$$f(\omega) = \int q(s)\rho(ds, \omega)$$

(the right-hand side is an integral with respect to the measure ρ), where the variable t has been replaced by s for convenience. Now view a stochastic process u as a family of random variables labeled by the

parameter t (i.e., there is a random variable u for every value of t) and apply the representation just derived at each value of t; so,

$$u(t,\omega) = \int q(t,s)\rho(ds,\omega).$$

Assume $u(t,\omega)$ is wide sense stationary. Then the covariance of u is

$$\begin{aligned} R(t_1 - t_2) &= E[u(t_1,\omega)\overline{u}(t_2,\omega)] \\ &= E\left[\int q(t_1,s_1)\rho(ds_1)\int \bar{q}(t_2,s_2)\bar{\rho}(ds_2)\right] \\ &= E\left[\int q(t_1,s_1)\bar{q}(t_2,s_2)\rho(ds_1)\bar{\rho}(ds_2)\right] \\ &= \int q(t_1,s_1)\bar{q}(t_2,s_2)E[\rho(ds_1)\bar{\rho}(ds_2)] \\ &= \int q(t_1,s)\bar{q}(t_2,s)\mu(ds). \end{aligned}$$

One can show that the converse is also true: If the last equation holds, then $u(t,\omega) = \int q(t,s)\rho(ds,\omega)$ with $E[\rho(ds)\bar{\rho}(ds)] = \mu(ds)$. Note that in all of the above, the equality holds in a mean square (L_2) sense and little can be said about the higher moments.

EXAMPLE. If $u = u(t,\omega)$ is a wide sense stochastic process, then it follows from Khinchin's theorem that

$$R(T) = E[u(t+T,\omega)\overline{u(t,\omega)}] \tag{4.3}$$

$$= \int e^{ikT}dF(k). \tag{4.4}$$

Conversely, if $E[\rho(dk)\overline{\rho(dk)}] = dF(k)$, we see that if

$$u(t,\omega) = \int e^{ikt}\rho(dk,\omega),$$

then

$$\begin{aligned} E[u(t+T,\omega)\overline{u(t,\omega)}] &= \int e^{ik(t+T-t)}E[\rho(dk)\overline{\rho(dk)}] \\ &= \int e^{ikT}dF(k). \end{aligned}$$

We have just shown that $dF(k)$ is the energy density in the interval dk. This $\rho(k)$ is the stochastic Fourier transform of u. The inverse Fourier transform does not exist in the usual sense (i.e., $\int u(t,\omega)e^{-ikt}\,dt$ for each ω does not exist), but for (4.4) to hold, it is sufficient for $E[|u(t)|^2]$ to exist for each t.

One can summarize the construction of the stochastic Fourier transform as follows: For the ordinary Fourier transform, the Parseval identity is a consequence of the definitions. To generalize the Fourier transform, we started from a general form of the Parseval identity and found a generalized version of the Fourier transform that satisfies it.

EXAMPLE. Suppose $dF(k) = \phi(k)\,dk$. Then

$$\int e^{ik(t_2-t_1)} dF(k) = \int e^{ikt_2}\sqrt{\phi(k)}e^{-ikt_1}\sqrt{\phi(k)}\,dk.$$

Recall that $\phi(k) \geq 0$. Write $\sqrt{\phi(k)} = \hat{h}(k) = \widehat{h(t)}(k)$, where $h(t)$ is the inverse Fourier transform of $\hat{h}(k)$, $\hat{h}(k) = \frac{1}{\sqrt{2\pi}}\int h(t)e^{-ikt}dt$; then

$$\begin{aligned} e^{-ikt_2}\sqrt{\phi(k)} &= e^{-ikt_2}\frac{1}{\sqrt{2\pi}}\int h(t)e^{-ikt}dt \\ &= \frac{1}{\sqrt{2\pi}}\int h(t)e^{-ik(t+t_2)}dt \\ &= \frac{1}{\sqrt{2\pi}}\int h(t-t_2)e^{-ikt}dt \\ &= \widehat{h(t-t_2)}(k), \end{aligned}$$

where the last parenthesis is there to remind you that $\widehat{h(t-t_2)}$ is a function of k. Since the Fourier transform preserves inner products, we find that

$$R(t_1,t_2) = \int \bar{h}(t-t_1)h(t-t_2)\,dt,$$

and by changing t to s, we obtain

$$R(t_1,t_2) = \int \bar{h}(s-t_1)h(s-t_2)\mu(ds),$$

where $\mu(ds) = ds$. Applying our representation, we get $u(t,\omega) = \int \bar{h}(s-t)\rho(ds)$, where $E[|\rho(ds)|^2] = ds$. The random measure constructed as increments of Brownian motion at instants ds apart has this property. Thus, any wide sense stationary stochastic process with $dF(k) = \phi(k)\,dk$ can be approximated as a sum of translates (in time) of a fixed function, each translate multiplied by independent Gaussian random variables. This is the "moving average" representation.

4.5. Prediction for Stationary Stochastic Processes

Suppose we are observing a stochastic process $u(t,\omega)$, have been observing it long enough to know it is stationary and to determine its temporal covariances, and suppose we are given observed values of $U(s)$

of $u(t,\omega)$ for for $s \leq t$ (we denote observed values by capital letters.) The question we address in this section is how to predict a value for $u(t+T,\omega)$ based on the information we have. For simplicity, we shall do so only for a stationary random sequence.

DEFINITION. A stationary random sequence is a collection $u(t,\omega)$ of random variables for $t = 0,1,2,3,\ldots$ as well as for $t = -1,-2,-3,\ldots$ such that the joint distribution of any subset is known, subject to the obvious compatibility conditions, and such that all the distributions are invariant under the transformation $t \to t+T$ for T integer. Such sequences are also known as "time series".

Assume $E[u(t)] = 0$. The covariance

$$R(T) = E[u(t+T)\overline{u(t)}],$$

where $T \in \mathbb{Z}$ satisfies, as before:

1. $R(0) \geq 0$.
2. $|R(T)| \leq R(0)$.
3. $R(-T) = \overline{R(T)}$.
4. $\sum_{i,j} R(i-j) z_i \overline{z_j} \geq 0$.

If $u(t,\omega) = \xi(\omega)h(t)$ is stationary we can repeat the arguments in Section 4.1. Since $R(0) = E[|u|^2] = E[|\xi|^2]|h(t)|^2$ we see that $h(t) = Ae^{i\phi(t)}$ for $t = 0, \pm 1, \ldots$. Since $R(1) = \overline{R(-1)}$ we obtain

$$\phi(t+1) - \phi(t) = -(\phi(t-1) - \phi(t)) \mod 2\pi$$

for $t = 0, \pm 1, \ldots$. Setting $\phi(0) = \alpha$ and $\phi(0) - \phi(-1) = \lambda$ we find by induction that $\phi(t) = \alpha + \lambda t \mod 2\pi$. Consequently $h(t) = Ae^{i(\alpha+\lambda t)} = Ce^{i\lambda t}$ for all integers t.

Define a periodic function f of the argument k by

$$f(k) = \frac{1}{2\pi} \sum_{T=-\infty}^{+\infty} R(T) e^{-iTk},$$

where T takes on integer values. Note that if $R(T)$ does not converge rapidly enough to 0 as $|T|$ increases, f may not be smooth. Then $R(T) = \int_{-\pi}^{\pi} e^{iTk} f(k) dk$. (The factor 2π of Fourier theory is broken up here differently than before.)

One can show that if $R(T)$ is a covariance for a time series, then $f \geq 0$. In particular, there exists a non-decreasing function $F(k)$ such that $R(T) = \int_{-\pi}^{\pi} e^{iTk} dF(k)$. Conversely, if $R(T)$ is given for all integers T, and if $\frac{1}{2\pi} \sum_T R(T) e^{-iTk} \geq 0$, then there exists a time series for which $R(T)$ is the covariance.

Consider the problem of finding a good estimate for $u(t+m,\omega)$ when we have values $u(t-n)$, $u(t-(n-1))$, $\dots$, $u(t-1)$. We would like to find a random variable $\hat{u}(t+m,\omega)$ with $m = 0, 1, 2, \dots$ such that

$$E\left[|u(t+m,\omega) - \hat{u}(t+m,\omega)|^2\right]$$

is as small as possible. We know from earlier work that

$$\hat{u}(t+m,\omega) = E[u(t+m,\omega)|u(t-1), u(t-2), \dots, u(t-n)].$$

The way to evaluate $\hat{u}$ is to find a basis $\{\phi_i\}$ in the space of functions of $\{u(t-n), \dots, u(t-1)\}$, expand $\hat{u}$ in this basis, i.e.,

$$\hat{u} = \sum_{j=1}^{n} a_j \phi_j(u(t-1), \dots, u(t-n)),$$

and calculate the coefficients a_j of the expansion. This is hard in general. We simplify the problem by looking only for the best approximation in the span of $\{u(t-1), \dots, u(t-n)\}$ i.e., we look for a random variable

$$\hat{u}(t+m,\omega) = \sum_{j=1}^{n} a_j u(t-j,\omega).$$

This is called linear prediction. The span L of the $u(t-j,\omega)$ is a closed linear space; therefore, the best linear prediction minimizes

$$E\left[|u(t+m,\omega) - \hat{u}(t+m,\omega)|^2\right]$$

for $\hat{u}$ in L. What we have to do is to find $\{a_j\}_{j=1}^{n}$, such that

$$E\left[\left|u(t+m,\omega) - \sum_{j=1}^{n} a_j u(t-j,\omega)\right|^2\right]$$

is as small as possible. We have

$$
\begin{aligned}
&E[|u-\hat{u}|^2] \\
&\quad = E\left[\left(u(t+m)-\sum_j a_j u(t-j)\right)\overline{\left(u(t+m)-\sum_l a_l u(t-l)\right)}\right] \\
&\quad = E\left[u(t+m)\overline{u(t+m)} - \sum_l \overline{a_l} u(t+m)\overline{u(t-l)}\right. \\
&\quad \left. - \sum_j a_j \overline{u(t+m)} u(t-j) + \sum_j \sum_l a_j \overline{a_l} u(t-j)\overline{u(t-l)}\right] \\
&\quad = R(0) - 2\mathrm{Re}\left(\sum_j \overline{a_j} R(m+j)\right) + \sum_j \sum_l a_j \overline{a_l} R(l-j),
\end{aligned}
$$

which is minimized when

$$
\frac{1}{2}\frac{\partial E\left[|u-\hat{u}|^2\right]}{\partial \overline{a_j}} = -R(m+j) + \sum_{l=1}^{n} a_l R(j-l) = 0 \tag{4.5}
$$

for $j = 1, \ldots, n$. Here we use the fact that if $f(x,y) = F(x+iy, x-iy) = F(z,\bar{z})$ is real, then $f_x = f_y = 0$ if and only if $F_{\bar{z}} = 0$ or $F_z = 0$ (see also exercise 5, chapter 1). The uniqueness of the solution of the system (4.5) and the fact that this procedure gives a minimum are guaranteed by the orthogonal projection theorem for closed linear spaces (see Section 1.1).

Rewrite (4.5) in terms of the Fourier transform. The spectral representation of $R(T)$ is

$$
R(T) = \int_{-\pi}^{\pi} e^{ikT} dF(k).
$$

Assume that $dF(k) = f(k)\, dk$. Then (4.5) becomes

$$
\int_{-\pi}^{\pi} \left(-e^{i(j+m)k} + \sum_{l=1}^{n} a_l e^{i(j-l)k}\right) f(k)\, dk = 0.
$$

Putting e^{ijk} outside the parentheses, we get

$$
\int_{-\pi}^{\pi} e^{ijk}\left(e^{imk} - \sum_{l=1}^{n} a_l e^{-ilk}\right) f(k)\, dk = 0. \tag{4.6}
$$

So far (4.6) is just a reformulation of (4.5). To continue we need an explicit representation of $f(k)$. Consider the example where $R(T) = Ca^{|T|}$ for $T = 0, \pm 1, \pm 2, \ldots$ where $C > 0$ and $0 < a < 1$. Is R the

covariance of a stationary process? It certainly satisfies conditions (1), (2), (3). To check (4) we compute

$$\begin{aligned} f(k) &= \frac{1}{2\pi}\sum_{n=-\infty}^{\infty} R(n)e^{-ink} \\ &= \frac{C}{2\pi}\left[\sum_{n=1}^{\infty}(ae^{-ik})^n + 1 + \sum_{n=1}^{\infty}(ae^{ik})^n\right] \\ &= \frac{C}{2\pi}\left[\frac{ae^{-ik}}{1-ae^{-ik}} + 1 + \frac{ae^{ik}}{1-ae^{ik}}\right] \\ &= \frac{C}{2\pi}\frac{1-a^2}{(1-ae^{-ik})(1-ae^{ik})} > 0. \end{aligned}$$

This shows that $R(T)$ is the Fourier transform of a non-negative function, and consequently the covariance of a stationary process.

Assume for simplicity that $C(1-a^2)/(2\pi) = 1$. We solve (4.6) using complex variables. Let $e^{ik} = z$. Then $\bar{z} = z^{-1}$, $dk = dz/(iz)$ and (4.6) becomes:

$$\frac{1}{2\pi}\int_{|z|=1} z^j\left(z^m - \sum_{\ell=1}^{n} a_\ell z^{-\ell}\right)\frac{1}{(z-a)\left(\frac{1}{z}-a\right)}\frac{dz}{iz} = 0$$

for $j = 1, 2, \ldots, n$. We must therefore determine $a_1 \ldots a_n$ such that

$$\sum_{\ell=1}^{n} a_\ell \frac{1}{2\pi i}\int_{|z|=1}\frac{z^{j-\ell}(1-az)^{-1}}{z-a}\,dz = \frac{1}{2\pi i}\int_{|z|=1}\frac{z^{j+m}(1-az)^{-1}}{z-a}\,dz.$$

We find the coefficients recursively by comparing two consecutive values of j, starting from the back. Let $j = n$ and $j = n-1$. Using residue theory we get

$$\sum_{\ell=1}^{n}\frac{a_\ell a^{n-\ell}}{1-a^2} = \frac{a^{n+m}}{1-a^2},$$

$$\sum_{\ell=1}^{n-1}\frac{a_\ell a^{n-1-\ell}}{1-a^2} + a_n\left[\frac{a^{-1}}{1-a^2} + \frac{(1-a\cdot 0)^{-1}}{0-a}\right] = \frac{a^{n-1+m}}{1-a^2}.$$

Multiplying the last equation by a and subtracting we get $a_n = 0$. This simplifies the next step with $j = n-1$ and $j = n-2$ substantially, and using similar arguments we obtain $a_{n-1} = 0$. In the last step

$$\frac{a_1}{2\pi i}\int_{|z|=1}\frac{z}{z}\frac{(1-az)^{-1}}{z-a}\,dz = \frac{1}{2\pi i}\int_{|z|=1}\frac{z^{1+m}(1-az)^{-1}}{z-a}\,dz$$

which yields: $a_1(1-a^2)^{-1} = a^{1+m}(1-a^2)^{-1}$ or $a_1 = a^{1+m}$. We have therefore shown that if $R(T) = Ca^{|T|}$ with $0 < a < 1$ then the best approximation of $u(t+m,\omega)$ for $m = 0, 1, \ldots$ is $a^{1+m}u(t-1,\omega)$.

4.6. Data Assimilation

We now turn to the topic of data assimilation, which could have been discussed at the end of chapter 3 but which has been set here so that the presentation can be read along with the related discussion of prediction for stationary processes in the previous section. As in the section on prediction, 4.5, we are given data and are trying to use them to make a prediction, but the setting is different.

There are many situations where one wants to draw conclusions on the basis of models that are not accurate enough, but that can be supplemented by current data. The canonical example is meteorology, where, at any one time, one has an incomplete description of the current weather, the equations of motion provide an incomplete description of the atmosphere, but data are coming in all the time. The use of data together with a model to assess the current state of a system and/or to make predictions is called "data assimilation," and the algorithms for doing that are called "filters." The availability of the model frees us from the restriction to stationary processes; if the data are accurate while the model is noisy we should be paying more attention to the data, and vice-versa if the model is less noisy. The question is how to best use the available information.

A useful model of a situation where data assimilation is needed consists of a stochastic differential equation

$$dx = f(x,t)\,dt + g(x,t)\,dw, \tag{4.7}$$

where $x = (x_1, x_2, \ldots, x_n)$ is an n-dimensional vector, dw is an n-dimensional BM, f is an n-dimensional vector function, and g is a scalar (i.e., an n by n diagonal matrix of the form gI, where g is a scalar and I is the identity matrix). The BM encapsulates all that is not known in this model. The initial state $x(0)$ is assumed given and may be random as well.

As the experiment unfolds, it is observed, and the values y_i of a "measurement process" are recorded at times $t_i = i\delta$, where δ is a fixed time interval; they are related to the evolving "state" $x(t)$ by

$$y_i = h(x_i, t_i) + GW_i, \tag{4.8}$$

where y_i is a k-dimensional vector, with in principle $k \le n$ (but in what follows we assume $k = n$), h is a nonlinear vector function, G is a diagonal matrix with nonzero diagonal terms, $x_i = x(i\delta)$, and W_i

is a vector whose components are independent Gaussian variables of mean 0 and variance 1, independent also of the BMs in the equation. Equation (4.8) says that the observations are noisy, with Gaussian noise. Now, the problem is to estimate x on the basis of (4.7) and the observations y_i. We are interested in cases where simply rerunning the problem with a different sample of BM will not do because the different samples differ more than we can tolerate; the observations should narrow down the range of possible $x(t)$. The solution of the problem is, of course, the process $\hat{x} = E[x(t)|\bar{y}(t)]$, where $\bar{y}(t)$ is the sequence $y_1, y_2, \dots$ for indices j such that $j\delta \le t$. This, as we know, is the best approximation of $x(t)$ given $\bar{y}$.

If the system (4.7) is linear and the data are Gaussian (or, as a special case, not random), then the solution of the problem is a Gaussian process. Its means and covariances can be calculated from those of the various functions that appear in the equation. This is the "Kalman filter" or "Kalman-Bucy filter," a mainstay of engineering. It provides a nonrandom solution of a random problem. This is not the place to present the algebra involved in deriving the full Kalman filter procedure, and we are content with a simple example as an illustration.

Suppose (4.7) is scalar and has the particularly simple form $dx = 0$ (i.e., the initial state $x(0)$, which we take as Gaussian with mean 0 and variance σ^2, does not change in time, $x(t) = x(0)$). Suppose the observation process is $y_i = x_i + gW_i$, with the W_i independent Gaussian variables of mean 0 and variance 1. The variance of y_1 is $\sigma^2 + g^2$, and the conditional expectation of x_1 given y_1 (i.e, the projection of x_1 on y_1), is $y_1\sigma^2/(\sigma^2 + g^2) = \hat{x}$, the filtered estimate. If the variance g^2 of the observation noise is large, the observation adds little to the accuracy of the simplest estimate unaided by observations in which the variable x_1 is estimated by its mean; on the other hand, if the variance of the observation noise is small, the observation is accurate and the estimate reduces to equating the estimate to the observation. Thus, the filter neatly blends in the information from the "equation" and the information from the observation, weighing their reliability as measured by the variances of the noises.

Now, consider the general nonlinear case of (4.7). We have to estimate the variables $x(t)$, and the natural thing to do is try to evaluate their probability density as it evolves in time. The initial state x is known and so is its probability density; all we have to do is evaluate sequentially the density of x_{i+1} assuming that we know the density of x_i.

Let P_i be the probability density of x at time $i\delta$ (taking into account the observations before that time and at that time). To find

the probability density of x at times $i\delta < t < (i+1)\delta$ (before any more observations come in), one can sample the density P_i, evolve the samples independently by (4.7), and, whenever needed, reconstruct a density by, for example, parametric estimation (see Section 2.4). The challenge is to modify the density at time $(i+1)\delta$ when new data must be taken into account. This can be done by Bayesian estimation (see Section 2.7).

Bayes' theorem says that

$$P(x|\bar{y}_{i+1}) = \frac{P(y_{i+1}|x(t))P(x|\bar{y}_i)}{\int P(y_{i+1}|x)P(x|\bar{y}_i)\,dx}, \tag{4.9}$$

where $P(x|\bar{y}_i)$ is the probability density determined from (4.7) taking into account the data up to and including time $i\delta$ but not the data at $(i+1)\delta$, $P(y_{i+1}|x(t))$ is the probability of finding the data if one knows the value $x((i+1)\delta)$ of the unknown vector x, and the integral in the denominator is what is needed to normalize the probabilities. The connection with Bayesian estimation is made by viewing $P(x|\bar{y}_i)$, the density in which the new data have not yet been taken into account, as the prior density, and then taking the density after the data have been used as the posterior density.

In words, (4.9) says that the new density given the new data is the product of the probability of getting the data if the values of the samples of the distribution were known, multiplied by the prior probability for samples of the distribution, the whole thing properly normalized. The probability of getting the data if the values of the samples were known can be obtained from the observation equation (4.8):

$$P(s_j \le x_j < s_j + ds_j) = \frac{1}{\sqrt{2\pi g_{jj}^2}} \exp\left(-\frac{(s_j - h_j(x,t))^2}{2g_{jj}^2}\right) ds_j, \tag{4.10}$$

where g_{jj} is a diagonal entry of the matrix G. Formula (4.9) can be evaluated as follows: We can find n samples of P_i and evolve them by (4.7). The density that one can, in principle, reconstruct from the positions of these samples after evolution does not take into account the new information at time $(i+1)\delta$, and we will use it as the prior density at the new time. The new information (i.e., the observation) at the time $(i+1)\delta$ makes it possible to assign a probability to each new sample position; if a sample is at the position x and the observation y_i is available, then its probability is given by (4.10).

Before taking the data into account, assign to each sample a weight, say the weight 1. Take the data into account by replacing these uniform weights by $Z^{-1}p_i^{\text{new}}$, where the p_i^{new} come from (4.10)—they take into account the fact that the new observations make some sample positions, far from the observation, unlikely, and those that are near the observation more likely. Choose Z so that the sum of the new weights is 1. The sample positions in x-space are unchanged. We now have samples whose positions have been determined by the prior density and that have weights that take the new observation into account. We can now estimate the new posterior density at time $(i+1)\delta$ from this information. One can estimate from these positions and weights the most likely state of the system given the observations.

One cannot use the same samples over and over, and at the beginning of each step one has to resample the new density P_{i+1} or else the algorithm goes haywire. This should be obvious: Some of the samples get very low weights after the new data have been taken into account; if the new densities are not resampled, then after a few steps one is dealing with samples all of which have weights near zero. Resampling is a key feature of many Monte Carlo schemes.

This is the Bayesian filter. The samples are often called "particles" and this filter is also known as a "particle filter."

4.7. Exercises

1. Find some way to show nonrigorously that the correlation function of white noise is a delta function. Suggestion: Approximate BM by a random walk with Gaussian increments of nonzero length, find the time series of the derivative of this walk, calculate the correlations, and take a (formal) limit.

2. Consider the stochastic process $u = \xi\cos(t)$, where ξ is a random variable with mean 0 and variance 1. Find the mean and the covariance functions. Obviously, this is not a stationary process. However, $\cos(t) = (e^{it} + e^{-it})/2$. How do you reconcile this with the construction we have of stationary processes as sums of exponentials?

3. Consider the differential equation $(u^2)_x = \epsilon u_{xx}$ on the real line, with the boundary conditions $u(-\infty) = u_0$, $u(+\infty) = -u_0$, where ϵ and u_0 are constants. Assume that u is a velocity, with units L/T, where L is a unit of length and T a unit of time. Find the units of ϵ. Because of the boundary conditions, u does not have a usual Fourier transform, but one can define one by taking the Fourier transform

of u' and dividing it by ik. Let $\hat{u}(k)$ be this Fourier transform of u. Define the energy spectrum by $E(k) = |\hat{u}(k)|^2$. Find the units of $E(k)$; show that the dimensionless quantity $E(k)k^2/u_0^2$ must be a function of the variable $k\epsilon/u_0$. Deduce that as you take the limit $\epsilon \to 0$, the spectrum converges to $E(k) \sim 1/k^2$.

4. Extend BM to $(-\infty, +\infty)$ by starting another BM from the origin backward in time. Consider the integral $\int_{-\infty}^{+\infty} \psi(s-t)\, dw(s)$, where w is the extended BM and ψ some nice smooth function that decays quickly to zero at infinity. Check that the integral makes sense and has the same value in the Stratonovich and Ito calculi. It defines a stochastic process $u(t)$. Calculate its mean and covariance. Is it stationary in the wide sense? The narrow sense?

5. Consider the wide sense stationary stochastic process $u = \xi e^{it}$, where ξ is a Gaussian variable with mean 0 and variance 1. What is its stochastic Fourier transform? What is the measure $\rho(dk)$?

6. Consider a stochastic process of the form $u(\omega, t) = \sum_j \xi_j e^{i\lambda_j t}$, where the sum is finite and the ξ_j are independent random variables with means 0 and variances v_j. Calculate the limit as $T \to \infty$ of the random variable $(1/T)\int_{-T}^{T} |u(\omega, s)|^2\, ds$. How is it related to the spectrum as we have defined it? What is the limit of $(1/T)\int_{-T}^{T} u\, ds$?

7. Suppose you have to construct on the computer (for example, for the purpose of modeling the random transport of pollutants) a Gaussian stationary stochastic process with mean 0 and a given covariance function $R(t_2 - t_1)$. Propose a construction.

8. Show that there is no stationary (wide sense) stochastic process $u = u(\omega, t)$ that satisfies (for each ω) the differential equation $y'' + 4y = 0$ as well as the initial condition $y(t = 0) = 1$.

9. Let η be a random variable. Its characteristic function is defined as $\phi(\lambda) = E[e^{i\lambda\eta}]$. Show that $\phi(0) = 1$ and that $|\phi(\lambda)| \leq 1$ for all λ. Show that if $\phi_1, \phi_2, \ldots, \phi_n$ are the characteristic functions of independent random variables $\eta_1, \ldots, \eta_n$, then the characteristic function of the sum of these variables is the product of the ϕ_i.

10. Show that if $\phi(\lambda)$ is the characteristic function of η, then

$$E[\eta^n] = (-i)^n \frac{d^n}{d\lambda^n}\phi(0)$$

provided both sides of the equation make sense. Use this fact to show that if ξ_i, $i = 1, ..., n$, are Gaussian variables with means 0,

not necessarily independent, then

$$E[\xi_1\xi_2\cdots\xi_n] = \begin{cases} \Sigma\Pi E[\xi_{i_k}\xi_{j_k}], & n \text{ even} \\ 0, & n \text{ odd.} \end{cases}$$

On the right-hand side, i_k and j_k are two of the indices, the product is over a partition of the n indices into disjoint groups of two, and the sum is over all such partitions (this is "Wick's theorem"). Hints: Consider the variable $\Sigma\lambda_j\xi_j$; its moments can be calculated from the derivatives of its characteristic function. By assumption, this variable is Gaussian and its characteristic function, i.e., the Fourier transform of its density, is given by a formula we have derived.

11. Consider the random differential equation

$$\frac{d}{dt}q(t) = -ibq(t), \quad q(0) = 1,$$

where b is a random variable and $i = \sqrt{-1}$. Define $Q(t) = E[q(t)]$. Show that $|Q(t)| \le Q(0)$. Suppose the distribution of b is not known but you know the moments of b; that is, you know $E[b^j]$ for all $j \le N$. Solve the equation by iteration: $q_0 = 1$ and $q_{j+1} = 1 - ib\int_0^t q_j(s)\,ds$ for $j \ge 0$, and then define $Q_j = E[q_j]$ for $j \le N$, thus using the information you have. Show that however large N may be, as long as it is finite, the approximation Q_j will violate the inequality $|Q(t)| \le Q(0)$ for t large and any $j > 1$.

12. Continue the example of data assimilation: Suppose $x(t) = x(0)$ is a scalar and you have observations $y_i = x_i + gW_i$, where g is a fixed constant. What is $\hat{x}_i = E[x_i|\bar{y}]$ for $i > 1$?

13. Consider the following functions $R(T)$; which ones are the covariance functions of some stationary stochastic process, and why? ($T = t_2 - t_1$ as usual).

    1. $R(T) = e^{-T^2}$.
    2. $R = Te^{-T^2}$.
    3. $R = e^{-T^2/2}(T^2 - 1)$.
    4. $R = e^{-T^2/2}(1 - T^2)$.

4.8. Bibliography

[1] G.I. Barenblatt, *Scaling*, Cambridge University Press, Cambridge, 2004.

[2] G.I. Barenblatt and A.J. Chorin, A mathematical model for the scaling of turbulence, Proc. Natl. Acad. Sci. USA, 101 (2004), pp. 15,023-15,026.

[3] A.J. Chorin, *Vorticity and Turbulence*, Springer-Verlag, New York, 1994.

[4] A.J. Chorin and P. Krause, Dimensional reduction for a Bayesian filter, Proc. Natl. Acad. Sci. USA, 101 (2004), pp. 15,013-15,017.

[5] C.K. Chui and G. Chen, *Kalman Filtering*, Springer-Verlag, Berlin, 1987.

[6] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte-Carlo Methods in Practice*, Springer-Verlag, New York, 2001.

[7] B. Efron, *The Jacknife, the Bootstrap, and Other Resampling Plans*, CBMS-NSF Regional Conference Series, SIAM, Philadelphia, 1982.

[8] M. Ghil and P. Melanotte-Rizzoli, Data assimilation in meteorology and oceanography, Adv. Geophys., 53 (1991), pp. 141-256.

[9] I. Gikhman and A. Skorokhod, *Introduction to the Theory of Random Processes*, Saunders, Philadelphia, 1965.

[10] P. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, Berlin, 1992.

[11] R. Kraichnan, Dynamics of nonlinear stochastic systems, J. Math. Phys. 2 (1961), pp. 124-148.

[12] R. Miller, *Introduction to the Kalman Filter*, Seminars on Data Assimilation, European Center for Medium-Range Weather Forecasts, Reading, UK (1997), pp. 47-60.

[13] A. Yaglom, *An Introduction to the Theory of Stationary Random Functions*, Dover, New York, 1962.

CHAPTER 5

Statistical Mechanics

5.1. Mechanics

We begin the discussion of statistical mechanics by a quick review of standard mechanics. Suppose we are given N particles whose position coordinates are given by a set of scalar quantities $q_1, \ldots, q_n$. In a d-dimensional space, one needs d numbers to specify a location, so that $n = Nd$. The rate of change of the position is

$$\frac{d}{dt} q_i = \dot{q}_i.$$

(This dot notation for the time derivative goes back to Newton and makes some of the formulas below look less cluttered.) A good way to write down the laws of motion is to specify a Lagrangian $\mathcal{L} = \mathcal{L}(q_i, \dot{q}_i, t)$ and follow the steps that will now be described; this procedure can be used for laws other than those of Newtonian mechanics as well. For any path $q(s)$, $t_0 \le s \le t$, that could take the particles from their locations at time t_0 to their locations at time t, we define an "action" by

$$A = \int_{t_0}^{t} \mathcal{L}(q(s), \dot{q}(s), s)\, ds,$$

and we require that the motion (according to the mechanics embodied in the Lagrangian) that takes us from $q(t_0)$ to $q(t)$ be along a path that is an extremal of the action. In other words, for the motion described by the functions $q(t)$ to obey the physics in the Lagrangian, it has to be such that perturbing it a little, say from $q(t)$ to $q(t) + \delta q(t)$, changes the action $A = \int_{t_0}^{t} \mathcal{L}\, ds$ very little. We simplify the analysis here by assuming that $\mathcal{L}$ does not explicitly depend on t. Then

$$\begin{aligned} \delta A &= \delta \int_{t_0}^{t} \mathcal{L}(q, \dot{q})\, ds = \int_{t_0}^{t} (\mathcal{L}(q + \delta q, \dot{q} + \delta \dot{q}) - \mathcal{L}(q, \dot{q}))\, ds \\ &= 0 + O(\delta q^2, \delta \dot{q}^2), \end{aligned}$$

where

$$\mathcal{L}(q+\delta q, \dot{q}+\delta\dot{q}) = \mathcal{L}(q_i, \dot{q}_i) + \sum \delta q_i \frac{\partial \mathcal{L}}{\partial q_i} + \sum \delta\dot{q}_i \frac{\partial \mathcal{L}}{\partial \dot{q}_i} + O(\delta q^2, \delta\dot{q}^2).$$

By integration by parts, we find

$$\begin{aligned}\delta \int_{t_0}^{t} \mathcal{L}\, ds &= \int_{t_0}^{t} \left(\sum \delta q_i \frac{\partial \mathcal{L}}{\partial q_i} + \sum \delta\dot{q}_i \frac{\partial \mathcal{L}}{\partial \dot{q}_i} + O(\delta q^2, \delta\dot{q}^2) \right) ds \\ &= \int_{t_0}^{t} \left(\sum \delta q_i \left(\frac{\partial \mathcal{L}}{\partial q_i} - \frac{d}{dt}\frac{\partial \mathcal{L}}{\partial \dot{q}_i} \right) + O(\delta q^2, \delta\dot{q}^2) \right) ds.\end{aligned}$$

For the path $q(t)$ to be extremal, the first term has to vanish, and we conclude that

$$\frac{\partial \mathcal{L}}{\partial q_i} - \frac{d}{dt}\frac{\partial \mathcal{L}}{\partial \dot{q}_i} = 0,$$

for all $i = 1, \ldots, n$. These are the Lagrange equations of motion.

EXAMPLE. Change notation so that $x = q$, $\dot{x} = \dot{q}$, and think of x as a coordinate in a one-dimensional space. Assume that a particle of mass m at x is acted on by a force F of the form $F = -\text{grad}\, V$, where $V = V(x)$ is a potential. Specify the laws of motion by setting $\mathcal{L} = \frac{1}{2}m\dot{x}^2 - V(x)$. The Lagrange equation of motion is

$$\frac{\partial \mathcal{L}}{\partial x} - \frac{d}{dt}\frac{\partial \mathcal{L}}{\partial \dot{x}} = 0$$

or, equivalently,

$$-\frac{\partial V}{\partial x} - \frac{d}{dt}(m\dot{x}) = 0,$$

which is Newton's second law, $F = m\ddot{x}$.

This formalism is also useful in quantum mechanics, where, in the notations of Section 3.6, the probability density of going from $q(t_0)$ to $q(t)$ is the square of the path integral

$$v(x,t) = \frac{1}{Z} \int e^{-\frac{i}{h}\int_0^t \left[\frac{1}{2}\left(\frac{dw}{ds}\right)^2 - U(x+w(s))\right]\, ds} \phi(x + w(t))[dw].$$

where the integration is over all paths that lead from $q(t_0)$ to $q(t)$; this expression is analogous to equation (3.30) of Section 3.6, except for the additional factor i/h in front of the integral, where i is $\sqrt{(-1)}$ and h is Planck's constant. One can see the action appear in the exponent. On scales where h cannot be viewed as very small, this is an oscillatory integral which produces wave-like motion; on scales where the h can be viewed as very small, the main contribution to this integral comes

from trajectories for which the exponent is stationary, leading back to the action formulation above.

We shall use the equations of motion mostly in their Hamiltonian form: Define a momentum p_i conjugate to q_i by $p_i = \partial\mathcal{L}/\partial\dot{q}_i$. The Hamiltonian function is

$$H = \sum p_i \dot{q}_i - \mathcal{L}.$$

By differentiating H with respect to $\dot{q}_i$ one sees that H is not a function of $\dot{q}_i$, and therefore it is a function of only the q_i, p_i; the equations of motion can be written as

$$\dot{q}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q_i}. \tag{5.1}$$

The proof that these equations are equivalent to the Lagrangian equations, under appropriate smoothness conditions, is just a manipulation of differentials, which we leave to the reader.

Example. Let $\mathcal{L} = \frac{1}{2}m\dot{x}^2 - V(x)$ as before, with $q = x$. Then $p = m\dot{x}$ and

$$H = p\dot{q} - \mathcal{L} = (m\dot{x})\dot{x} - \left(\frac{1}{2}m\dot{x}^2 - V(x)\right) = \frac{1}{2}\frac{(m\dot{x})^2}{m} + V.$$

The Hamiltonian equations of motion are

$$\dot{x} = \frac{\partial H}{\partial p} = \frac{p}{m}$$

and

$$\dot{p} = m\frac{d^2x}{dt^2} = -\frac{\partial H}{\partial q} = -\frac{\partial V}{\partial x} = F.$$

If the Hamiltonian does not depend explicitly on time, then it is a constant during the motion; indeed,

$$\begin{aligned}\frac{dH}{dt} &= \sum_{i=1}^{n} \frac{\partial H}{\partial p_i}\frac{dp_i}{dt} + \sum_{i=1}^{n} \frac{\partial H}{\partial q_i}\frac{dq_i}{dt} \\ &= \sum_{i=1}^{n} \frac{\partial H}{\partial p_i}\left(-\frac{\partial H}{\partial q_i}\right) + \frac{\partial H}{\partial q_i}\frac{\partial H}{\partial p_i} \\ &= 0.\end{aligned}$$

The constant value of the Hamiltonian is the energy E of the system. A system of equations that can be put into the form (5.1) is a Hamiltonian system.

As an illustration, consider the harmonic oscillator, which is a particle of mass m that can move on the line, with a rubber band anchoring

it to the origin. The force on the particle is $F = -Kx$, where K measures the elasticity of the band and x is the position of the particle. The momentum of the particle is $p = m\dot{x}$, and the equation of motion is $\dot{p} = -Kx$. These equations are reproduced if one sets $H = \frac{1}{2m}p^2 + \frac{1}{2}Kq^2$, where the variable x has been renamed q to conform with the general notation above. (Quite often the energy, i.e., the value of the function H, is the sum of a contribution that is quadratic in the momenta p (the "kinetic energy") and a second contribution that is a function of the positions q (the "potential energy"). The Lagrangian $\mathcal{L}$ is then the kinetic energy minus the potential energy.)

The equations of motion for the harmonic oscillation can be solved explicitly. Set $\omega = \sqrt{K/m}$ (not to be confused with a point in probability space); the solution is $q(t) = A\cos(\omega t) + B\sin(\omega t)$; $p(t) = -Am\omega\sin(\omega t) + Bm\omega\cos(\omega t)$, where the coefficients A, B are determined by the initial values of q, p. With a suitable change of units one make K, m have the numerical values $1, 1$, and then $\omega = 1$ and $H = q^2/2 + p^2/2$.

5.2. Statistical Mechanics

Consider a Hamiltonian system with n degrees of freedom (q_1, p_1), $\ldots$,(q_n, p_n), where H does not depend explicitly on the time t. From now on, we will denote the vector of positions by q and the vector of momenta by p so that $H = H(q, p)$. A microscopic state of the system (a "microstate" for short) is a set of values of the $q_1, \ldots q_n, p_1, \ldots, p_n$. The system evolves in a $2n$-dimensional space, which is denoted by Γ and is often called the phase space. The sequence of points in Γ (microstates) that the system visits as it evolves from an initial condition is called a trajectory.

If the system has many degrees of freedom, it is impossible to follow its exact evolution in time, since specification of all the initial conditions is impossible and the numerical solution of the very large systems that arise in practice is also out of reach. So we settle for a more modest approach. We assume that the initial data $q(0)$ and $p(0)$ are drawn from a probability density W. Then, instead of considering single trajectories, we look at the collection, or "ensemble," of trajectories that are initially distributed according to W.

As the trajectories evolve individually, the probability density naturally changes; let the density of microstates at time t be $W(t)$, where each microstate is the location of a trajectory at that time. $W(t)$ describes the ensemble at time t; it is the "macrostate" of the ensemble. Thus, the microstate is a list of numbers, or a vector in Γ, and the

macrostate is a probability density in Γ. The set of all macrostates corresponds to Ω, the sample space of our earlier discussion.

We now derive an equation of motion for $W(t) = W(q, p, t)$. Consider the vector $u = (\dot{q}_1, \dots, \dot{p}_n)$. First, note that its divergence is zero:

$$\begin{aligned}\operatorname{div} u &= \sum_{i=1}^{n} \frac{\partial}{\partial q_i}\left(\frac{dq_i}{dt}\right) + \sum_{i=1}^{n} \frac{\partial}{\partial p_i}\left(\frac{dp_i}{dt}\right) \\ &= \sum_{i=1}^{n} \frac{\partial}{\partial q_i}\left(\frac{\partial H}{\partial p_i}\right) + \sum_{i=1}^{n} \frac{\partial}{\partial p_i}\left(-\frac{\partial H}{\partial q_i}\right) \\ &= 0.\end{aligned}$$

This vector field can be said to be "incompressible," in analogy with fluid dynamics.

Consider a volume V in Γ-space and a density of systems W. The number of microstates in V at a given time t is, on average, $\int_V W\, dV$, where dV is the element of volume in Γ; when the position variables q are cartesian coordinates $dV = dqdp$ (where $dq = dq_1 \cdots dq_n$ and similarly for dp). If microstates neither appear nor disappear, then the only change in the number of systems in V can come from the inflow/outflow of systems across the boundary of V. Therefore, as in fluid mechanics,

$$\frac{d}{dt}\int_V W dq\, dp = -\int_{\partial V} Wu \cdot n\, dS = -\int_V \operatorname{div}(Wu)\, dV,$$

where n is the outer normal to the boundary ∂V of V. If we assume that the density is smooth, we can deduce from the above that

$$\frac{\partial W}{\partial t} + \operatorname{div}(Wu) = 0, \tag{5.2}$$

and, using the incompressibility of u,

$$\frac{\partial W}{\partial t} + u \cdot \operatorname{grad} W = 0. \tag{5.3}$$

This last equation is known as the Liouville equation. One can define a linear differential operator (the Liouville operator)

$$L = \sum_{i=1}^{n} \frac{\partial H}{\partial p_i}\frac{\partial}{\partial q_i} - \frac{\partial H}{\partial q_i}\frac{\partial}{\partial p_i}$$

and then (5.3) becomes

$$\frac{\partial W}{\partial t} = -LW. \tag{5.4}$$

This equation is linear even when the original system is not. Inasmuch as it is an equation for the evolution of a probability density function, it is analogous to the Fokker-Planck equation; this analogy will be pursued in the next chapter.

Once we have the density $W(t)$, we can define physical observables for the ensemble, which are averages of physical quantities over the ensemble. The energy of each microstate is the value of the Hamiltonian H for that microstate; the energy of the ensemble is

$$E(t) = E[H(t)] = \int_\Gamma H(q,p)W(q,p,t)\,dV,$$

where dV is an element of volume in the phase space Γ. Similarly, if $\Phi = \Phi(q,p)$ is a property of a microstate,

$$\bar{\Phi} = E[\Phi] = \int_\Gamma \Phi(q,p)W(q,p,t)\,dV.$$

A probability density W is invariant in time if it is a stationary solution of (5.2); that is, if we draw the initial data from W, solve the equations for each initial datum, and look at the density of solutions at some later time t, it is still the same W. In other words, sampling the density and evolving the microstates commute. We now give two examples of time invariant densities for a Hamiltonian system.

Suppose that initially W is zero outside a region V and suppose that the system has no way of leaving V. Further suppose that W is constant inside V. Then from (5.3), we conclude that W is invariant. We apply this in the following construction. Consider in Γ-space a surface $H = E_0$ as well as the surface $H = E_0 + \Delta E_0$, where $E_0, \Delta E_0$ are a constants. The volume enclosed between these two surfaces is called an energy shell. Consider the following initial density:

$$W(q,p) = \begin{cases} (\text{volume of shell})^{-1}, & (q,p) \in \text{shell} \\ 0, & \text{otherwise.} \end{cases}$$

Since no systems can leave the energy shell (because the energy is a constant of the motion), this density is invariant in time. If we let the thickness ΔE_0 of the energy shell go to zero, we get a "microcanonical" density. The resulting surface density on the energy surface $H = E_0$ need not be constant.

Suppose $\phi(H)$ is a function of H such that $\int_\Gamma \phi(H)\,dq\,dp = 1$ and $\phi(H) \geq 0$. Then $W(q,p) = \phi(H)$ is invariant in time. Note first that

$u \cdot \operatorname{grad} W$ vanishes. Indeed,

$$\begin{aligned} u \cdot \operatorname{grad} W &= \sum_{i=1}^{n} \frac{dq_i}{dt}\frac{\partial W}{\partial q_i} + \sum_{i=1}^{n} \frac{dp_i}{dt}\frac{\partial W}{\partial p_i} \\ &= \frac{\partial \phi}{\partial H}\left(\sum_{i=1}^{n} \frac{dq_i}{dt}\frac{\partial H}{\partial q_i} + \sum_{i=1}^{n} \frac{dp_i}{dt}\frac{\partial H}{\partial p_i}\right) \\ &= 0. \end{aligned}$$

Therefore, from (5.3), $\partial W/\partial t = 0$. In particular, one can choose as an invariant density $W(q,p) = Z^{-1}\exp(-\beta H(q,p))$, where $\beta > 0$ is a constant and $Z = \int_\Gamma \exp(-\beta H)\, dq\, dp$. A density of this form is called canonical.

A property of the Liouville operator that will be used later is the following: Let $E[\cdot]$ is the expectation with respect to a canonical density; we have seen that if u, v are two functions defined on the relevant probability space, then $E[uv]$ defines an inner product, $(u,v) = E[uv]$, and then

$$(Lu, v) = E[(Lu)v] = -E[u(Lv)] = -(u, Lv)$$

(i.e., L is skew-symmetric). This can be checked by writing down the definitions and integrating by parts.

5.3. Entropy and Equilibrium

Consider a probability space where Ω consists of a finite number of points $\omega_1, \omega_2, \ldots, \omega_n$ with probabilities $P_1, P_2, \ldots, P_n$ (whose sum must be 1) (note that P, a probability, is not related to the p of the preceding sections). We now want to define a quantity called "entropy" on that space, to be denoted by S. S will be a function of the P_i: $S = S(P_1, \ldots, P_n)$ and we will consider the case where n may vary. We want S to be a measure of the uncertainty in the probability density and, to that end, satisfy the following axioms:

1. For each n, S is a continuous function of all its arguments.
2. If all of the P_i are equal ($P_i = 1/n$ for all i), one can define $S_n = S(1/n, \ldots, 1/n)$ and require that S_n be a monotonically increasing function of n (the more points in Ω, the more uncertainty if all points are equally likely).
3. Let $1 \le k_1 < k_2 < k_2 < \cdots < k_M = n$ be a partition of $[1,n]$ and let $Q_j = P_{k_{j-1}+1} + \cdots + P_{k_j}$ (i.e., $Q_1 = P_1 + \cdots + P_{k_1}$ with

$k_0 = 0$, $Q_2 = P_{k_1+1} + \cdots + P_{k_2}$, etc). Then

$$S(P_1, \ldots, P_n) = S(Q_1, \ldots, Q_M) + \sum_{j=1}^{M} Q_j S\left(\frac{P_{k_{j-1}+1}}{Q_j}, \ldots, \frac{P_{k_j}}{Q_j}\right).$$

In other words, the uncertainty is the sum of the uncertainties inherent in any grouping of points plus the average of the uncertainties within each grouping.

A function S with these properties should be small if all the probability is concentrated at a few points and should become ever larger as there is more doubt as to where an arbitrary point would lie. One can prove that a function S that satisfies these requirements is determined uniquely up to a multiplicative constant and is

$$S = -\sum_i P_i \log P_i.$$

This is the entropy associated with the probability space we started from. In physics, one multiplies this expression for S by the constant k (Boltzmann's constant). The entropy associated with a pdf (probability density function) f is, similarly, $S = -\int f(x) \log f(x)\, dx$. The entropy is a number attached to the pdf that measures, in the way described above, the uncertainty implicit in the pdf. If $S = 0$ and one makes the experiment that defines the density f, one knows in advance what the result will be; the larger S, the less one knows in advance.

Now consider a set of microstates (or, equivalently, the sample space for an evolving statistical mechanics system), with some reasonable σ-algebra of events. Suppose we have measured some physical, macroscopic, quantities, say $\bar{\Phi}_1, \bar{\Phi}_2, \ldots, \bar{\Phi}_m$, for some finite m. These are averages with respect to a density W of a set of microscopic (i.e., relating to each microstate) quantities Φ_i. We now ask the question: What pdf W compatible with these measurements (i.e., such that $\bar{\Phi}_i = \int \Phi_i(q,p) W(q,p)\, dV$) has maximum entropy? We now show the following: If there exists a vector $\beta = (\beta_1, \ldots, \beta_n)$ and a number $Z > 0$ such that

$$W_\beta = Z^{-1} \exp\Big(-\sum \beta_i \Phi_i(q,p)\Big)$$

is a probability density compatible with the measurements ("admissible" for short), then W_β is the admissible density that has the largest entropy among all admissible densities.

We now prove that W_β maximizes the entropy. It is an exercise in calculus to show that $\psi(x) = x \log x - x + 1 \geq 0$ for $x \geq 0$, with

equality only for $x = 1$. Put $x = W/W_\beta$ in this inequality, where W is an arbitrary admissible density. Then

$$-W \log W + W \log W_\beta \leq W_\beta - W.$$

Integrate this inequality over Γ and use the fact that both W and W_β are densities; this gives

$$-\int_\Gamma W \log W \, dV \leq -\int_\Gamma W \log W_\beta \, dV.$$

However, from the definition of W_β, we find that $-\log W_\beta = \log Z + \sum \beta_i \Phi_i$, and since both W and W_β are compatible with the measurements $\bar{\Phi}_i$, it follows that

$$-\int_\Gamma W \log W_\beta \, dV = \log Z + \sum \beta_i \bar{\Phi}_i = -\int_\Gamma W_\beta \log W_\beta \, dV,$$

because the integral of any density is 1; therefore the entropies of all the W's are less than the entropy of W_β:

$$S(W) \leq S(W_\beta),$$

where $S(W)$ is the entropy associated with a density W. Furthermore, the inequality is strict unless $W = W_\beta$.

As an example, suppose one has a single measurement, that of E, the energy of the ensemble, $E = E[H]$; then $W_\beta = Z^{-1} e^{-\beta H}$, where the β in the exponent is a scalar, and $Z = \int_\Gamma e^{-\beta H} \, dV$. The parameter β is determined from the equation

$$E = E[H] = \int_\Gamma Z^{-1} H e^{-\beta H} \, dV = -\frac{\partial}{\partial \beta} \log Z.$$

With this density, the entropy is $S = \beta E + \log Z$. A similar calculation, which we omit, produces the microcanonical density in the absence of any measurements.

It is a physical principle that the entropy of a physical system always increases, so it is reasonable to assume that any density for a physical system will evolve in time into one that maximizes the entropy. We already know that a canonical density is time invariant, so the canonical density is a good candidate for an asymptotic, invariant density, which is called in physics a "thermal equilibrium." This is particularly satisfying from the point of view of statistics as well: One can show that estimates based on partial measurements are unbiased if one assumes that the density that gives rise to them maximizes the entropy.

The temperature T of a system is defined by the equation

$$T^{-1} = \frac{\partial S}{\partial E};$$

one can check that if the density is the canonical density above, then $T = 1/\beta$ (in physics, there is an additional factor of k from the physicists' definition of entropy). Then the canonical density can be written as $W = Z^{-1}\exp(-H/T)$. For a system of N non-interacting particles, T/m can be seen to be the variance of the velocity of each particle (m is the mass of each particle). The canonical density has T as a fixed parameter and is the right density to use when the system under study allows no exchange of mass through its walls and has walls kept at a fixed temperature T. For the sake of simplicity, in this volume we shall always place ourselves in this case.

One can now proceed to derive all of thermodynamics from our definitions but we forbear to do so. We merely pause to note that the normalization constant Z varies when T varies, and is known in physics as the "partition function."

We now perform some useful calculations for a system of non-interacting particles. Consider N particles of mass m in a cube of side L (and volume $V = L^3$). Make the system periodic in space, so that if there is a particle at the point $x_1, x_2, x_3, 0 \leq x_i < L$, there are particles with the same mass and momenta at the points $x_i + k_i L$ for any integers k_i (and we use the letter x rather than q to denote location). If a particle leaves the box another particle enters from the opposite side. The Hamiltonian is $H = \frac{1}{2m}\sum_1^{3N} p_i^2$, where the momenta p have been relabeled consecutively regardless of the particle to which they belong. The partition function Z is $Z = \int\int\cdots\int dx_1 dx_2 \cdots dx_{3N} \int\cdots\int dp_1 \cdots dp_{3N} e^{-\beta H}$; the x integrations are trivial and yield V^N; the p integrals can be factored into a product of the $3N$ integrals $\int dp e^{-\beta p^2/2m} = \sqrt{2\pi m/\beta}$, so that $Z = V^N(2\pi m/\beta)^{3N/2}$ and $E = E[H] = -\frac{\partial}{\partial\beta}\log(Z) = \frac{3N}{2}T$. In a system of non-interacting particles, the energy is the number of degrees of freedom (i.e., the number of parameters required to specify the spatial configuration) times $T/2$ (in physics conventions, the Boltzmann constant k appears as an added factor). This is the "equipartition theorem".

Consider next the pdf f_H of the energy H. H is a constant on each trajectory but it is a random variable because the initial values are random; we have just calculated its mean. To see what is going on, consider three independent Gaussian variables ξ_1, ξ_2, ξ_3, each of mean 0 and variance 1; what is the pdf of $h = \frac{1}{2}(\xi_1^2 + \xi_2^2 + \xi_3^2)$? The probability

$f(\epsilon)d\epsilon$ that $\epsilon < h \leq \epsilon + d\epsilon$ is the probability that the point (ξ_1, ξ_2, ξ_3) lie in the shell between the sphere of radius $\sqrt{2\epsilon}$ and the sphere of radius $\sqrt{2(\epsilon + d\epsilon)}$, and is $(2/\sqrt{\pi})\,\sqrt{\epsilon}\,e^{-\epsilon}d\epsilon$. The geometric factor $\sqrt{\epsilon}$ increases because the annular becomes larger, while the exponential factor decreases.

Similarly, let x_i and p_i be as in the calculation of $E = E[H]$ above. The pdf f_H of $H = \frac{1}{2m}\sum_1^{3N} p_i^2$ is $C(N,\beta)e^{-\beta H}H^{3N/2-1}V^N/Z$, where $C(N,\beta)$ is independent of H, and can be determined from the fact that the integral of the density is 1.

If one plots $f_H/E[H]$ as a function of H for various values of N one finds that the graphs become increasingly concentrated around $H = E$, so that for large N the microcanonical and the canonical densities become indistinguishable. This remark is known as the "equivalence of ensembles".

If one plots the trajectory of a system in Γ-space one often notices that it wanders all over the surface of $H = E$ in a seemingly random fashion, and it is plausible that the average of a smooth function of the q, p on that surface with respect to the microcanonical density equals the average of the function along any trajectory (with a similar statement for the canonical density as a result of the equivalence of ensembles). If such a statement can be proved (it very occasionally can) it is called an "ergodic theorem", otherwise it is an "ergodicity assumption"; either one greatly simplifies computations. An example of an ergodic system is the system where Γ is the interval $[0, 1)$, and the equation of motion is $x_n = (x_{n-1} + \gamma) \bmod 1$, with x_0 given. One can readily check that if γ is irrational, then the average of any continuous function F on $[0, 1)$ equals its average over the trajectory that starts at any x_0.

5.4. The Ising Model

We now introduce the Ising model in two space dimensions, which is widely used as a model problem in statistical mechanics. Consider an $N \times N$ regular lattice in the plane with lattice spacing 1, and at each node (i, j), set a variable $s_{i,j}$ (a "spin") that can take only one of two values: $s_{i,j} = 1$ ("spin up") or $s_{i,j} = -1$ ("spin down"). Make the problem periodic, so that $s_{i+N,j} = s_{i,j}$ and $s_{i,j+N} = s_{i,j}$. Associate with this problem the Hamiltonian

$$H = -\frac{1}{2}\sum s_{i,j}(s_{i+1,j} + s_{i-1,j} + s_{i,j+1} + s_{i,j-1})$$

(i.e., minus the sum of the products of each spin with its four nearest neighbors). This "Hamiltonian" does not include any momenta, and

the variables take integer values only, so there is no time evolution, and "thermal equilibrium" here is meaningful only in the sense that the probability density we use maximizes the entropy. This is a reasonable generalization of the previous discussion, because once equilibrium has been reached the dynamics are no longer important.

The microstates of the system are the 2^{N^2} ways of arranging the up and down spins. We assign to each microstate the probability $Z^{-1}\exp(-H/T)$, where, as above, T is the temperature and Z is a normalization factor. A function of the microstates that is of interest is the "magnetization"

$$\mu = \frac{1}{N^2}\sum_{i,j} s_{i,j}.$$

Clearly, if all the spins are aligned, $\mu = +1$ or $\mu = -1$. With the above definitions, $E[\mu] = 0$ because a microstate with a given set of values for the spins and a microstate with exactly the opposite values have equal probabilities.

The covariance function is

$$\mathrm{Cov}(i', j') = E\left[(s_{i,j} - E[\mu])(s_{i+i',j+j'} - E[\mu])\right],$$

where the expected value of μ has been taken into account in preparation for the possibility, soon to be discussed, that it may be nonzero. The correlation length is a number ξ such that for $\sqrt{i'^2 + j'^2} > \xi$, the covariance is not significant (and we do not explain further how large "significant" is).

One can show, and check numerically as explained below, that the Ising model has the following properties:

1. For T very large or very small, ξ is small, of the order of 1. There is an intermediate value T_c of T for which ξ is very large.
2. The behavior of the magnetization μ is very different when $T < T_c$ and when $T > T_c$. In the former case, the likely values of μ hover around two nonzero values $\pm\mu_*$; if one adds dynamics to this problem (as we shall do with Monte Carlo sampling in the next section), one sees that the system is very unlikely to move from $+\mu_*$ to $-\mu_*$ or vice versa. For very large values of N, the phase space Γ separates into two mutually inaccessible regions that correspond to μ positive and μ negative. The averages of μ over each region then have one sign. On the other hand, when $T > T_c$, this separation does not occur. The value $T = T_c$ is a "critical value" of T and the parameter $E[\mu]$ is an "order parameter" that can be used to detect the partial order in which spins are aligned in each of the two mutually inaccessible regions

of Γ. As T passes from above this value T_c to below the critical value T_c, one has a "phase transition" in which the system goes from a disordered "phase" to a partially ordered phase. If one averages μ for $T < T_c$ only over the appropriate part of the phase space, one finds that, when $|T - T_c|$ is small, $E[\mu]$ is proportional to $|T_c - T|^\alpha$, where $\alpha = 1/6$ is an instance of a "critical exponent."

Some explanation for the splitting of the phase space into two mutually exclusive parts in two (and higher) dimensions is provided by a comparison of the Ising models in one space dimension and in two space dimensions. A one dimensional Ising models is just the obvious one-dimensional analog of what we have just discussed—a periodic chain of spins indexed by a single integer variable with a Hamiltonian involving near-neighbor interactions. In either dimension the microstates where all the spins point in the same directions are minima of the energy. Suppose β is large (T is small); suppose you are in one dimension and all the spins point up; how much energy do you have to invest to flip all the spins from up to down? Clearly, to flip one spin you must add to the system the energy 2β; once you have flipped one of the spins, you can flip its neighbors one after the other without a cost, until there is only one pointing up; then you flip that last holdout and recover your energy investment. The conclusion is that in one dimension these energy minima are not very deep.

By contrast, to flip all the spins in two dimensions on an N by N lattice you have to invest at least $2N\beta$ units of energy; thus the energy minima in two dimension are deep and get deeper as N increases, to the point of mutual unreachability as $N \to \infty$.

5.5. Markov Chain Monte Carlo

Let $\phi(q, p)$ be a scalar function of the q's and p's (i.e., $\phi : \Gamma \to \mathbb{R}$). We want to compute the expectation value of ϕ with respect to the canonical density:

$$E[\phi] = \int_\Gamma \phi(q, p) \frac{e^{-H(q,p)/T}}{Z} \, dq \, dp.$$

The estimation of such integrals is difficult because the number of variables is typically huge, and $e^{-\beta H(q,p)}$ is usually very small except on a very small part of Γ, which sampling at random will rarely find. The computation can become exceedingly long.

Indeed, consider a one-dimensional Ising model. The spins now live on a one-dimensional lattice. The Hamiltonian H associated with a

microstate is

$$H = -\sum_{i=1}^{n} s_i s_{i+1},$$

where, as before, the domain is periodic so that $s_{i+n} = s_i$. Take the case $n = 4$. There are $2^4 = 16$ possible microstates of the chain; for instance, one possible microstate is $S = (+1, -1, -1, +1)$. The possible values of the Hamiltonian are -4, 0, and 4. There are 2 microstates with $H = -4$ (these are the microstates for which all s_i's are of the same sign), 12 microstates with $H = 0$, and 2 microstates with $H = 4$ (the microstates with alternating signs). Suppose the temperature is $T = 1$; then, using (5.5), the two microstates with all s_i's of the same sign have probability of about 0.45. Together they have probability 0.9 of appearing. The next most likely microstate has a probability of only 0.008. The situation becomes even more dramatic as the number of sites in the Ising lattice increases. In general, there will be a very small number of microstates with significant probabilities and a very large number of microstates with probabilities near zero. Thus, if we want to compute the average of some random variable $\phi(S)$, it would not make sense to sample each site with equal frequency. It would be good to construct a chain that visits the sites with probability approximately equal to

$$\pi_i = \frac{1}{Z} e^{-H_i/T}.$$

This is what we called importance sampling in chapter 2.

An excellent method for doing this is "Markov chain Monte Carlo" or "metropolis sampling" or "rejection sampling," which will now be explained. The idea is to walk among the microstates and learn along the way how to reach the probable microstates. Consider a system with a finite number of microstates $S_1, S_2, \ldots, S_n$ (as in the Ising case). Each microstate consists of a list of +1s and −1s. To each such microstate we assign a value $H_i = H(S_i)$ of a Hamiltonian H and a probability

$$P_i = P(S_i) = \frac{e^{-H_i/T}}{Z}, \tag{5.5}$$

where

$$Z = \sum_{i=1}^{n} e^{-H_i/T}.$$

Suppose $\phi = \phi(S)$ is a function on the space $\Gamma = \{S_1, \dots, S_n\}$. We have

$$E[\phi] = \sum_{i=1}^{n} \phi(S_i) P_i = \sum_{i=1}^{n} \phi(S_i) \frac{e^{-H_i/T}}{Z}.$$

DEFINITION. Consider a space Γ containing the microstates S_1, $S_2, \dots, S_n$. A random chain on Γ is a discrete time stochastic process (see chapter 4) such that for each instant t, $X_t = S_j$ for some j such that $1 \leq j \leq n$.

DEFINITION. The probability

$$P(X_t = S_j | X_{t-1} = S_{j_1}, X_{t-2} = S_{j_2}, \dots)$$

is called the transition probability of the chain. The chain is a Markov chain if

$$P(X_t = S_j | X_{t-1} = S_i, X_{t-2} = S_{i_2}, \dots) = P(X_t = S_j | X_{t-1} = S_i).$$

For a Markov chain, we write

$$P(X_t = S_j | X_{t-1} = S_i) = p_{ij} = P(S_i \to S_j),$$

where $\sum_j p_{ij} = 1$ and $p_{ij} \geq 0$. The matrix M with elements p_{ij} is called the transition (or Markov) matrix.

Suppose that we know $P(S_i \to S_j) = p_{ij}$. We have

$$\begin{aligned} P(X_t = S_j | X_{t-2} = S_i) &= \sum_k P(S_i \to S_k) P(S_k \to S_j) \\ &= \sum_k p_{ik} p_{kj}, \end{aligned}$$

which are the entries of the matrix M^2. If $M^{(2)}$ is the matrix whose entries are the probabilities that we go from S_i to S_j in two steps, then $M^{(2)} = M^2$.

DEFINITION. A Markov chain is ergodic in Γ if if given any two microstates S_i and S_j in Γ (where we may have $i = j$), there is a nonzero probability of the chain going from S_i to S_j in n steps for some n.

In other words, a chain is ergodic if the ij element of M^n is, for every pair i, j, non-zero for some n.

The following theorem holds, but we provide no proof here.

THEOREM 5.1. *If a Markov chain is ergodic in Γ, then there exist numbers π_i such that $\pi_i > 0$, $\sum_i \pi_i = 1$, and $\pi_j = \sum_i \pi_i p_{ij}$.*

It is easy to check that if the probability of the microstate S_i is π_i, this probability is unchanged by a step in the chain. This probability density is also attractive, so that as N increases, the fraction of times that the chain visits S_i converges to π_i. The condition $\pi_j = \sum_i \pi_i p_{ij}$ may be hard to check, but it is implied by the simpler sufficient condition $\pi_i p_{ij} = \pi_j p_{ji}$, known as the "detailed balance condition," which is usually quite easy to check. The set of probabilities $\{\pi_i\}$ is the analog, for a Markov chain, of an invariant density for a Hamiltonian system. Consider now the expression $\frac{1}{N}\sum_{t=1}^{N} \phi(S_t)$, where ϕ is a function we are trying to average and the sequence of S_t belongs to the ergodic random chain we are constructing. As $N \to \infty$ this expression converges to the sum $\sum_{i=1}^{n} \phi(S_i)\pi_i$, i.e. the average over the chain converges to the average over the equilibrium density. This conclusion resembles an ergodic theorem, as such theorems were presented at the end of Section 5.3; the reason we can assert the theorem here, while in that previous section we could in general only assume it, is that here we have asserted as fact that the chain was ergodic.

All we have to do now is figure out transition probabilities for which the invariant weights π_i are the given weights $e^{-\beta H}/Z$. The resulting algorithm would be an almost perfect instance of importance sampling, because every microstate would eventually be visited with a frequency equal to its probability. The construction has two steps.

Step 1. Construct an arbitrary ergodic symmetric Markov chain (a Markov chain is symmetric if $p_{ij} = p_{ji}$). For example, in the Ising case, start with an arbitrary microstate, and at each time step, pick a number i between 1 and n with equal probability and change the value s_i associated with the site i to the opposite value: $s_i \to -s_i$. By definition, $\sum_j p_{ij} = 1$ for all i; this first Markov chain is designed to sample all the microstates but can be very inefficient (i.e., spend a lot of time in unimportant microstates). It has an invariant set of π_i, but not the one we want.

Step 2. Let the Markov process defined above have transition probabilities p_{ij}. Construct a modified Markov chain by defining new transition probabilities p^*_{ij} as follows:
If $i \neq j$:

$$p^*_{ij} = \begin{cases} p_{ij}\dfrac{\pi_j}{\pi_i}, & \text{if } \dfrac{\pi_j}{\pi_i} < 1 \\ p_{ij}, & \text{if } \dfrac{\pi_j}{\pi_i} \geq 1. \end{cases}$$

If $i = j$:

$$p^*_{ii} = p_{ii} + \sum p_{ik} \left(1 - \frac{\pi_k}{\pi_i}\right),$$

where the sum is over all k such that $\pi_k/\pi_i < 1$.

We claim that on the average the modified process will visit the microstate S_j $100\pi_j$ percentage of the time. This is a consequence of the equilibrium condition

$$\pi_j = \sum_i \pi_i p^*_{ij},$$

which is a consequence of the detailed balance equation $\pi_i p^*_{ij} = \pi_j p^*_{ji}$

How to apply this result: Let M be the transition matrix of an ergodic Markov process on the microstates $\{S_j\}$. Suppose that we are currently in the microstate S_i. We use M to pick the next microstate S_j; the transition probability of this is p_{ij}. Having picked S_j in this way, we calculate the ratio π_j/π_i. If $\pi_j/\pi_i \geq 1$, we accept S_j as the new microstate. On the other hand, if $\pi_j/\pi_i < 1$, then, with probability π_j/π_i, we accept S_j as the new microstate, and with probability $1 - \pi_j/\pi_i$, we take the old microstate S_i to be the new microstate. This procedure gives the transition probabilities p^*_{ij} defined above.

Two things should be noted. First, for an Ising model,

$$\frac{\pi_j}{\pi_i} = \exp\left(-\frac{H(S_j)}{T} + \frac{H(S_i)}{T}\right) = \exp\left(-\frac{\Delta H}{T}\right),$$

where ΔH is the difference in energy between the microstates S_i and S_j, so that the value of Z is never needed. Second, the change ΔH can be computed if one knows the values of the neighboring spins and is therefore very inexpensive to find.

This construction is easy to program and quite efficient in general. The exception is in more than one space dimension for T near the critical value T_c. We have seen that the error in Monte Carlo methods depends on the number of samples used, and estimated it on the assumption that these samples were independent. However, the successive microstates generated by the Markov chain are not independent. This is OK if the covariances between the successive microstates die out quickly, and they usually do. However, near T_c in two or more dimensions, the spatial correlation length is very large and so is the temporal correlation time of the Monte Carlo samples—more and more metropolis moves are needed to obtain a spin microstate independent of the previous one, and the cost of the calculation diverges (this is known as "critical slowing down"). A cure will be described in the next section.

5.6. Renormalization

As we have just seen, Monte Carlo methods fail when T is near T_c for the Ising model in two or more dimensions. We now sketch a methodology that makes it possible to study the neighborhood of T_c. The idea is to construct a sequence of auxiliary problems, each with fewer variables than the one before, while preserving the pdfs of the variables that remain. The properties of the neighborhood of the critical point will emerge from the examination of the relations between these auxiliary problems. We begin by a definition:

DEFINITION. Suppose you have a set of random variables, $\xi_1, \xi_2, \ldots, \xi_n$. Let their pdf (probability density function) be $f(x_1, \ldots, x_n)$. Pick a subset of the variables, say $\xi_1, \ldots, \xi_m$, with $m < n$. Call this subset $\hat{s}$. The pdf of $\hat{s}$ is $\int f dx_{m+1} \cdots dx_n$. This pdf, the pdf of a subset of variables, is called a marginal of f. In the case of discrete variables the integral is replaced by a sum in the obvious way.

We now set out to calculate marginals for subsets of spins in Ising models. Assume from now on that n, the number of spins, is a power of 2, and that the array of spins is continued to the whole line by periodicity, $s_{i+n} = s_i$. Consider again the one-dimensional case, with spins indexed by the integer i, and let $\hat{s}$ be the set of spins with odd indices, i.e., $s_1, s_3, ...$; let the set of spins not in $\hat{s}$ be called $\tilde{s}$. The computation of the marginal for $\hat{s}$ requires in principle a huge summation, but it can in fact be performed easily and exactly. It is convenient at this point to introduce a new notation. Let $W = -\beta H$, where H is the Hamiltonian; (this W is not to be confused with the probability density in the discussion of the Fokker-Planck equation in chapter 3). The introduction of W frees us of the need to keep track of β and of stray minus signs in the calculations to come. We shall also refer to W as a Hamiltonian. First, note that one can add a constant to the Hamiltonian with impunity: let A be a constant; $e^{W+A} = e^W e^A$, but Z is also multiplied by e^A and in computing the probability of a microstate the factors e^A cancel. Then add such a harmless constant to W, call it nA_0, and give β the new name K_0. Now $W = W^{(0)} = nA_0 + K_0 \sum_i s_i s_{i+1}$, where a superscript (0) has been added to W. Write the marginal for $\hat{s}$ in the form as $e^{W^{(1)}}$ (this can be done because in the Ising model the pdf is always positive). When $e^{W^{(1)}}$ is a marginal of $e^{W^{(0)}}$ the Hamiltonian $W^{(1)}$ is said to be renormalized from $W^{(0)}$. Assume that $W^{(1)}$ also involves products of spins with their new neighbors (spins two locations away in the original problem), so that the probability of

a state of $\hat{s}$ is $\exp W^{(1)} = \exp\left(\frac{n}{2}A_1 + K_1 \sum_i s_i s_{i+2}\right)/Z_1$, where $W^{(1)}$ is the new Hamiltonian, A_1, K_1 are constants, the sum in the exponentials is over next-nearest neighbors (which are the nearest members of $\hat{s}$), and Z_1 is the new normalization factor. The assumption is correct if the constants A_1, K_1 satisfy

$$e^{(n/2)A_1 + K_1 \sum s_i s_{i+2}}/Z_1 = \sum_{\tilde{s}} e^{nA_0 + K_0 \sum s_i s_{i+1}}/Z, \tag{5.6}$$

where $\sum_{\tilde{s}}$ is the sum over $s_2 = -1, +1$, $s_4 = -1, +1, \ldots$. This last equation is satisfied if the following equations hold for all values of s_1, s_3:

$$e^{A_1 + K_1 s_1 s_3} = \sum_{s_2 = \pm 1} e^{A_0 + K_0(s_1 s_2 + s_2 s_3)} \tag{5.7}$$

and if $Z_1 = Z$ as well. This gives four equations for the two parameters A_1, K_1, one for each pair of values of s_1, s_3, but miraculously one gets the same equation for both cases where $s_1 s_3 = 1$, and again for the cases where $s_1 s_3 = -1$. These equations yield:

$$K_1 = \frac{1}{2}\log\cosh(2K_0), \tag{5.8}$$

$$A_1 = \log 2 + \frac{1}{2}A_0 + K_1. \tag{5.9}$$

The marginal has been computed with little pain. It is important to observe that to make the iteration work one had to introduce a seemingly redundant constant A_0, whose value was zero, but that the successive renormalized Hamiltonians acquired non-zero, but still harmless, values of the constant.

The process can be repeated an arbitrary number of times, with the parameters A_n, K_n after n iterations computable from A_{n-1}, K_{n-1}. One obtains a nested sequence of smaller subsystems, with the probabilities of the configurations in each equal to their marginals in the original spin system. The calculus inequality $\log\cosh x - x < 0$ for $x > 0$ shows that the K_n decrease as n increases, so that the variables in the subsystems become more and more independent, until, in the limit, $K_n \to 0$ and the variables are completely independent. If initially $\beta = 0$ the successive Hamiltonians are all zero, and if initially $T = 0$ the successive Hamiltonians have $K_n = \infty$; these values of β or T are invariant under marginalization or renormalization.

Suppose ξ is the correlation length in any one of these systems, defined as the distance such that the covariance of s_i, s_j is negligible if $|i - j| > \xi$ but not if $|i - j| < \xi$. Each time we marginalize, the correlation length in units of interspin separation decreases by a factor

of 2. Indeed, one can start from the original spin problem, marginalize once, and then move the 3d spin to location 2 — no harm is done, because spin 2 is out of the game- then move spin 5 to location 3, etc. Now one has a smaller system, identical to the marginalized system apart from the labeling of the spins, whose the correlation length is obviously half of the original one. As the correlation length shrinks to zero, the system approaches a system of independent spins, as we already know.

The preceding construction makes it possible to sample the one-dimensional Ising model effectively without Markov chains. Indeed, decimate until the reduced system has 2 spins per period. This reduced problem has four states with probabilities p_1, p_2, p_3, p_4 which add up to 1. It is easy to set up a sampling scheme which samples each one of these states with a frequency equal to its probability. Once this is done, go to the subset with 4 spins; 2 of these have been sampled, each of the remaining 2 has known neighbors and only 2 possible states, which again can be sampled with a frequency equal to their probability; keep on filling the system in. The result is a sample of the original system, each state is sampled with a frequency equal to its probability, and two successive samples are independent. There is no Markov chain and no memory.

Unfortunately, in more than one dimension marginalization is not so trivial to achieve, but it can still be done approximately quite well. For example, in two dimensions, choose $\hat{s}$, the set of variables to keep, to be the spins with indices i, j both of which are even. The pdf we start with is $e^W/Z = e^{W^{(0)}}/Z$; try to find a marginal of the form $e^{W^{(1)}}/Z$ (same Z). If we try a $W^{(1)}$ as simple as before we will fail, but a slightly more complicated $W^{(1)}$, with extra terms that contain products of non-neighboring spins, with additional coefficients, say, $L, M, \ldots$, will do quite well as long as one is happy with approximations; for example one can try

$$\begin{aligned} W^{(1)} &= K_1 \sum s_{i,j}(s_{i+2,j} + s_{i,j+2}) + L_1 \sum s_{i,j}(s_{i+2,j+2} + s_{i+2,j-2}) \\ &\quad + M_1 \sum s_{i,j}(s_{i+2,j} + s_{i-2,j} + s_{i,j+2} + s_{i,j-2})^2; \end{aligned} \tag{5.10}$$

one can view $W^{(0)}$ as having the same form with $K_0 = \beta$ and $L_0 = M_0 = 0$. The operation can be repeated. The numerical values of the constants $K, L...$ are determined by requiring, as least approximately, that $e^{W^{(1)}}$ be a marginal of $e^{W^{(0)}}$; we omit the details, which are messy but not hard in principle. The operation can be repeated. The successive Hamiltonians are described by the values of a finite number

of constants, say $\mathbf{K}_n = (K_n, L_n, M_n)$ for the n-th member of the sequence; in the "parameter space", where the coordinates are $L, M...$, the parameters evolve and undergo a "parameter flow", along trajectories that can be labeled by the value of T (and of possibly other parameters) at their starting point.

DEFINITION. The mapping from one pdf to the next is a marginalization, the mapping from one Hamiltonian to the next (i.e., from the parameters $\mathbf{K}$ that describe one Hamiltonian to the parameters that describe the next) is a renormalization, the set of renormalization transformations is called the renormalization group, RNG for short.

Denote the transformation of the parameters $\mathbf{K}$ (and therefore of the Hamiltonian W) via a marginalization/renormalization by $\mathbf{R}$, so that $W^{(1)} = \mathbf{R}W^{(0)}, W^{(2)} = \mathbf{R}W^{(1)}$ etc.

The moral is that marginals for subsets of variables that form some regular pattern in the plane (or in space) can be computed quite effectively by guessing a suitable form for the corresponding Hamiltonians. Just as in one dimension, each successive Hamiltonian describes a system with a smaller correlation length. The (parameters in the) successive Hamiltonians can be thought of as "flowing" in the space of Hamiltonians. At this point is not clear what this has to do with critical points.

Before proceeding, some remarks: the main point about marginalization/renormalization is that it produces a system with a shorter correlation length while preserving the pdf of the variables that remain. This can be generalized in various ways. For example, one can divide the set of spins into 2 by 2 or 3 by 3 blocks and devise some rule for attaching a new spin (a variable that takes on only the values ± 1) to each block; one can then guess a Hamiltonian that describes the pdf of the new variables along the lines just sketched. This is what is usually meant by renormalization in the literature. What has been described above is the special case where the value of the spin attached to a block is just the value of one of the spins in the block. This is often called a decimation.

The reader is invited to note the resemblance of renormalization, as it was just generalized, the the second construction of the central limit theorem in Section 2.5. Indeed, one can view renormalization as a generalization of the central limit theorem to certain sets of non-independent variables.

We now turn to a description of the critical point at T_c in, say, two space dimensions, where the mean value m of the magnetization goes from being non-zero (for $T < T_c$) to zero. In the limit of very

large array size, the properties of the system of spins are described by "critical exponents"; for example, for T smaller than T_c but close to it, m is proportional to $|T - T_c|^b$, where $b = 1/6$; b is a critical exponent. One can have such exponents only if the correlation length ξ is infinite; if ξ is finite, one can calculate m by adding up a finite number of exponential functions; such sums are analytic functions and no non-integer power can appear. Indeed, one can show that ξ is proportional to $|T - T_c|^{-\nu}$, where ν is another exponent. These exponents cannot be effectively calculated by Markov chain Monte Carlo; one needs a very large array of spins to approximate a system of spins with long correlations, and in addition, near T_c, Markov chain Monte Carlo has very long correlation times which make the computations extremely expensive. The exponents are of great interest, and we now sketch how they can be found from the RNG.

The fact that at T_c the correlation length is infinite means that the renormalization flow has a fixed point at $T = T_c$ (because a fraction of infinity is still infinity). This does not mean that the point $K_0 = 1/T_c, L_0 = 0, M_0 = 0$ is invariant, only that if one starts from that point one converges to a fixed point. This fixed point is unstable under the renormalization flow, because if one starts a flow off that point the correlation length is in general finite and each RNG transformation will make it smaller, i.e., take one further away.

Let W^* be the Hamiltonian evaluated at the fixed point, so that $W^* = \mathbf{R}(W^*)$; Let $W^{(0)}$ be close to the W^*; one can write $W^{(0)} = W^* + \delta W$ (what varies are of course the coefficients $\mathbf{K}$; δW is a vector of increments $\delta K, \delta L....$) Apply the RNG: $W^{(1)} = \mathbf{R}(W^* + \delta W) = \mathbf{R}(W^*) + A\delta W$, where the matrix A is the matrix of the derivatives of the coefficients in $W^{(1)}$ with respect to the coefficients of $W^{(0)}$, evaluated at W^*; it is a fixed matrix, a property of the RNG. The claim is that the critical exponents can be found if A is known.

We demonstrate this in the special case of the exponent ν in the relation $\xi = constant \cdot |T - T_c|^{-\nu}$. Suppose you find yourself near the unstable fixed point on some trajectory that started from a system with T near T_c but not equal to T_c. Your correlation length is some finite number ξ. Start renormalizing. At each step ξ is reduced by a factor d (defined by your choice of grouping in the RNG); you leave the neighborhood of the fixed point when $\xi/d^n = u$, where u is a number of order 1 (the eventual formula will be such that the exact value of u does not matter). Now find the eigenvalues and eigenvectors of the matrix A. A calculation we do not reproduce reveals that λ, the largest eigenvalue in absolute value, is real and larger than one; let e be the corresponding eigenvector of length one. Write the coefficients

of δW that you started from in a basis where e is one of the basis vectors, $W = W^* + A\delta W = W^* + A\Delta e + ... = W^* + \lambda\Delta e + \ldots$, where Δ is the component of δW along e, and the dots denote terms that will be relatively small in what follows. Apply $\mathbf{R}$; the new W is $\mathbf{R}(W^* + A\Delta e) = W^* + \lambda^2\Delta e$; after n steps you will be at $W^* + \lambda^n\Delta e$, and you if you leave the neighborhood after n steps, the quantity $\lambda^n\Delta$ should be of order one.

The coefficient Δ depends on the trajectory on which you are, and therefore depends on the temperature T at its starting point, $\Delta = \Delta(T)$; if you start at T_c the trajectory enters the fixed point and has no component along the vector that leaves the fixed point, $\Delta(T_c) = 0$. Assuming some smoothness in $\Delta(T)$, we can write $\Delta = c(T - T_c)$, where one can view c as constant as long as one stays in the neighborhood of the fixed point. Taking logs of the two equations that characterize n (one in terms of ξ and the other in terms of Δ), and assuming that the the analysis is the same just above T_c and just below T_c so that $\Delta = c|T - T_c|$, we find that $\xi = constant \cdot |T - T_c|^{-\nu}$, where $\nu = \log d/\log\lambda$, an expression that depends only on the matrix A.

Finally, try to obtain ξ, the correlation length for the Ising model, by dimensional analysis. ξ can depend on ℓ, the interspin distance, s, the magnitude of the spins (which so far has always been 1), and the temperature T (or $T - T_c$ for convenience). In dimensionless variables, we find $\xi/\ell = \Phi(|T - T_c|/s^2)$, where Φ is an unknown dimensionless function, and we assumed again that the relationship is symmetric about T_c. Note that we find again that ξ should be measured in units of interspin distance.

Make now the complete similarity assumption $\Phi(0) = B$, where B is a non-zero finite constant (see Section 4.3). The result is $\xi/\ell = B$ at T_c, which we know is not true. Try then an incomplete similarity assumption, $\Phi(|T - T_c|/s^2) = (|T - T_c|/s^2)^\gamma \Phi_1(|T - T_c|/s^2)$, where $\Phi_1(0)$ is a non-zero constant and γ is an anomalous exponent. We find that this fits the previous analysis, with $\gamma = -\nu$. Thus the exponents revealed by the RNG are anomalous exponents in the sense of dimensional analysis.

5.7. Exercises

1. Consider complex variables $u_j = q_j + ip_j$ (j is an index, i is $\sqrt{-1}$) at the points jh, where j takes the values $..., -1, 0, 1, 2, ..., h = 2\pi/n$,

n is an integer, and $u_{j+n} = u_j$. Consider the Hamiltonian

$$H = \frac{1}{2}\sum_{j=1}^{n}\left[\left(\frac{q_{j+1}-q_j}{h}\right)^2 + \left(\frac{p_{j+1}-p_j}{h}\right)^2 + \frac{1}{2}(q_j^4+p_j^4)\right].$$

Treat the q, p as conjugate variables (i.e. p_j is the momentum associated with the position variable q_j) and derive the equations of motion. Check formally that as $h \to 0$, these equations converge to the nonlinear Schroedinger equation $iu_t = -u_{xx} + q^3 + ip^3$. Suppose the initial data for the equation are picked from the density $Z^{-1}e^{-H/T}$ for some $T > 0$. By comparing the Hamiltonian with the Feynman-Kac formula in the physicists' notation, deduce that a typical solution of the equation with this kind of data has no derivatives in x (and is therefore a "weak" solution). Check that as $h \to 0$, the Hamiltonian times h converges to the integral

$$\frac{1}{2}\int_0^{2\pi}\left(q_x^2 + p_x^2 + \frac{1}{2}(q^4+p^4)\right)dx.$$

You may notice that as the mesh size $h \to 0$, the temperature T must also tend to zero.

2. Compute the magnetization m in the Ising model in 2 dimensions by Markov chain Monte Carlo, on a 30 by 30 lattice, for $\beta = 1$ and $\beta = 0.2$, and compare with the exact answer: $m = [1-\sinh(2\beta)^{-4}]^{1/8}$ for β larger than the critical value $1/T_c = .4408$ and $m = 0$ otherwise.

3. Write a program that generates all 2^{n^2} microstates of a two-dimensional periodic Ising model. Define the magnetization μ as the sum of all the spins divided by n^2. For $n = 3$ and $\beta = 1, \beta = 0.01$, make a histogram of the probabilities of the various values of μ; note the different qualitative behavior at low β and at high β. Estimate the fraction of microstates which have probabilities less than 10^{-6}. Observe that it is difficult to estimate the histogram above by a Monte Carlo program where microstates are sampled at random rather than examined one after the other. Note that the large probabilities of extreme values at high β (small T) come from the fact that the probabilities of the extreme microstates are very high; at low β each microstate with small $|\mu|$ is still less likely than a microstate with an extreme value, but the small values of $|\mu|$ win because there are many such microstates. Programming notes: A relatively easy way to generate all microstates of the Ising model is to write the integers in a base 2 representation and then map the $0, 1$ coefficients onto $-1, +1$. You have to be careful because you

may generate high values of $e^{-\beta H}$ even when $n = 3$ — one way out is to use double precision. To find the fraction of microstates with very low probabilities you need the partition function, which has to be computed; you may therefore need to run your program twice.

4. Calculate the entropy of the pdf $f(x) = e^{-x^2}/\sqrt{\pi}$. Do the same for the microcanonical density for the Hamiltonian $H = \sum_i p_i^2/2m$, where m is a (constant) mass (the second part is not trivial)..

5. Consider a particle with position q, momentum p, and Hamiltonian $H = (1/2)(q^2+p^2)$. Derive the equations of motion and the Liouville equation. Then derive a Fokker-Planck equation for the equations of motion by the methods of chapter 3 and check that it coincides with the Liouville equation.

6. Verify the identity $(Lu, v) = -(u, Lv)$ at the end of Section 5.2.

7. Consider two particles moving on the line with positions q_1, q_2 and momenta p_1, p_2. The Hamiltonian is $H = (1/2)(q_1^2 + q_2^2 + q_1^2 q_2^2 + p_1^2 + p_2^2)$. Write down the equations of motion. Suppose the inital data are drawn at random from the pdf (probability density function) $Z^{-1}e^{-H/T}$; check that this density is invariant in time.

8. Consider again the Hamiltonian $(1/2)(q_1^2 + q_2^2 + p_1^2 + p_2^2 + q_1^2 q_2^2)$. Suppose you are unable to solve the equations for q_2, p_2 (maybe because you don't have initial data). The equations for q_1, p_1 have the form $dq_1/dt = R_1(q_1, q_2, p_1, p_2)$, and similarly $dp_1/dt = R_2$. At time t you don't have values of q_2, p_2 to substitute into R_1, R_2. The best you can do is replace R_1, R_2 by their best approximation by a function of q_1, p_1 only, i.e., replace R_1 by $E[R_1|q_1, p_1]$, and the same for R_2, using the invariant probability density (this amounts to averaging the equations over the uncomputable parameters, as is often done in various contexts). Show that the new system you got is also Hamiltonian, with Hamiltonian $\hat{H} = -T \log \int e^{-H(q_1,q_2,p_1,p_2)/T} dq_2 dp_2$. (Note: This construction is very useful in some contexts, but should be treated with caution. One has to consider carefully what the solution $q_1(t), p_1(t)$ represents. It is obviously not the same as the $q_1(t), p_1(t)$ of the original four-by-four system. In the next chapter we shall see that it is not the best estimate of $q_1(t), p_1(t)$ in the original system either. Averaging over unknown quantities is dangerous.)

5.8. Bibliography

[1] G.I. Barenblatt, *Scaling, Self-Similarity, and Intermediate Asymptotics*, Cambridge University Press, Cambridge, 1996.
[2] J. Binney, N. Dowrick, A. Fisher, and M. Newman, *The Theory of Critical Phenomena, an Introduction to the Renormalization Group*, the Clarendon Press, Oxford, 1992.
[3] A.J. Chorin, Conditional expectations and renormalization, Multiscale Model. Simul. 1 (2003), pp. 105-118.
[4] A.J. Chorin, Monte Carlo without chains, Comm. Appl. Math. Comp. Sc. 3 (2008), pp. 77-93.
[5] H. Dym and H. McKean, *Fourier Series and Integrals*, Academic Press, New York, 1972.
[6] L.C. Evans, Entropy and Partial Differential Equations, Lecture notes, UC Berkeley Mathematics Department, 1996.
[7] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group*, Perseus, Reading, MA, 1992.
[8] J. Hammersley and D. Handscomb, *Monte-Carlo Methods*, Methuen, London, 1964.
[9] E.T. Jaynes, *Papers on Probability, Statistics and Statistical Physics*, Kluwer, Boston, 1983.
[10] G. Jona-Lasinio, The renormalization group–a probabilistic view, Nuovo Cimento 26 (1975), pp. 99-118.
[11] L. Kadanoff, *Statistical Physics: Statics, Dynamics, and Renormalization*, World Scientific, Singapore, 1999.
[12] D. Kandel, E. Domany, and A. Brandt, Simulation without critical slowing down–Ising and 3-state Potts model, Phys. Rev. B 40 (1989), pp. 330-344.
[13] J. Lebowitz, H. Rose, and E. Speer, Statistical mechanics of nonlinear Schroedinger equations, J. Stat. Phys. 54 (1988), pp. 657-687.
[14] J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer, NY, 2002.
[15] C. Thompson, *Mathematical Statistical Mechanics*, Princeton University Press, Princeton, NJ, 1972.

CHAPTER 6

Time-Dependent Statistical Mechanics

6.1. More on the Langevin Equation

We now turn to the statistical mechanics of systems not in equilibrium. The first few sections are devoted to special cases, which will be used to build up experience with the questions one can reasonably ask and the kinds of answer one may expect. A general formalism will follow, with applications.

Consider first the Langevin equation, already discussed in Section 3.8, which we now write as

$$du = -au\,dt + \sqrt{2D}\,dw, \tag{6.1}$$

where w is a Brownian motion, as before. A constant factor $\sqrt{2D}$ has been inserted in front of the noise. Equation (6.1) models the dynamics of a heavy particle bombarded by light particles. This equation will now be fully solved. (See Section 3.8 for an earlier solution.)

After multiplication of (6.1) by e^{at}, we get

$$d(ue^{at}) = \sqrt{2D}e^{at}\,dw. \tag{6.2}$$

Integrating both sides from 0 to t gives

$$\int_0^t d(ue^{as}) = \sqrt{2D}\int_0^t e^{as}\,dw.$$

Let $u(0) = b$. Then

$$u(t)e^{at} - b = \sqrt{2D}\int_0^t e^{as}\,dw.$$

After multiplying both sides by e^{-at}, we obtain

$$u(t) - be^{-at} = \sqrt{2D}\int_0^t e^{a(s-t)}dw.$$

The last integral may be rewritten in the form

$$\int_0^t e^{a(s-t)}dw = \lim_{\Delta\to 0}\sum_{j=0}^{n-1} e^{a(j\Delta-t)}(w((j+1)\Delta) - w(j\Delta))$$

(where one does not have to worry about the Ito/Stratonovich dichotomy because the coefficient is non-random, and the two formalisms are equivalent). The summands of the last sum are independent Gaussian variables with mean 0. The variance of the sum is the sum of variances of its summands; i.e.,

$$\mathrm{Var}\left(\sum_{j=0}^{n-1} e^{a(j\Delta-t)}(w((j+1)\Delta)-w(j\Delta))\right)=\sum_{j=0}^{n-1}\Delta e^{2a(j\Delta-t)};$$

taking the limit $\Delta\to 0$ we find

$$\mathrm{Var}\left(\int_0^t e^{a(s-t)}dw\right)=\int_0^t e^{2a(s-t)}ds=\frac{1}{2a}-\frac{1}{2a}e^{-2at}.$$

As $t\to\infty$, this variance tends to $1/(2a)$. Also, as $t\to\infty$, be^{-at} tends to zero. Therefore, the solution $u(t)$ of the Langevin equation (6.1) tends to a Gaussian variable with mean 0 and variance D/a.

If the particle we are observing has mass m and if we interpret u as its velocity, then its energy is $\frac{1}{2}mu^2$. According to what we found in chapter 5, the probability that the particle has velocity u is proportional to $\exp(-mu^2/2T)$. Thus, we must have

$$a=\frac{Dm}{T}.$$

The coefficient a is a friction coefficient, and the relation between the friction and the temperature is an instance of a "fluctuation-dissipation theorem"; it is a consequence of the requirement that the system tend to equilibrium for long times and it relates the rate of dissipation of energy to the amplitude T of the "thermal fluctuations."

Note that we have not provided an explanation of the conditions for the validity of our modeling of the motion of a heavy particle under the influence of many others, in particular for the validity of the modeling of the interactions as white noise. This looks plausible, but should be discussed further.

The discussion of the fluctuation-dissipation theorem can be also presented in terms of the Fokker-Planck equation associated with the problem. We do that in a slightly more general case. Consider a particle of mass m subjected to noise, with the following equations of motion:

$$\begin{aligned}\dot q&=\frac{\partial H}{\partial p},\\ \dot p&=-\frac{\partial H}{\partial q}-a\frac{\partial H}{\partial p}+\sqrt{2D}\,\frac{dw(t)}{dt},\end{aligned}$$

where $H = p^2/2m + Kq^2/2$ is the Hamiltonian (making the system a harmonic oscillator), a and D are constants as above, and $w(t)$ is Brownian motion (BM). Substitution of the specific Hamiltonian into this equation yields

$$\dot{q} = \frac{p}{m}, \tag{6.3}$$

$$\dot{p} = -qK - a\frac{p}{m} + \sqrt{2D}\,\frac{dw}{dt}. \tag{6.4}$$

Note that we still have offered no motivation for the use of white noise. The presence of an extra term in addition to the usual derivatives of H and to the noise is motivated by the discussion earlier in this section, where a dissipation term appeared, and will be fully explained by the result below.

A slight generalization of the argument in Section 3.8 yields the following Fokker-Planck equation for the probability density $W(p, q, t)$ of p, q:

$$\frac{\partial W}{\partial t} = \frac{\partial J_1}{\partial q} + \frac{\partial J_2}{\partial p},$$

where (J_1, J_2) is the probability flux vector

$$J_1 = -p\frac{W}{m}, \quad J_2 = KqW + ap\frac{W}{m} + D\frac{\partial W}{\partial p}.$$

This equation allows $W = Z^{-1}e^{-H/T}$ as a stationary density provided $a = D/T$, in agreement with the result above (in (6.4), the coefficient a has already been divided by m).

The interpretation of the fluctuation-dissipation theorem in this case is straightforward. The higher the driving force on the particles, the higher the temperature at the eventual equilibrium; the bigger the friction, the lower the temperature. In the present case, the friction and the dissipation come from a single cause—the interaction between the particle we are considering in detail and other, unexamined, particles, just as a runner who runs into a group of milling people will be diverted from his straight path and also slowed down by a single cause—the collisions with those people. The two effects of the single cause at equilibrium are related to each other by the fluctuation-dissipation relation.

6.2. A Coupled System of Harmonic Oscillators

In the previous section, we considered a particle acted upon by noise; the noise presumably represents an interaction with other particles, but the properties of the interaction and the validity of its description as noise were not considered. In this section, we consider, in a simple case, the interaction of a singled-out particle, the "tagged" or "resolved" particle, with other particles in the framework of a Hamiltonian description of the entire system.

The particles are all in a one-dimensional space; the resolved particle is located at x, has velocity v and unit mass, and is acted on by a potential $U(x)$. It interacts with n other particles, located at q_j and having momenta p_j, with $j = 1, ..., n$. The Hamiltonian is

$$H = \frac{1}{2}v^2 + U(x) + \frac{1}{2}\sum_j p_j^2 + \frac{1}{2}\sum_j f_j^2 \left(q_j - \frac{\gamma_j}{f_j^2}x\right)^2, \tag{6.5}$$

where the f_j and γ_j are constants. The γ_j are "coupling constants," and one can check that in the absence of interaction (i.e., if one sets the coupling constants to zero), the f_j would be the frequencies of oscillation of the various particles. This Hamiltonian is quadratic (except for the term in U) so that the equations of motion for the nonresolved particles are linear; this is what makes the problem solvable explicitly. The particles with linear equations of motion of the form implied in this Hamiltonian for the unresolved particles are linear oscillators.

The equations of motion are

$$\begin{aligned}
\dot{x} &= v, \\
\dot{v} &= -\frac{dU}{dx} + \sum_j \gamma_j \left(q_j - \frac{\gamma_j}{f_j^2}x\right), \\
\dot{q}_j &= p_j, \\
\dot{p}_j &= -f_j^2 q_j + \gamma_j x.
\end{aligned}$$

The equations of motion for the unresolved particles can be solved explicitly by the method of variation of constants:

$$q_j(t) = q_j(0)\cos(f_j t) + p_j(0)\frac{\sin(f_j t)}{f_j} + \frac{\gamma_j}{f_j}\int_0^t x(s)\sin(f_j(t-s))\,ds,$$

where $q_j(0)$ and $p_j(0)$ are initial conditions (about which nothing has been said as yet). The integral term in this equation can be rewritten

after integration by parts as

$$\frac{\gamma_j}{f_j^2}(x(t) - x(0)\cos(f_j t)) - \gamma_j \int_0^t v(s)\frac{\cos(f_j(t-s))}{f_j^2}\, ds.$$

Collecting terms and inserting them into the equations for x and v, one finds

$$\dot{x}(t) = v(t), \quad \dot{v}(t) = -U'(x) + \int_0^t K_n(t-s)v(s)\, ds + F_n(t), \tag{6.6}$$

where

$$K_n(t) = -\sum_j \frac{\gamma_j^2}{f_j^2}\cos(f_j t)$$

and

$$F_n(t) = \sum_j \gamma_j \left(q_j(0) - \frac{\gamma_j}{f_j^2}x(0) \right) \cos(f_j t) + \sum_j \gamma_j p_j(0)\frac{\sin(f_j t)}{f_j}.$$

Suppose that the goal is to follow the motion of the resolved particle (the one at x with velocity v) without following the motion of all of the others. Specific initial values $q_j(0)$, $p_j(0)$ cannot be taken into account. The best one can do is sample these initial values for the unresolved particles from some acceptable density, which makes the whole evolution stochastic. The first term on the right-hand side of (6.6) is the effect of a potential that acts on the resolved particle alone at the time t, and it has no analog in the Langevin equations of the previous section. The second term on the right-hand side of (6.6) is analogous to the dissipation term $-au$ in the the previous Langevin equation and represents not only dissipation but also a memory, inasmuch as through this term, the velocity at previous times impacts the current velocity. That a reduced description of the motion of the resolved variable involves a memory should be intuitively obvious: Suppose you have $n > 3$ billiard balls moving about on top of a table and are trying to describe the motion of just three; the second ball may strike the seventh ball at time t_1 and the seventh ball may then strike the third ball at a later time. The third ball then "remembers" the state of the system at time t_1, and if this memory is not encoded in the explicit knowledge of where the seventh ball is at all times, then it has to be encoded in some other way. The analog of this term in the following sections will be called a "memory" term, to emphasize the possibly unfamiliar memory effect. The kernel of this integral term, K_n, does not depend on the initial data and, therefore, this term is not random.

The last term involves the random initial data and is a random function, analogous to the white noise in the Langevin equation of

Section 6.1. Equation (6.6) generalizes the Langevin equation and we shall call the last term the noise term; in general, it is not white noise. White noise can be expanded in terms of sines and cosines, but except under very special conditions, the coefficients in this expansion will not be the ones in the above expression for F_n.

Finally, suppose the initial density W is $W = Z^{-1}e^{-H/T}$, with H given by (6.5). One can readily check that with this choice, $E\left[p_j(0)p_k(0)\right] = T\delta_{jk}$, where δ_{jk} is the Kronecker δ symbol. Also,

$$E\left[\left(q_j(0) - \frac{\gamma_j}{f_j^2}x(0)\right)\left(q_k(0) - \frac{\gamma_k}{f_k^2}x(0)\right)\right] = \frac{T\delta_{jk}}{f_j^2},$$

where $x(0)$ is the nonrandom initial value of $x(t)$. With this choice of initial W, one can also check that

$$E[F_n(t)F_n(t-t')] = -TK_n(t').$$

This is the fluctuation-dissipation theorem relevant to the present problem. It emerges simply as a consequence of the equations of motion combined with the canonical choice of initial density.

It should be noted that the problem in this section is not an equilibrium problem because the Hamiltonian depends on the variable x and changes in time. As time advances, the values of the variable x become increasingly uncertain and the system "decays" to equilibrium; this decay is accomplished by the memory and the noise.

6.3. Mathematical Addenda

A pattern has emerged in the questions asked so far in the present chapter: We consider problems with many variables where thermal equilibrium has not been established (i.e., where there is no probability density invariant in time). Such a density may be established in the future of the systems under study, this fact has present consequences, but there is no universal recipe for the evolution of the probability density and no analog of an ergodic hypothesis to simplify calculations. What one strives for is a reduced, practical description of key variables—the analog of what was called renormalization in the equilibrium case. The reduced equations we have derived replace those parts of the system that are not fully described by a pair of matched terms, a stochastic term that can be called "noise" and a damping, or "memory," term; they have to be matched to preserve the possibility of future equilibrium. The matching conditions are called "fluctuation-dissipation theorems." We now propose to derive these results in some generality; however, before we can embark on this analysis, some mathematical addenda are needed.

6.3.1. How to write a nonlinear system of ordinary differential equations as a linear partial differential equation. Consider a system of ordinary differential equations

$$\frac{d}{dt}\phi(x,t) = R(\phi(x,t)), \quad \phi(x,0) = x, \tag{6.7}$$

where R, ϕ, and x are (possibly infinite dimensional) vectors with components R_i, ϕ_i, and x_i, respectively.

We claim that this nonlinear system can be rewritten as a linear partial differential equation. This is not an approximation, but an exact representation; the cost of getting a linear system is the greater conceptual and practical complexity of having to deal with a partial differential equation.

Define the Liouville operator (as in Section 5.2):

$$L = \sum_i R_i(x)\frac{\partial}{\partial x_i}.$$

It is not assumed here that the system (6.7) is Hamiltonian, so that the coefficient functions in L are not derivatives of some H, as in Section 5.2. The variables in the coefficients and in the differentiations belong to a space with as many dimensions as the space of initial data for (6.7). Now form the differential equation

$$u_t = Lu, \tag{6.8}$$

with initial data $u(x,0) = g(x)$. This is called a Liouville equation, although the sign of the right-hand side is the opposite of the one in front of the right-hand side of the Liouville equation for the probability density in chapter 5. The claim is that the solution of this equation is $u(x,t) = g(\phi(x,t))$, where $\phi(x,t)$ is the solution of the system (6.7) with initial data x. If this is true, one can clearly solve the partial differential equation (6.8) if one can solve the system of ordinary differential equations. Conversely, if one can solve (6.8) for every g, one can set $g(x) = x_j$ and obtain the j-th component of ϕ, for every j.

To prove this claim, we first prove the following useful identity:

$$R(\phi(x,t)) = D_x\phi(x,t)R(x). \tag{6.9}$$

In this formula, $D_x\phi(x,t)$ is the Jacobian of $\phi(x,t)$:

$$D_{x_j}\phi_i(x,t) = \frac{\partial\phi_i}{\partial x_j},$$

and the multiplication on the right-hand side is a matrix vector multiplication; the left-hand side is the vector R evaluated when the argument is ϕ, while on the right the argument of R is x, the initial datum of ϕ; ϕ is assumed to satisfy (6.7).

Define $F(x,t)$ to be the difference of the left-hand side and the right-hand side of (6.9):

$$F(x,t) = R(\phi(x,t)) - D_x\phi(x,t)R(x).$$

Then, at $t = 0$, we have

$$\begin{aligned} F(x,0) &= R(\phi(x,0)) - D_x\phi(x,0)R(x) \\ &= R(x) - D_x(x)R(x) \\ &= R(x) - IR(x) \\ &= 0. \end{aligned} \tag{6.10}$$

Differentiating F with respect to t, and using the chain rule repeatedly, we get

$$\begin{aligned} \frac{\partial}{\partial t}F(x,t) &= \frac{\partial}{\partial t}R(\phi(x,t)) - \frac{\partial}{\partial t}(D_x\phi(x,t)R(x)) \\ &= (D_xR)(\phi(x,t))\frac{\partial}{\partial t}\phi(x,t) - D_x\left(\frac{\partial}{\partial t}\phi(x,t)\right)R(x) \\ &= (D_xR)(\phi(x,t))\frac{\partial}{\partial t}\phi(x,t) - D_x(R(\phi(x,t)))R(x) \\ &= (D_xR)(\phi(x,t))R(\phi(x,t)) - (D_xR)(\phi(x,t))D_x\phi(x,t)R(x) \\ &= (D_xR)(\phi(x,t))\left(R(\phi(x,t)) - D_x\phi(x,t)R(x)\right) \\ &= (D_xR)(\phi(x,t))F(x,t). \end{aligned} \tag{6.11}$$

From (6.10) and (6.11) one can conclude that $F(x,t) \equiv 0$. Indeed, the initial value problem defined by (6.10) and (6.11) has a unique solution given that R and ϕ are smooth. Since $F(x,t) = 0$ solves this problem, we have proved (6.9).

Take an arbitrary smooth function $g(x)$ on Γ and form the function

$$u(x,t) = g(\phi(x,t)).$$

Clearly, $u(x,0) = g(x)$. Differentiate this function with respect to t using the chain rule:

$$\frac{\partial u}{\partial t} = \sum_i \frac{\partial g(\phi(x,t))}{\partial x_i}\frac{\partial \phi_i(x,t)}{\partial t} = \sum_i R_i(\phi(x,t))\frac{\partial g(\phi(x,t))}{\partial x_i}.$$

Using (6.9), this last expression becomes

$$\begin{aligned}\sum_i \left(\sum_j \frac{\partial \phi_i(x,t)}{\partial x_j} R_j(x) \right) \frac{\partial g(\phi(x,t))}{\partial x_i} & \\ &= \sum_j R_j(x) \left(\sum_i \frac{\partial g(\phi(x,t))}{\partial x_i} \right) \frac{\partial \phi_i(x,t)}{\partial x_j} \\ &= \sum_j R_j(x) \frac{\partial g(\phi(x,t))}{\partial x_j} \\ &= Lu. \end{aligned} \tag{6.12}$$

Hence, $u(x,t) = g(\phi(x,t))$ is the (unique) solution of the equation

$$u_t = Lu, \;\; u(x,0) = g(x). \tag{6.13}$$

Clearly, if one can solve the system (6.7) for all x, one can solve the Liouville equation (6.13) for any initial datum g. Conversely, suppose one can solve the Liouville equation for all initial data g and pick $g(x) = x_j$; the solution of the Liouville equation is then $\phi_j(x,t)$, the j^{th} component of the solution of the system of ordinary differential equations (6.7).

If L is skew-symmetric, the Liouville equation for the probability density in chapter 5 and the Liouville equation here, which is equivalent to the original system, differ by a sign, i.e., the system and its probability density move in opposite directions. By judicious insertions of factors of complex i, one can get these two Liouville equations to be adjoint; the two equations are then related like the Schrödinger and Heisenberg representations in quantum mechanics.

6.3.2. More on the semigroup notation. In Section 3.9 we introduced the semigroup notation, according to which the solution of (6.13) is denoted by $e^{tL}g$; the time dependence is explicitly marked, and the value of this solution at a point x is denoted by $e^{tL}g(x)$. With this notation, the formula for the solution $u(x,t) = g(\phi(x,t))$ of (6.13) becomes

$$e^{tL}g(x) = g(e^{tL}x). \tag{6.14}$$

Note that $e^{tL}x$ is not e^{tL} evaluated at x but e^{tL} acting on the vector whose components are the functions x_i; the time propagation of a function g commutes with the time propagation of the initial conditions x_i. This equality has an obvious interpretation: If g is a fixed, time-invariant function of the variables that describe a mechanical system,

it changes in time only inasmuch as these variables change. Equation (6.12) becomes, in the semigroup notation, simply

$$Le^{tL} = e^{tL}L. \tag{6.15}$$

The analogous formula for matrices is, of course, well known.

Let A, B be two operators; the following formula holds for their exponentiations, as defined above via the corresponding equations:

$$e^{t(A+B)} = e^{tA} + \int_0^t e^{(t-s)(A+B)} B e^{sA}\, ds. \tag{6.16}$$

The best way to see that this identity holds is to form the difference $z(t)$ between the right-hand side and the left-hand side:

$$z(t) = e^{t(A+B)} - e^{tA} + \int_0^t e^{(t-s)(A+B)} B e^{sA}\, ds \tag{6.17}$$

and check that $z(0) = 0$ and $z'(t) = (A+B)z(t)$; by the uniqueness of the solution of the ordinary differential equation, $z(t) = 0$ for all t. This formula is often called the "Duhamel formula" or, in physics, the "Dyson formula."

6.3.3. Hermite polynomials and projections. The polynomials orthonormal with respect to the inner product

$$(u, v) = \int_{-\infty}^{+\infty} \frac{e^{-x^2/2}}{\sqrt{2\pi}} u(x)v(x)\, dx$$

are called the Hermite polynomials. One can generalize them to spaces with more dimensions: If one defines the inner product

$$(u, v) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} (2\pi)^{-n/2}\, e^{-(\sum x_i^2)/2} u(x)v(x)\, dx_1 \cdots dx_n,$$

then one finds that the following polynomials form an orthonormal family: First the constant polynomial 1; then the n linear monomials $x_1, x_2, \ldots, x_n$; then the polynomials of degree 2: $p_{ij} = x_i x_j$ for $i \neq j$, etc. More generally, if $H(q, p)$ is a Hamiltonian, one can define a family of polynomials in the variables q, p that are orthonormal with respect to the canonical density $Z^{-1}e^{-H/T}$. We shall still call these polynomials "Hermite polynomials," and we shall do the same for polynomials orthonormal with respect to an inner product with a weight W.

Consider an n-dimensional space Γ with a probability density W. Divide the coordinates into two groups: $\hat{x}$ and $\tilde{x}$. Let g be a function of x; then $\mathbb{P}g = E[g|\hat{x}]$ is an orthogonal projection onto the subspace of functions of $\hat{x}$. One can perform this projection by spanning that subspace by those Hermite polynomials that are functions of $\hat{x}$ and

using them as in Section 1.2 (see also Exercise 2.2). One can then approximate the "conditional expectation" projection by a "finite-rank" projection in which one uses only a fixed set of Hermite polynomials. A particular finite-rank projection widely used in physics is the one in which only the Hermite polynomials of degree 1 are used; this is also known as the "linear" projection (as if other projections were not linear). We have already used a linear projection implicitly in the "linear" prediction method in Section 4.5.

6.4. The Mori-Zwanzig Formalism

Return now to the system

$$\frac{d\phi_i(x,t)}{dt} = R(\phi(x,t)), \quad \phi_i(x,0) = x_i, \quad 1 \le i \le n. \tag{6.18}$$

Suppose one is interested only in the first m variables $\phi_1, \ldots, \phi_m$, with $m < n$; partition the vector ϕ as in Section 5.6 into "resolved" variables $\hat{\phi}$ and "unresolved" variables $\tilde{\phi}$ so that

$$\phi = (\hat{\phi}, \tilde{\phi}), \quad \hat{\phi} = (\phi_1, \ldots, \phi_m), \quad \tilde{\phi} = (\phi_{m+1}, \ldots, \phi_n),$$

and similarly $x = (\hat{x}, \tilde{x})$ and $R = (\hat{R}, \tilde{R})$. We now look for equations for the components $\hat{\phi}(t)$ with the initial conditions $\hat{\phi}(0) = \hat{x}$. We further assume that at time $t = 0$ we know the joint pdf of all the variables x; once the initial data $\hat{x}$ are given, the pdf of the variables in $\tilde{x}$ is the joint pdf of all the x variables conditioned by $\hat{x}$. Something has to be assumed about the missing variables $\tilde{x}$ lest the problem become meaningless; the assumptions here are often realistic, but one should be careful not to use what is now coming when these assumptions do not hold.

This is the third time in this book that we are trying to make a prediction or draw a conclusion on the basis of uncertain or statistical information. In Section 4.5 we made predictions on the assumption that we had some observations for a process which we knew to be stationary and whose covariances were known. In Section 4.6 we made predictions or drew conclusions for a process for which we had a noisy model and a stream of noisy observations. Now we want to make prediction from a model which we can implement only in part, and/or for which we have only partial initial data.

Form the Liouville equation $u_t = Lu$ as above in Section 6.3; the components $\hat{\phi}$ are $\hat{\phi}_j(x,t) = e^{tL}x_j$ (note that $\hat{\phi}_j$ depends on all the data x when the system is not linear; if $\tilde{x}$ is random, $\hat{\phi}$ is random as

well). By definition,

$$\frac{\partial}{\partial t}e^{tL}x_j = Le^{tL}x_j = e^{tL}Lx_j, \tag{6.19}$$

where the last equality is the commutation rule (6.15). Let $\mathbb{P}$ be the conditional expectation projection $\mathbb{P}g(x) = E[g|\hat{x}]$, where the probability density is the assumed density for the initial conditions. We shall use the same notation even when we later approximate the conditional expectation by a finite-rank projection. Note that the $\mathbb{P}$ here is a projection on a space of functions of a fixed set of variables and is, therefore, time independent. Furthermore, objects such as $\mathbb{P}\hat{\phi}(t) = E[\hat{\phi}(t)|\hat{x}]$ are of great interest: they are the best estimates of the future values of a reduced system of variables given partial information about the present. This is the kind of thing a meteorologist, for example, wants to calculate: a best prediction of a set of interesting features of the future weather given our limited information about the present state of the atmosphere.

Define, furthermore, $\mathbb{Q} = I - \mathbb{P}$ and keep in mind that $\mathbb{P}^2 = \mathbb{P}$, $\mathbb{Q}^2 = \mathbb{Q}$, and $\mathbb{P}\mathbb{Q} = 0$, as must be true for any projection. Equation (6.19) can be rewritten as

$$\frac{\partial}{\partial t}e^{tL}x_j = e^{tL}\mathbb{P}Lx_j + e^{tL}\mathbb{Q}Lx_j. \tag{6.20}$$

Consider the first term. We have

$$Lx_j = \sum_i R_i(\partial/\partial x_i)x_j = R_j(x);$$

so $\mathbb{P}Lx_j = E[R_j(x)|\hat{x}] = \bar{R}_j(\hat{x})$ is a function of the reduced set of variables $\hat{x}$, the average of right-hand-sides of the j^{th} equation at the initial time, conditioned by the partial knowledge embodied in $\hat{x}$; $e^{tL}\mathbb{P}Lx_j = \bar{R}_j(\hat{\phi}(x,t))$ by the commutation rule (6.14). If one replaces the projection used here by the projection of Section 5.6, one finds that equation (6.20) coincides with the equation we used to renormalize in Section 5.6 provided $\mathbb{Q} = 0$; however, $\mathbb{Q} \neq 0$ unless $\hat{\phi}$ coincides with ϕ, which explains why the "equilibrium" renormalization of chapter 5 has to be reformulated here before one can deal with time-dependent statistics.

We now split the second term in (6.20) using Dyson's formula with $A = \mathbb{Q}L$ and $B = \mathbb{P}L$ (the reasons for the split will emerge soon):

$$e^{tL} = e^{t\mathbb{Q}L} + \int_0^t e^{(t-s)L}\mathbb{P}Le^{s\mathbb{Q}L}\,ds. \tag{6.21}$$

Here, the linearity of the Liouville equation is being used—this step is the reason for the introduction of that equation into the analysis.

Using (6.21), (6.20) becomes

$$\frac{\partial}{\partial t}e^{tL}x_j = e^{tL}\mathbb{P}Lx_j + e^{t\mathbb{Q}L}\mathbb{Q}Lx_j + \int_0^t e^{(t-s)L}\mathbb{P}Le^{s\mathbb{Q}L}\mathbb{Q}Lx_j\,ds. \quad (6.22)$$

This is the Mori-Zwanzig equation. This equation is exact and is an alternative way of writing the original system (6.18). It is an equation for each one of the $\phi_j(x,t) = e^{tL}x_j$, $j = 1,\dots,m$.

Now examine the different terms that appear on the right-hand side of (6.22). The first term is a function only of $\hat{\phi}(x,t)$ and represents the self-interaction of the resolved variables; it is a Markovian term, inasmuch as it is evaluated at the same time t as the left-hand side of the equation.

To decode the second term, write

$$e^{t\mathbb{Q}L}\mathbb{Q}Lx_j = w_j.$$

The function $w_j(x,t)$ satisfies, by definition, the equations

$$\begin{aligned} &\frac{\partial}{\partial t}w_j(x,t) = \mathbb{Q}Lw_j(x,t), \\ &w_j(x,0) = \mathbb{Q}Lx_j = (I-\mathbb{P})R_j(x) = R_j(x) - \bar{R}_j(\hat{x}). \end{aligned} \quad (6.23)$$

If one identifies $\mathbb{P}Lx_j$ as the "average of the initial data," then $w_j(x,0)$ is a "fluctuating part of the initial data" (according to the often used terminology in which a "fluctuating part" of a random variable η is $\eta - E[\eta]$). Obviously, $\mathbb{P}w_j(x,0) = 0$. If one took this initial function and applied the operator e^{tL} to it (i.e., solved the Liouville equation starting from this initial function), the result would, in general, have a nontrivial mean part (one not in the null space of $\mathbb{P}$); the evolution equation for w_j removes the "mean part" at each instant of time. As a result, $\mathbb{P}w_j(x,t) = 0$ for all time t.

Call the space of functions of $\hat{x}$ the "resolved subspace" and its orthogonal complement (with respect to the inner product defined by the initial density) the "noise subspace." $\mathbb{P}$ applied to any element of the noise subspace gives zero, and, similarly, $\mathbb{Q}$ applied to any element of the resolved subspace gives zero. The functions $w_j(x,t) = e^{t\mathbb{Q}L}\mathbb{Q}Lx_j$ are in the noise space; we shall call the vector of which they are the components the "noise" for short. The noise is determined by the initial data and by the system (6.18) and does not have to be white noise. Equation (6.23) is the "orthogonal dynamics" equation.

The third term in (6.22) is the "memory" term because it involves integration of quantities that depend on the state of the system at earlier times. To see what this term does, approximate the projection $\mathbb{P}$ by a finite-rank projection in terms of Hermite polynomials $(H_1,\dots,H_p)$

(whose arguments belong to $\hat{x}$). We have

$$\begin{aligned}\mathbb{P}Le^{s\mathbb{Q}L}\mathbb{Q}Lx_j &= \mathbb{P}L(\mathbb{P}+\mathbb{Q})e^{s\mathbb{Q}L}\mathbb{Q}Lx_j \\ &= \mathbb{P}L\mathbb{Q}e^{s\mathbb{Q}L}\mathbb{Q}Lx_j \\ &\sim \sum_{k=1}^{p}(L\mathbb{Q}e^{s\mathbb{Q}L}\mathbb{Q}Lx_j, H_k(\hat{x}))H_k(\hat{x}).\end{aligned}$$

To simplify the analysis, assume that L is skew-symmetric, $(u, Lv) = -(Lu, v)$; we have seen that this includes the case where the system (6.18) we started from was Hamiltonian. Then we find

$$\begin{aligned}(L\mathbb{Q}e^{s\mathbb{Q}L}\mathbb{Q}Lx_j, H_k(\hat{x})) &= -(\mathbb{Q}e^{s\mathbb{Q}L}\mathbb{Q}Lx_j, LH_k) \\ &= -(e^{s\mathbb{Q}L}\mathbb{Q}Lx_j, \mathbb{Q}LH_k).\end{aligned}$$

Both $\mathbb{Q}Lx_j$ and $\mathbb{Q}LH_k$ are in the noise subspace, and $e^{s\mathbb{Q}L}\mathbb{Q}Lx_j$ is a solution at time s of the orthogonal dynamics equation with data in the noise subspace; $\mathbb{P}Le^{s\mathbb{Q}L}\mathbb{Q}Lx_j$ is then a sum of temporal covariances of "noises" (i.e., of functions in the noise subspace). The operator $e^{(t-s)L}$ commutes with each $(L\mathbb{Q}e^{s\mathbb{Q}L}\mathbb{Q}Lx_j, H_k(\hat{x}))$ because the latter expression is an inner product that does not evolve in time, and by the rule (6.14), one finds $e^{(t-s)L}H_k(\hat{x}) = H_k(\hat{\phi}(t-s))$; if one makes the change of variables $t' = t - s$ and drops the prime, one finds that the memory integral has an integrand which is a sum of terms; each one of them is the product of a temporal covariance of a noise (i.e., a variable that lives in the null space of $\mathbb{P}$), evaluated at the time difference $(t-s)$, multiplied by a variable that depends on the state of the system at the time s. Such terms of course represent a memory. The split (6.21) has divided the non-Markovian term in the equation of motion for the $\hat{\phi}$ into a noise and a memory that depends on the temporal covariances of the noise.

One can bring in an apparent simplification by multiplying (6.22) by the projection $\mathbb{P}$; remember that $\mathbb{P}$ is time invariant, so that $\mathbb{P}(\partial/\partial t)\hat{\phi}$ becomes $(\partial/\partial t)E[\hat{\phi}|\hat{x}]$—the best estimates of the future values of the variables we are following, given the data; these quantities are often of interest. Using the fact that $\mathbb{P}$ operating on the noise term is zero, one finds

$$\frac{\partial}{\partial t}\mathbb{P}e^{tL}x_j = \mathbb{P}e^{tL}\mathbb{P}Lx_j + \int_0^t \mathbb{P}e^{(t-s)L}\mathbb{P}Le^{s\mathbb{Q}L}\mathbb{Q}Lx_j\,ds, \tag{6.24}$$

where $\mathbb{P}e^{tL}x_j = E[\hat{\phi}(x,t)|\hat{x}]$ by definition. However, the Markovian term is now more complicated: we have seen that $e^{tL}\mathbb{P}Lx_j$ is, in general,

a nonlinear function $\bar{R}(\hat{\phi}(t))$; therefore, $\mathbb{P}\bar{R}(\hat{\phi}(t))$ is, in general, not equal to $\bar{R}(\mathbb{P}\hat{\phi}(t))$ and some approximation scheme must be devised.

To make contact with the literature, one has to make some drastic simplifications. Assume that the "linear" projection will do the job (this is generally true if the processes ϕ have small amplitude). Suppose that the initial probability density W is such that $E[x_i x_j] = (x_i, x_j) = \delta_{ij}$. Assume that the noise $e^{t\mathbb{Q}L}\mathbb{Q}Lx_j$ is white noise (occasionally this is a good assumption; see the next few sections). Then the correlations that appear in the integrand of the memory term are delta functions, and in this case, the integral in the memory term becomes a simple function evaluation. With some further assumptions about the original equations (6.18), explained in the next two sections, one recovers as special cases the systems of the first two sections of this chapter. Thus equations (6.22) are general Langevin equations, generalizing what we have earlier called the Langevin equation.

The Mori-Zwanzig equations (6.22) are exact. If one has a system of equations for ϕ, a pdf for the initial data, specific initial data for $\hat{\phi}(t = 0)$ and one wants to find $\hat{\phi}(t)$, one can either sample the whole vector of initial data, solve for $\phi(t)$, and throw away all that is not $\hat{\phi}$, or one can try to solve (6.22). One can average in either case. Equations (6.22) are fewer in number, but this advantage is outweighed by the need to find the noise and its temporal covariances. What equations (6.22) do provide is a starting point for approximation.

What are such approximations needed for? There are two settings in which they can be useful:

(1) The analysis of how large mechanical systems converge to the kind of equilibrium discussed in the previous chapter. If one sets the values of some initial data to fixed values but lets the other initial values be picked at random from a canonical density, one in fact takes the mechanical system out of equilibrium at time $t = 0$. An ergodic Hamiltonian system will then see its entropy increase and it will tend toward equilibrium; it is often of interest to see how this happens and this can, in principle, be done by approximating equations (6.22).
(2) Suppose one wants to make predictions on the basis of partial data (as, for example, in weather forecasting). One can assume something reasonable about the statistics of missing information (e.g., on the basis of previous experience), and turn to (6.22). Prediction methods based on the Mori-Zwanzig formalism also go under the name "optimal prediction."

Finally, some words on the long-time behavior of the solutions of (6.22). Suppose the system (6.18) is Hamiltonian and ergodic. If the initial data are nonequilibrium data (not sampled from a canonical density, for example, some of them are given numbers $\hat{x}$), then as time unfolds, the system will approach equilibrium (i.e., the joint density of the $\phi(t)$ will approach a canonical density as the entropy increases). The averages $\mathbb{P}\phi(t) = E[\phi|\hat{x}]$ will converge to the averages of ϕ with respect to the canonical density—the predictive power of initial data decays to zero with time (for example, one can make decent one-day weather forecasts on the basis of today's observations, but very poor one-year forecasts). The solutions of the equation for $\mathbb{P}\hat{\phi}(t)$ tend to constants (often zero) independent of the data. The Markovian term in (6.22) tends to zero as well, and one is left with an equation that merely balances noise and memory.

6.5. More on Fluctuation-Dissipation Theorems

We have established a relation between kernels in the memory term and the noise (the former is made up of covariances of the latter). This is the mathematical content of some of the "fluctuation-dissipation theorems" in physics. However, under some specific restricted circumstances, the relation between noise and memory takes on more intuitively appealing forms, which we now briefly describe. As has already been mentioned, physicists often take a restricted basis in the range of P consisting of the coordinate functions $x_1, ..., x_m$ (the components of $\hat{x}$). The resulting projection is the linear projection. The use of this projection is appropriate when the amplitude of the functions $\hat{\phi}(t)$ is small. One then has $H_k(\hat{x}) = x_k$ for $k \leq m$. The correlations in (6.22) are then simply the temporal covariances of the fluctuations in the same quantities as the ones that are being solved for. This is known as the fluctuation-dissipation theorem of the second kind. If the noise can be viewed as white, the covariances are delta functions, the memory is instantaneous, and the equations look like the Langevin equation of Section 6.1.

Specialize further to a situation where there is a single resolved variable, say ϕ_1, so that $m = 1$ and $\hat{\phi}$ has a single component. The Mori-Zwanzig equation becomes

$$\frac{\partial}{\partial t}e^{tL}x_1 = e^{tL}PLx_1 + e^{tQL}QLx_1 + \int_0^t e^{(t-s)L}PLe^{sQL}QLx_1 ds,$$

or

$$\frac{\partial}{\partial t}\phi_1(x,t) = (Lx_1,x_1)\phi_1(x,t) + e^{tQL}QLx_1 + \int_0^t (LQe^{sQL}QLx_1,x_1)\phi_1(x,t-s)\,ds$$

$$= (Lx_1,x_1)\phi_1(x,t) + e^{tQL}QLx_1 - \int_0^t (e^{sQL}QLx_1,QLx_1)\phi_1(x,t-s)\,ds, \tag{6.25}$$

where we have again inserted a harmless factor Q in front of e^{QL}, assumed that L was skew-symmetric as above, and for the sake of simplicity also assumed $(x_1,x_1)=1$ (if the last statement is not true, the formulas can be adjusted by the insertion of a suitable constant). Taking the inner product of (6.25) with x_1, you find

$$\begin{aligned}\frac{\partial}{\partial t}(\phi_1(x,t),x_1) &= (Lx_1,x_1)(\phi_1(x,t),x_1) + (e^{tQL}QLx_1,x_1) \\ &\quad - \int_0^t (e^{sQL}QLx_1,QLx_1)(\phi_1(x,t-s),x_1)\,ds \\ &= (Lx_1,x_1)(\phi_1(x,t),x_1) \\ &\quad - \int_0^t (e^{sQL}QLx_1,QLx_1)(\phi_1(x,t-s),x_1)\,ds, \end{aligned} \tag{6.26}$$

because $Pe^{tQL}QLx_1 = (e^{tQL}QLx_1,x_1)x_1 = 0$ and, hence, $(e^{tQL}QLx_1,x_1)$ vanishes as well. Multiply equation (6.26) by x_1, and remember that $P\phi_1(x,t) = (\phi_1(x,t),x_1)x_1$. You find

$$\frac{\partial}{\partial t}P\phi_1(x,t) = (Lx_1,x_1)P\phi_1(x,t) - \int_0^t (e^{sQL}QLx_1,QLx_1)P\phi_1(x,t-s)\,ds. \tag{6.27}$$

You observe that the covariance $(\phi_t(x,t),x_1)$ and the projection of ϕ_1 on x_1 obey the same homogeneous linear integral equation. This is a "fluctuation-dissipation theorem of the first kind", which embodies the Onsager principle, according to which spontaneous fluctuations in a system decay at the same rate as perturbations imposed by external means when both are small (so that the linear projection is adequate). This reasoning can be extended to cases where there are multiple resolved variables, and this is often done with the added simplifying assumption that $(x_i,x_j)=0$ when $i\neq j$. We omit the details.

6.6. Scale Separation and Weak Coupling

There are situations where one knows that the noise term in the Mori-Zwanzig equations can be approximated by white noise, and then the memory term becomes local in time and everything is simpler. This happens in particular when there is scale separation between the resolved and unresolved variables. This means that there is a significant gap between the frequencies of the resolved components $\hat{\phi}$ and the frequencies of the unresolved components $\tilde{\phi}$. The heuristic reason is clear: If the resolved variables take a time Δt to vary significantly, during this time interval the unresolved variables make many uncorrelated contributions to the motion of the resolved variables, whose effect can be described by a sum of independent Gaussian variables (by the central limit theorem) and, hence, summarized as the effect of a white noise. A closely related situation is that of "weak coupling," where the variations of $\tilde{\phi}$ affect $\hat{\phi}$ by a small amount; it takes many of them to have a significant effect and their cumulative effect is that of a large number of independent contributions. The detailed description of these situations requires asymptotic solutions of a singular perturbation problems, as we illustrate by an example.

Consider a particle at a point x whose velocity v can be either $+1$ or -1; it jumps from one value to the other in every short time interval dt with a probability dt, with independent probabilities for a jump on two disjoint intervals. Let the position x of the particle be given by

$$\dot{x} = \epsilon v(t),$$

or

$$x(t) = \epsilon \int_0^t v(s)\, ds.$$

The presence of the parameter ϵ, which will soon be made small, embodies a weak coupling assumption. The variable x is analogous to a resolved variable; for simplicity, we present a model in which the unresolved, "fast," variable v is not determined by an equation but rather by fiat.

The probability density function $W(x, \pm 1, t)$ is the probability that the particle be between x and $x + dx$ while v is either $+1$ or -1. It can be thought of as a vector $W = (W^+, W^-)$, where $W^+(x, t)$ is the probability that the particle be between x and $x+dx$ with $v = +1$ with a similar definition for W^-. $W^+(x, t+\delta t)$ equals $(1-\delta t)$ (the probability that there is no change in velocity) times $W(x-\epsilon\delta t)$ (because particles moving at speed ϵ go from $x - \epsilon\delta t$ to x in a time δt), plus $\delta t W^-(x, t)$ (because of jumps from the minus state). Collecting terms, expanding

$W(x-\epsilon\delta t)$, dividing by δt, and letting $\delta t \to 0$, as in Section 3.9, yields

$$W_t^+ = -\epsilon W_x^+ + W^- - W^+,$$

and, similarly,

$$W_t^- = \epsilon W_x^- + W^+ - W^-,$$

where the subscripts x and t denote differentiation. Define

$$U = W^+ - W^-, \;\; V = W^+ + W^-;$$

one finds

$$U_t = -\epsilon V_x - 2U, \;\; V_t = -\epsilon U_x,$$

and, hence,

$$U_{tt} = \epsilon^2 U_{xx} - 2U_t.$$

Once U is found, V, W^+, and W^- follow immediately.

One does not expect, with the weak coupling when ϵ is small, to have a significant displacement x of a particle when t is of order 1. We therefore introduce a slow timescale such that, when a unit time has passed on this slower scale, one can expect a significant displacement to have occurred; we do this by setting $\tau = \epsilon^2 t$. The equation for $U = U(x,\tau)$ becomes

$$\epsilon^2 U_{\tau\tau} = U_{xx} - 2U_\tau,$$

and, in the limit $\epsilon \to 0$, we obtain $U_\tau = \frac{1}{2}U_{xx}$, a heat equation that can be solved by examining particles undergoing BMs, as promised. This is of course just a reflection of the fact that by the central limit theorem, the sum of the independent contributions to x due to the assumed velocity adds up over time to a Gaussian variable.

Similarly, one can see that a heavy particle bombarded by lighter particles undergoes a displacement that, over the proper time scales, satisfies the Langevin equation as written in the first section of this chapter. The ratio of masses provides the needed ϵ.

6.7. Long Memory and the t-Model

In the previous section we considered situations where the memory could be viewed as having zero range (i.e., where the covariances that appear in the Mori-Zwanzig identity can be viewed as delta functions). We saw that this happened when there was separation of scale, i.e., when the variables omitted were much faster than the ones resolved. We now consider the opposite situation, where the unresolved variables are in fact slow. This corresponds in particular to the situation in hydrodynamics, where persistent shocks and vortices on all scales slow

down the decay of temporal correlations. If the unresolved variables are slow one can expect the memory to be long and the integrand in the memory term to vary slowly; under these conditions the Mori-Zwanzig formula may simplify again, in a way different from what happens in the short-memory case, leading to what is known as the "t-model".

The easiest derivation of the t-model is as follows. If the integrand in the memory term varies slowly (in s, the variable of integration in that term), one can expand the integrand in powers of s and keep only the first, constant term. This leads to

$$\int_0^t e^{(t-s)L}\mathbb{P}Le^{s\mathbb{Q}L}\mathbb{Q}Lx_j\,ds \cong \int_0^t e^{tL}\mathbb{P}L\mathbb{Q}Lx_j\,ds = te^{tL}\mathbb{P}L\mathbb{Q}Lx_j. \quad (6.28)$$

The integral has disappeared, all that remains is the coefficient t. The origin of the time t is of course the time where the partial data were known; we have always assumed that such a privileged time existed. Note also that higher order approximations, all without an integral term, can be obtained by keeping terms beyond the first in the expansion of the memory integrand. Another derivation of the same result starts from the approximation

$$e^{t\mathbb{Q}L} \cong e^{tL}, \quad (6.29)$$

in the memory term. The rationale for this approximation is that, if the "noise" varies slowly, it cannot be much influenced on a short time scale by the variation of the resolved variables; if that is the case $e^{t\mathbb{Q}L} \cong e^{tL}$ when applied to vectors in the noise space. A little algebra then recovers the expression (6.28) for the memory term.

The replacement of the memory integral by a Markovian term is a major simplification, but not the end of all difficulties. In particular, one still has to derive and sample the noise term, with or without the assumption (6.29). We now make the further assumption that the noise is small enough so that, for reasonable functions $F(\phi)$, one has approximately

$$\mathbb{P}F(\phi) = F(\mathbb{P}\phi). \quad (6.30)$$

(remember that $\mathbb{P}$ denotes averaging with respect to partial initial data.) This assumption has to be checked in each application of the resulting formulas; in physics this kind of approximation is called a "mean-field approximation".

With these assumptions, and after multiplication by $\mathbb{P}$, the Mori-Zwanzig equations reduce to:

$$\frac{d}{dt}\mathbb{P}e^{tL}\hat{x} = \mathbb{P}e^{tL}\mathbb{P}L\hat{x} + t\mathbb{P}e^{tL}\mathbb{P}L\mathbb{Q}L\hat{x}. \quad (6.31)$$

This is the "t-model", which is deterministic (no random term) and Markovian.

As an example, consider a model consisting of two linear oscillators with a nonlinear coupling, whose Hamiltonian is

$$H(\phi) = \frac{1}{2}\left(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2 + \phi_1^2\phi_3^2\right),$$

with (ϕ_1, ϕ_2) and (ϕ_3, ϕ_4) canonical pairs of coordinates. We encountered this model in exercises with the notations $\phi_1 = q_1, \phi_2 = p_1, \phi_3 = q_2, \phi_4 = p_2$. The resulting equations of motion are

$$\begin{aligned} \frac{d}{dt}\phi_1 &= \phi_2, \quad \frac{d}{dt}\phi_2 = -\phi_1(1+\phi_3^2), \\ \frac{d}{dt}\phi_3 &= \phi_4, \quad \frac{d}{dt}\phi_4 = -\phi_3(1+\phi_1^2). \end{aligned}$$

Suppose one wants to solve only for $\hat{\phi} = (\phi_1, \phi_2)$, with initial data $\hat{x} = (x_1, x_2)$. Assume the initial data (x_3, x_4) are sampled from a canonical density with temperature $T = 1$. A quick calculation yields

$$E[x_3^2 | x_1, x_2] = \frac{1}{1+x_1^2}.$$

The advance in time described by the multiplication by e^{tL} requires just the substitution $\hat{x} \to \hat{\phi}$. If one commutes the nonlinear function evaluation and the conditional averaging as suggested above, and writes,

$$\Phi(t) = \mathbb{P}\hat{\phi} = E[\hat{\phi}|\hat{x}],$$

one finds

$$\begin{aligned} &\mathbb{P}e^{tL}\mathbb{P}Lx_1 = \mathbb{P}e^{tL}x_2 = \Phi_2, \\ &\mathbb{P}e^{tL}\mathbb{P}Lx_2 = \mathbb{P}e^{tL}\bar{R}_2(\hat{x}) = \mathbb{P}\bar{R}_2(\hat{\phi}(\hat{x},t)) \sim -\Phi_1\left(1 + \frac{1}{1+\Phi_1^2}\right), \end{aligned}$$

where $\bar{R}(\hat{x}) = E[R(x)|\hat{x}]$, so that

$$\begin{aligned} &\frac{d}{dt}\Phi_1 = \Phi_2, \\ &\frac{d}{dt}\Phi_2 = -\Phi_1\left(1 + \frac{1}{1+\Phi_1^2}\right) - 2t\frac{\Phi_1^2\Phi_2}{(1+\Phi_1^2)^2}. \end{aligned} \tag{6.32}$$

For the sake of completeness, we present intermediate steps in this calculation for the second equation:

$$\begin{aligned}
Lx_2 &= -x_1(1+x_3^2),\\
\mathbb{P}Lx_2 &= -x_1\left(1+\frac{1}{1+x_1^2}\right),\\
\mathbb{Q}Lx_2 &= -x_1(1+x_3^2)+x_1\left(1+\frac{1}{1+x_1^2}\right),\\
L\mathbb{Q}Lx_2 &= x_2\left[-(1+x_3^2)+\left(1+\frac{1}{1+x_1^2}\right)+\frac{2x_1^2}{(1+x_1^2)^2}\right]-2x_1x_3x_4,\\
\mathbb{P}L\mathbb{Q}Lx_2 &= -\frac{2x_1^2x_2}{(1+x_1^2)^2}.
\end{aligned}$$

The last term in (6.32) represents the damping due to the loss of predictive power of partial data; the coefficient of the last term increases in time and one may worry that this last term eventually overpowers the equations and leads to some odd behavior. This does not happen. Indeed, one can prove the following: If the system (6.18) one starts from is Hamiltonian with Hamiltonian H, if the initial data are sampled from an initial canonical density conditioned by partial data $\hat{x}$, and if $\widehat{H}$ is the "renormalized' Hamiltonian $\hat{H} = -T\log\int e^{-H(\phi_1,\phi_2,\phi_3,\phi_4)/T}d\phi_3 d\phi_4$ (see problem 5.9), then with the t-model $(d/dt)\widehat{H} \le 0$, showing that the components of $\hat{\phi}$ decay as they should. The proof requires a minor technical assumption (that the Hamiltonian H can be written as the sum of a function of p and a function of q, a condition commonly satisfied) and we omit it.

There are other, more sophisticated and more complicated, ways of improving predictions when memory is short, long, or neither. Their analysis is beyond the scope of this volume.

6.8. Exercises

1. Find the Fokker-Planck equation that corresponds to the Langevin equations as it is written in Section 6.1, and then show that the canonical density is a stationary solution of the Fokker-Planck equation only if the fluctuation-dissipation relation holds.

2. Derive the Fokker-Planck equation at the end of Section 6.1, and then derive the fluctuation-dissipation relation as sketched in the text.

3. Check the relationships

$$E\left[\left(q_j(0) - \frac{\gamma_j x(0)}{f_j^2}\right)\left(q_k(0) - \frac{\gamma_k x(0)}{f_k^2}\right)\right] = \frac{\delta_{jk} T}{f_j^2},$$

$$E[p_j(0)p_k(0)] = T\delta_{jk},$$

$$E[F_n(t)F_n(t-t')] = -TK_n(t')$$

at the end of Section 6.2.

4. Consider the ordinary differential equation $d\phi(x,t)/dt = 1, \phi(x,0) = x$, construct the corresponding Liouville equation, solve this Liouville equation explicitly when the initial datum is $u(x,0) = x$, and verify that $u(x,t) = \phi(x,t)$. Suppose the density of the initial values x is $f(x)$. Find the density $f(\phi,t)$ of $\phi(x,t)$ at time t, and verify the remark at the end of Section 6.3.1.

5. Check the derivation of the approximation of the model in Section 6.7.

6. For the previous problem, determine $\Phi_{1,2}(t)$ for $0 \leq t \leq 2$ in the following two ways: (i) approximately by solving the approximating two-by-two system of ordinary differential equations, and (ii) by repeatedly sampling x_3, x_4 given x_1, x_2, solving the four-by-four system and averaging. Compare the two solutions.

6.9. Bibliography

[1] B. Alder and T. Wainwright, Decay of the velocity correlation function, Phys. Rev. A 1 (1970), pp. 1-12.

[2] R. Balescu, *Statistical Dynamics, Matter out of Equilibrium*, Imperial College Press, London, 1997.

[3] S. Chandrasekhar, Stochastic problems in physics and astronomy, Rev. Mod. Phys. 15 (1943), pp. 1-88; reprinted in N. Wax, *Selected Papers on Noise and Stochastic Processes*, Dover, New York, 1954.

[4] A.J. Chorin, O.H. Hald, and R. Kupferman, Optimal prediction and the Mori-Zwanzig representation of irreversible processes, Proc. Natl. Acad. Sci. USA 97 (2000), pp. 2968–2973.

[5] A.J. Chorin, O.H. Hald, and R. Kupferman, Optimal prediction with memory, Physica D, 166 (2002), pp. 239-257.

[6] A.J. Chorin and P. Stinis, Problem reduction, renormalization, and memory, Comm. Appl. Math. Comp. Sci. 1 (2005), pp. 1-27.
[7] W. E and B. Engquist, The heterogeneous multiscale methods, Comm. Math. Sci. 1 (2003), pp. 87-133.
[8] D. Evans and G. Morriss, *Statistical Mechanics of Nonequilibrium Liquids*, Academic, New York, 1990.
[9] G. Ford, M. Kac, and P. Mazur, Statistical mechanics of assemblies of coupled oscillators, J. Math. Phys. 6 (1965), pp. 504-515.
[10] D. Givon, R. Kupferman, and A. Stuart, Extracting macroscopic dynamics, model problems and algorithms, Nonlinearity 17 (2004), pp. R55-R127.
[11] S. Goldstein, On diffusion by discontinuous movements and on the telegraph equation, Q. J. Mech. Appl. Math. 4 (1951), pp. 129-156.
[12] O.H. Hald and P. Stinis, Optimal prediction and the rate of decay for solutions of the Euler equations in two and three dimensions, Proc. Natl. Acad. Sci. USA 104 (2007), pp. 6527-6532.
[13] P. Hohenberg and B. Halperin, Theory of dynamical critical phenomena, Rev. Mod. Phys. 49 (1977) pp. 435-479.
[14] I. Horenko and C. Schutte, Likelihood-based estimation of multidimensional Langevin models and its application to biomolecular dynamics, Multiscale Model. Simul. 7 (2008), pp. 731-773.
[15] M. Kac, A stochastic model related to the telegrapher's equation, Rocky Mountain J. Math. 4 (1974), pp. 497-509.
[16] S.S. Ma, *Modern Theory of Critical Phenomena*, Benjamin, Boston, 1976.
[17] A. Majda and X. Wang, *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*, Cambridge University Press, 2006.
[18] A. Majda, I. Timofeyev, and E. Vanden-Eijnden, A mathematical framework for stochastic climate models, Commun. Pure Appl. Math. 54 (2001), pp. 891-947.
[19] G. Papanicolaou, Introduction to the asymptotic analysis of stochastic equations, in *Modern Modeling of Continuum Phenomena*, R. DiPrima (ed.), Providence RI, 1974.
[20] C. Schutte, J. Walters, C. Hartmann, and W. Huisinga, An averaging principle for fast degrees of freedom exhibiting long-term correlations, Multiscale Model. Simul. 2 (2004), pp. 501-526.
[21] P. Stinis, Stochastic optimal prediction for the Kuramoto-Sivashinksi equation, Multiscale Model. Simul. 2 (2004), pp. 580-612.

[22] K. Theodoropoulos, Y.H. Qian and I. Kevrekidis, Coarse stability and bifurcation analysis using timesteppers: a reaction diffusion example, Proc. Natl. Acad. Sci. USA 97 (2000), pp. 9840-9843.

[23] R. Zwanzig, Problems in nonlinear transport theory, in *Systems Far from Equilibrium*, L. Garrido (ed.), Springer-Verlag, New York, 1980.

[24] R. Zwanzig, Nonlinear generalized Langevin equations, J. Statist. Phys. 9 (1973), pp. 423-450.

Index

anomalous exponent, 90, 131

Bayes' theorem, 41–43, 103
Box-Muller algorithm, 29, 44
branching Brownian motion, 73, 74
Brownian motion (BM), 47–50, 54, 55, 62, 66, 68, 69, 72–74, 76–79, 85, 91, 96, 103, 135, 137

canonical density, 115, 117–119, 121, 144, 149, 150, 155–157
central limit theorem, 34, 36, 37, 53, 54, 129, 152, 153
Chapman-Kolmogorov equation, 67, 68, 70
conditional expectation, 37, 39, 102, 145, 146
conditional probability, 37, 41, 42
correlation length, 120, 125, 128–130
covariance, 25, 45, 83, 85–88, 91, 95–97, 100, 102, 104–106, 120, 125, 128, 145, 148–151, 153
critical exponent, 121, 130
critical point, 126, 129, 130
critical slowing down, 126

decimation, 129
dimensional analysis, 89, 90, 131

energy shell, 114
ensemble, 112, 114, 117, 119
entropy, 115–118, 120, 132, 133, 149, 150
ergodic, 119, 123–125, 140, 149, 150

Feynman diagram, 59, 62, 64, 75
Feynman-Kac formula, 59, 61, 132
filter, 101, 102, 104
fluctuation-dissipation theorem, 136, 137, 140, 150, 151
Fokker-Planck equation, 68, 69, 72, 73, 79, 114, 126, 133, 136, 137, 156, 157

Glimm's method, 78
Gram-Schmidt process, 8

Hald model, 133, 155–157
harmonic oscillator, 111, 112, 137, 138
Hermite polynomial, 44, 144, 145, 147

incomplete similarity, 90, 131
Ising model, 119–122, 125, 126, 128, 131–133

Khinchin's theorem, 87, 95
Kolmogorov equation, 69
Kolmogorov-Obukhov scaling law, 90

Lagrangian, 65, 109–112
Langevin equation, 68, 69, 72, 75, 135, 136, 139, 140, 149, 150, 153, 156
least squares, 1, 5, 39, 40
linear prediction, 98, 145
Liouville equation, 113, 133, 141, 143, 145–147, 157
long memory, 153, 154, 156

macrostate, 112, 113
marginalization, 127–129
Markov chain, 121–125, 128, 130, 132, 133
maximum likelihood, 33, 44
memory term, 139, 140, 147, 149, 150, 152, 154
metropolis sampling, 122
microcanonical density, 114, 117, 119, 133

microstate, 112–114, 116, 120–126, 132, 133
Monte Carlo, 28–32, 44, 79, 104, 120–122, 125, 126, 130, 132, 133
Mori-Zwanzig, 145, 147, 149, 150, 152–155

optimal prediction, 149
order parameter, 121
Ornstein-Uhlenbeck, 47, 69
orthogonal dynamics, 147, 148

path integral, 59, 65, 110
phase space, 112, 114, 120, 121

QR decomposition, 9

random choice, 78
rejection sampling, 122
renormalization, 90, 126, 127, 129, 130, 140, 146

scale separation, 152
Schroedinger equation, 56, 132
semigroup notation, 74, 143, 144
short memory, 154, 156
similarity, 89, 90, 131
stationary stochastic process, 83–87, 96, 105, 106
stochastic differential equation, 69, 72, 73, 75, 77, 79, 101
stochastic Fourier transform, 91, 95, 96, 105

t-model, 153–156
transition kernel, 67
turbulence, 88, 90, 91

vorticity, 72, 73

weak coupling, 152, 153
white noise, 48, 68, 85, 87, 104, 136, 137, 140, 147, 149, 150, 152
Wick's theorem, 106
Wiener integral, 62, 63, 78, 79, 133
Wiener measure, 54–56, 63, 65
Wiener process, 47