



UNIVERSITY OF CAMBRIDGE

Faculty of Mathematics

MATHEMATICAL TRIPPOS

2025-26

GUIDE TO COURSES IN PART II

This booklet contains informal and non-technical descriptions of courses to be examined in Part II in the academic year 2025-26, as well as a summary of the overall structure of Part II.

This *Guide to Courses* is intended to supplement the more formal descriptions contained in the booklet *Schedules of Lecture Courses and Form of Examinations*.

These and other Faculty documents for students taking the Mathematical Tripos are available from the undergraduate pages on the Faculty's website at <https://www.maths.cam.ac.uk/undergrad/>

1 Introduction

Each lecture course in the Mathematical Tripos has an official syllabus, or *schedule*, that sets out formally, and in technical terms, the material to be covered. The *schedules* are listed in the booklet *Schedules of Lecture Courses and Form of Examinations* that is available for download at <https://www.maths.cam.ac.uk/undergrad/course/schedules.pdf>. The *Schedules* booklet is the definitive reference for matters of course content and assessment, for students, lecturers and examiners.

The present guide, by contrast, provides an *informal* description of each lecture course in Part II. These descriptions are intended to be comprehensible without much prior knowledge, and to convey something of the flavour of each course, and to suggest some preparatory reading, if appropriate.

A summary of the overall structure of Part II is also given below, including the distribution of questions on examination papers.

Changes to lecture courses since last year

The wording of the schedules of IIC Statistical Modelling, IIC Coding & Cryptography and IID Principles of Statistics have changed for 2025-26.

2 Structure of Part II

The structure of Part II may be summarised as follows:

- There are two types of lecture courses, labelled C and D. C-courses are all 24 lectures, D-courses may be 16 or 24 lectures. This year there are 10 C-courses and 27 D-courses. There is in addition a Computational Projects course (CATAM).
- C-courses are intended to be straightforward, whereas D-courses are intended to be more challenging.
- There is no restriction on the number or type of courses you may present for examination.
- The examination consists of four papers, with questions on the courses spread as evenly as possible over the four papers subject to:
 - each C-course having four Section I ('short') questions and two Section II ('long') questions;
 - each 24-lecture D-course having no Section I questions and four Section II questions;
 - each 16-lecture D-course having no Section I questions and three Section II questions.
- Only six questions from Section I may be attempted on each paper.
- Each Section I question is marked out of 10 with one beta quality mark, while each Section II question is marked out of 20 with one quality mark, alpha or beta. Thus each C-course and 24-lecture D-course carries 80 marks and a number of quality marks, while each 16-lecture D-course carries 60 marks and a number of quality marks. The Computational Projects course carries 150 marks and no quality marks.

3 Distribution of Questions on the Examination Papers

The distribution of Section II ('long') questions on the examination papers is as follows:

C-Courses	Paper 1	Paper 2	Paper 3	Paper 4
Number Theory			*	*
Topics in Analysis		*		*
Coding and Cryptography	*	*		
Automata and Formal Languages	*		*	
Statistical Modelling	*			*
Mathematical Biology			*	*
Further Complex Methods	*	*		
Classical Dynamics		*		*
Cosmology	*		*	
Quantum Information and Computation		*	*	
D-Courses	Paper 1	Paper 2	Paper 3	Paper 4
Logic and Set Theory	*	*	*	*
Graph Theory	*	*	*	*
Galois Theory	*	*	*	*
Representation Theory	*	*	*	*
Number Fields	*	*		*
Algebraic Topology	*	*	*	*
Linear Analysis	*	*	*	*
Analysis of Functions	*	*	*	*
Riemann Surfaces	*	*	*	
Algebraic Geometry	*	*	*	*
Differential Geometry	*	*	*	*
Probability and Measure	*	*	*	*
Applied Probability	*	*	*	*
Principles of Statistics	*	*	*	*
Stochastic Financial Models	*	*	*	*
Mathematics of Machine Learning	*	*		*
Asymptotic Methods		*	*	*
Dynamical Systems	*	*	*	*
Integrable Systems	*	*	*	
Principles of Quantum Mechanics	*	*	*	*
Applications of Quantum Mechanics	*	*	*	*
Statistical Physics	*	*	*	*
Electrodynamics	*		*	*
General Relativity	*	*	*	*
Waves	*	*	*	*
Fluid Dynamics	*	*	*	*
Numerical Analysis	*	*	*	*

4 Informal Description of Courses

C-Courses

Number Theory

Michaelmas, 24 lectures

Number Theory is one of the oldest subjects in mathematics and contains some of the most beautiful results. This course introduces some of these beautiful results, such as a proof of Gauss's Law of Quadratic Reciprocity, and a proof that continued fractions give rise to excellent approximations by rational numbers. The new RSA public codes familiar from Part IA Numbers and Sets have created new interest in the subject of factorisation and primality testing. This course contains results old and new on the problems.

On the whole, the methods used are developed from scratch. You can get a better idea of the flavour of the course by browsing Davenport *The Higher Arithmetic* CUP, Hardy and Wright *An introduction to the theory of numbers* (OUP, 1979) or the excellent *Elementary Number Theory* by G A and J M Jones. (Springer 1998).

Automata and Formal Languages

Michaelmas, 24 lectures

The lecture course deals with three basic ideas: the ideas of computability and decidability and how what is computable or decidable depends on the model of computation. The notion of computability is one of the fundamental concepts of modern science, independent of the concrete technology that the current generation of computers uses. It is directly related to Turing's famous limitative theorem that established that the power of computation based on any given model of computation is limited.

In the lecture course, we discuss automata and register machines as models of computation, giving rise to three different levels of computational power:

1. regular languages, recognised by finite automata,
2. context-free languages, recognised by push-down automata,
3. computably enumerable languages, recognised by register machines.

Prerequisites are in IA Numbers and Sets, including the notions of set, function, relation, product, partial order, and equivalence relation.

Coding and Cryptography

Lent, 24 lectures

When we transmit any sort of message errors will occur. Coding theory provides mathematical techniques for ensuring that the message can still be read correctly. Since World War II it has been realised that the theory is closely linked to cryptography – that is to techniques intended to keep messages secret. This course will be a gently paced introduction to these two commercially important subjects concentrating mainly on coding theory.

Discrete probability theory enters the course as a way of modelling both message sources and (noisy) communication channels. It is also used to prove the existence of good codes. In contrast the construction of explicit codes and cryptosystems relies on techniques from algebra. Some of the algebra should already be familiar – Euclid's Algorithm, modular arithmetic, polynomials and so on – but there are no essential prerequisites. IB Linear Algebra would be useful. IB Groups, Rings and Modules is very useful.

The book by Welsh recommended in the schedules (*Codes and Cryptography*, OUP), although it contains more than is in the course, is a good read.

Topics in Analysis

Lent, 24 lectures

Some students find the basic courses in Analysis in the first two years difficult and unattractive. This is a pity because there are some delightful ideas and beautiful results to be found in relatively elementary Analysis. This course represents an opportunity to learn about some of these. There are no formal prerequisites: concepts from earlier courses will be explained again in detail when and where they are needed. Those who have not hitherto enjoyed Analysis should find this course an agreeable revelation.

Statistical Modelling

Michaelmas, 24 lectures

This course is complementary to Part II Principles of Statistics, but takes a more applied perspective. By embedding statistical theory with real datasets and practical problems, students will learn to think like a statistician and gain confidence in solving real world problems with statistics.

There will be approximately 16 hours of lectures and eight hours of practical classes. The bulk of the lectures will focus on classical statistical models including linear and generalised linear models, which provide a powerful and flexible framework for the study of the relationship between a response (e.g. alcohol consumption) and one or more explanatory variables (age, sex etc.). We will also spend a few lectures on alternative perspectives of statistical modelling that focus on predictive performance or causality. In the practical classes, we will learn how to implement the techniques and ideas covered in the lectures by analysing several real data sets. We will be making extensive use of the statistical computer programming language **R**, which can be downloaded free of charge and for a variety of platforms from <https://cran.r-project.org/>. Most students will find it useful to write and execute their code in **RStudio**, an integrative development environment for **R** that can be downloaded from <https://posit.co/products/open-source/rstudio>.

This course should appeal to a broad range of students, including those considering further research in any aspect of Statistics and those considering careers in data-intensive industries (finance, manufacturing, healthcare, etc.). Those interested might like to try downloading **R** and experimenting with one of the excellent tutorials in

- <https://cran.r-project.org/doc/manuals/r-release/R-intro.html> (maintained by **R** Core Team)
- <https://rforcats.net/>
- <https://www.dpmms.cam.ac.uk/~pmea/>

Cosmology

Michaelmas, 24 lectures

This course presents a mathematically rigorous description of 13.8 billion years of history, from the Big Bang to the present day and beyond. The course starts by deriving the equations which describe an expanding universe. The need to include a number of surprising and mysterious ingredients, such as dark energy, dark matter, and inflation, will be discussed. Subsequently, the course will introduce the mathematics necessary to understand the first few minutes after the Big Bang, when the universe was very hot and the elements were forged. The course ends by explaining how small perturbations in the early universe subsequently grew into the glorious galaxies and structures that we see today. You will need to be comfortable with Newtonian dynamics, special relativity and some basic facts about quantum mechanics. No knowledge of astrophysics or general relativity is needed.

Classical Dynamics

Michaelmas, 24 lectures

This course follows on from the dynamics sections of Part IA Dynamics and Relativity and also uses the Euler–Lagrange equations from Part IB Variational Principles. The laws of motion for systems of particles and for rigid bodies are derived from a Lagrangian (giving Lagrange’s equations) and from a Hamiltonian (giving Hamilton’s equations) and are applied, for example, to the axisymmetric top.

One advantage of the formalism is the use of generalised coordinates; it is much easier to find the kinetic and potential energy in coordinates adapted to the problem and then use Lagrange’s equations than to

work out the equations of motion directly in the new coordinates. At a deeper level, the formalism gives rise to conserved quantities (generalisations of energy and angular momentum), and leads (via Poisson brackets) to a system which can be used as a basis for quantization.

The material in this course will be of interest to anyone planning to specialise in the applied courses. It is not used directly in any of the courses but an understanding of the subject is fundamental to Theoretical Physics.

Mathematical Biology

Lent, 24 lectures

The aim of the course is to explain from a mathematical point of view some underlying principles of biology, ranging from biochemistry and gene regulation to population dynamics and spread of infectious disease. In particular we examine mechanisms for feedback control, sensitivity amplification, oscillations, developmental instabilities, pattern-formation, competitive growth, and predator-prey interactions.

The material should be of interest to anyone who is fascinated by the richness of biological dynamics, but has been discouraged by too detail-oriented biological explanations. Mathematical methods include basic stochastic theory, nonlinear dynamics, differential equations, and numerical analysis. The concepts and techniques are not very difficult, and intuitive guiding principles and illustrative examples will be favoured over rigorous proofs. This is an exciting field with large unexplored territories for applied mathematicians.

While this course fits well with Dynamical Systems and there are places where understanding in one course will help the other, it is not essential and this course does not rely on any other Part II course, assuming concepts only from Part IA (Differential Equations and Probability) and IB (Methods).

Quantum Information and Computation

Lent, 24 lectures

Quantum processes can provide enormous benefits for information processing, communication and security, offering novel features beyond the possibilities of standard (classical) paradigms. These benefits include (i) new kinds of algorithms (so-called quantum algorithms) providing an exponentially faster method for some computational tasks, (ii) new modes of communication such as quantum teleportation, and (iii) the possibility of unconditionally secure communication in quantum cryptography. Most of these exciting developments have occurred in just the past few decades and they underpin striking applications of quantum technologies that are currently being developed.

This course will provide an introduction to these topics. No previous contact with the theory of computation or information will be assumed. IB Quantum Mechanics is essential, but only to provide prior exposure to basic ideas. This course rests on quantum theory in just a finite-dimensional setting, so the principal mathematical ingredients (from finite-dimensional linear algebra) will be readily accessible. We will begin by expounding the postulates of quantum mechanics in this setting (using Dirac notation) and then immediately make connections to information (quantum states viewed as information carriers, quantum teleportation) and computation (notion of qubits and quantum gates). Then we will discuss quantum cryptography (quantum key distribution), and quantum computing, culminating in an exposition of principal quantum algorithms, including the Deutsch–Jozsa algorithm, Grover’s searching algorithm and an overview of Shor’s quantum factoring algorithm. The course is cross-disciplinary in its conceptual ingredients and will be of interest to pure and applied mathematicians alike.

Further Complex Methods

Lent, 24 lectures

This course is a continuation in both style and content of Part IB Complex Methods, which is the only prerequisite. It will appeal to anyone who enjoyed that course. The material is classical — much of it can be found in Whittaker and Watson’s ‘Modern Analysis’, written in 1912. The passage of time has not diminished the beauty of material, though the Faculty Board decided against naming the course ‘Modern Analysis’.

The course starts with revision of Complex Methods and continues with a discussion of the process of analytic continuation, which is at the heart of all modern treatments of complex variable theory. There follows a section on special functions, including the Gamma function (which is basically the factorial function when looked at on the real line, but on the complex plane it really blossoms) and the Riemann zeta and its connection with number theory. Then the theory of series solutions of differential questions in the complex plane is developed, and suddenly the treatment given in Part IA Differential Equations makes sense. Naturally, the messy business of actually solving specific equations by series is not in the style of the course. Particularly important are those equations that have exactly three singular points, all regular. This leads to a study of the properties of the delightful hypergeometric function, of which almost every other function you know can be thought of as a special case. This is the high point of the course, involving nearly all the theory that has preceded it.

D-Courses

Graph Theory

Michaelmas, 24 lectures

Discrete mathematics is commonplace in modern mathematics, both in theory and in practice. This course provides an introduction to working with discrete structures by concentrating on the most accessible examples, namely graphs. After a discussion of basic notions such as connectivity (Menger's theorem) and matchings (Hall's marriage theorem), the course develops in more detail the theory of extremal graphs, ideas of graph colouring, and the beautiful theorem of Ramsey. A significant feature is the introduction of probabilistic methods for tackling discrete problems, an approach which is of great importance in the modern theory.

There are no formal prerequisites but it will be helpful to recall some of the elementary definitions from the Part IA Probability course. The attractions and drawbacks of Graph Theory are similar to those of that course and of the Part IA Numbers and Sets course; whilst the notions are not conceptually difficult, the problems might on occasion require you to think a little.

The text *Modern Graph Theory* by Bollobás is an excellent source and contains more than is needed for the course. For a lighter introduction try Wilson's *Introduction to Graph Theory*, or for a little more look at Bondy and Murty's old but now online *Graph Theory with Applications*.

Representation theory

Michaelmas, 24 lectures

This course, suitable for pure and applied mathematicians, is an introduction to the basic theory of linear (matrix) actions of finite groups on vector spaces. The key notion we define is the character of a linear representation: this is a function on conjugacy classes of the group which determines the representation uniquely. Orthogonality relations between characters lead to a convenient and efficient calculus with representations, once the basic character table of the group has been computed. Later in the course 'finite' is replaced by 'compact' generalising the results with little extra effort.

The Linear Algebra course is essential and Groups, Rings and Modules is helpful.

Algebraic Topology

Michaelmas, 24 lectures

Topology is the abstract study of continuity: the basic objects of study are metric and topological spaces, and the continuous maps between them (concepts which were introduced in Part IB Analysis and Topology). One important difference between topology and algebra is that in constructing continuous maps one has vastly more freedom than in constructing algebraic homomorphisms; thus problems which involve proving the *non-existence* of continuous maps with particular properties (e.g. the problem of showing that \mathbb{R}^m and \mathbb{R}^n are not homeomorphic unless $m = n$) are hard to solve using purely topological methods. The technique that has proved most successful in tackling such problems is that of developing *algebraic invariants*, which assign to every topological space (in a suitable class) an algebraic structure such as a group or vector space, and to every continuous map a homomorphism of the appropriate kind. Thus questions of the non-existence of continuous maps are reduced to questions of non-existence of homomorphisms, which are easier to solve.

Two particular algebraic invariants are studied in this course: the fundamental group, and the simplicial homology groups. Of these, the former is easier to define, but hard to calculate except in a few particular cases; the latter requires the erection of a considerable amount of machinery before it can even be defined, but once this is done it becomes relatively easy to calculate. The course concludes with a classic example of the application of simplicial homology: the classification of all compact 2-manifolds up to homeomorphism.

Apart from Part IB Analysis and Topology, the only prerequisite is a modicum of geometrical intuition. Concepts and techniques of algebraic topology are used almost everywhere in mathematics where topological spaces occur; they also arise in many research areas in mathematical physics, as well as pure mathematics.

For introductory reading, browse *Basic Topology* by M.A. Armstrong (Springer-Verlag).

Galois Theory

Michaelmas, 24 lectures

The most famous application of Galois theory – discussed at the end of this course – is the proof that the general quintic equation with rational coefficients cannot be solved by radicals. Apart from this, Galois theory plays an indispensable role in algebraic number theory and several other areas of pure mathematics. It is a subject which (in favourable circumstances) allows one to handle given polynomials elegantly and with a minimum of algebraic manipulation.

Familiarity with the material concerning field extensions and the polynomial ring $K[t]$ from Part IB Groups, Rings and Modules is essential, while Part IB Linear Algebra is useful. The most closely related Part II courses are Representation Theory and Number Fields. The book *Galois Theory* by I. Stewart (Chapman and Hall, 1989) gives a very readable introduction to the subject.

Linear Analysis

Michaelmas, 24 lectures

Functional Analysis provides the framework, and a great deal of machinery, for much of modern mathematics: not only for pure mathematics (such as harmonic analysis and complex analysis) but also for the applications of mathematics, such as probability theory, the ordinary and partial differential equations met in applied mathematics, and the mathematical formulation of quantum mechanics.

The basic idea of Functional Analysis is to represent functions as points in an infinite-dimensional vector space. Since the space is infinite-dimensional, algebraic arguments are not enough, and it is necessary and appropriate to introduce the idea of convergence by a norm, which in turn defines a metric on the space.

In this course, most attention is paid to two sorts of spaces. The first consists of spaces of continuous functions: here the appropriate convergence is uniform convergence. The second is Hilbert space (particularly important in Quantum Mechanics) which provides an infinite-dimensional analogue of Euclidean space, and in which geometrical ideas and intuitions are used.

Riemann Surfaces

Lent, 16 lectures

A Riemann surface is the most general abstract surface on which one can define the notion of an analytic function, and hence study complex analysis. Roughly speaking, a surface is made into a Riemann surface when the change from one local coordinate system to another system is given by an analytic function. Not every Riemann surface has a global coordinate system; this accounts for both the interesting and the difficult parts of the theory.

The course begins with a study of the Riemann sphere (which is just the complex plane with infinity attached) and of elliptic functions (that is to say, doubly periodic analytic functions) which are the analytic functions defined on a torus. Abstract Riemann surfaces and holomorphic maps are then introduced and some of the results already studied in earlier courses on complex analysis are extended to this more general context.

Another view on Riemann surfaces comes from Riemann's original idea that the so-called 'multivalued functions' are just considered on a wrong domain: the natural domain is a surface covering the complex plane several (possibly infinitely many) times. This surface is called the Riemann surface of an analytic function and is obtained by the process of analytic continuation, extending the function (while keeping it analytic) in a maximal way from a domain in \mathbb{C} .

The last part of the course shows that most Riemann surfaces carry their own intrinsic non-Euclidean geometry; thus complex analysis is much more closely connected to non-Euclidean geometry than to Euclidean geometry (despite the fact that it is first studied in the Euclidean plane).

Prerequisite for this course is IB Complex Analysis (some knowledge of IB Analysis and Topology will also be useful, especially for elliptic functions). Related Part II courses include those on Algebraic Topology, Algebraic Geometry and Differential Geometry. As a preliminary reading, consider the early parts of G.A. Jones and D. Singerman, *Complex functions* CUP, 1987, and of A.F. Beardon, *A primer on Riemann surfaces* CUP, 1984.

Algebraic Geometry

Lent, 24 lectures

Algebraic geometry is a branch of mathematics which, as the name suggests, combines techniques of abstract algebra, especially commutative algebra, with the language and the problematics of geometry. It occupies a central place in modern mathematics and has multiple conceptual connections with such diverse fields as complex analysis, topology and number theory. Initially a study of polynomial equations in many variables, the subject of algebraic geometry starts where equation solving leaves off, and it becomes at least as important to understand the totality of solutions of a system of equations, as to find some solution; this does lead into some of the deepest waters in the whole of mathematics, both conceptually and in terms of technique.

This course is an introduction to the basic ideas of algebraic geometry (affine and projective spaces, varieties), followed by a more detailed study of algebraic curves. We will develop the basic tools for understanding the properties of algebraic curves, and apply these at the end of the course to the beautiful theory of elliptic curves, which among other things played an essential part in the proof of Fermat's Last Theorem! You will find it highly advantageous to have attended the Part IB course Groups, Rings and Modules. Part II courses with which this course is related include Galois Theory, Differential Geometry, and Algebraic Topology. Students wishing to do some preliminary reading could browse the books of Reid or Kirwan noted in the schedules.

Differential Geometry

Lent, 24 lectures

A manifold is a space that looks locally like euclidean space. The surface of a sphere or the surface of a torus are natural examples, but manifolds often arise indirectly, for example, as the space of solutions of some set of conditions, the parameter space of a family of mathematical objects, configuration spaces in mechanical systems and so on. Manifolds provide the appropriate arena on which one can explore interactions between various branches of Mathematics and Theoretical Physics.

Manifolds often come endowed with geometric structures, for example, with a way of measuring the length of a curve (Riemannian metrics) as in the case of a surface in 3-space. One can then define geodesics and curvature and study how these objects influence one another and interact with the topology of the manifold. A key illustration of this interplay, and a central result in this course, is the Gauss-Bonnet theorem, which shows that the average curvature determines the topological type of the surface.

Rather than worrying about how to define abstract manifolds (which you will see in Part III), we will study manifolds as objects already embedded in euclidean space. This will allow us to have a very short working definition of manifold and get fairly quickly into examples and the basic notions in Differential Topology (such as regular values, degree and transversality), giving a measure of how much one space folds onto another. Once we have set up the framework we will study the (Riemannian) geometry of curves and surfaces in euclidean space and we will prove at the end of the course that curvature can detect knottedness.

IB Geometry provides useful examples and an introduction to some of the ideas that we will develop and IB Analysis and Topology will be very useful when we set up the manifold framework in the first few lectures.

Analysis of Functions

Lent, 24 lectures

The analysis of functions has its roots in the rigorous study of the equations of mathematical physics, and is now a key part of modern mathematics. This course builds on the Part II courses Linear Analysis and Probability & Measure, which are pre-requisites, applying the theory of integration and the tools of functional analysis to explore such topics as Lebesgue and Sobolev spaces, the Fourier transform and the generalised derivative.

These topics are important and interesting in themselves, but the emphasis is on their use in other areas of mathematics (for instance in the representation of functions and in partial differential equations), rather than their maximal generalisation. You can get an idea of the flavour of the course by browsing *Analysis* by Lieb & Loss (Springer).

Logic and Set Theory

Lent, 24 lectures

The aim of this course is twofold: to provide you with an understanding of the logical underpinnings of the pure mathematics you have studied in the last two years, and to investigate to what extent, if any, the ‘universe of sets’ can be considered as a structure in its own right. As such, it has few formal prerequisites: some familiarity with naive set theory, as provided by the IA Numbers and Sets course, is helpful, but no previous knowledge of logic is assumed. On the other hand, the course has links to almost all of pure mathematics, and examples will be drawn from a wide range of subjects to illustrate the basic ideas.

The course falls into three main parts. One part develops the notions of validity and provability in formal logic, culminating in the Completeness Theorem, which asserts that these two notions coincide. Another part is concerned with ordinals and cardinals: these are notions that generalise the ideas of size and counting to the infinite. The final part is an introduction to formal set theory, where one makes precise the idea of a ‘universe of sets’, and studies its structure.

The book ‘Notes on Logic and Set Theory’ by P.T. Johnstone (C.U.P., 1987) covers most of the material of the course, and is suitable for preliminary reading.

Number Fields

Lent, 16 lectures

Number theory studies properties of the integers and the rationals. The questions that arise are usually very simple to state but their solutions are often very deep and involve techniques from many branches of mathematics. It is also a subject where numerical experiments have proved useful as a guide to the sort of result one might seek to establish. Diophantine equations constitute a central theme; they are basically polynomial in form and lead to the study of integer or rational points on algebraic varieties. Two particularly famous problems here, Fermat’s Last Theorem and the Catalan Conjecture, have only recently proved amenable to solution.

The course provides an introduction to algebraic number theory – it arose historically from investigations of reciprocity laws and attempts to solve the Fermat problem and it now forms one of the nicest and most fundamental topics in mathematics. Knowledge of the course on Groups, Rings and Modules is desirable.

The book *Problems in algebraic number theory* by J. Esmonde and R. Murty contains most of the material covered in the course. For a historical introduction to the subject, see also Chapter 1 of D. Cox’s *Primes of the form $x^2 + ny^2$* .

Probability and Measure

Michaelmas, 24 lectures

Measure theory is basic to some diverse branches of mathematics, from probability to partial differential equations. This course combines a systematic introduction to measure theory with an account of some of the main ideas in probability. You will be familiar with the Riemann integral from Parts IA and IB and have done some elementary probability in IA. The expectation operator of probability behaves somewhat like the integral, and in this course we see that they are both examples of some more general integral. These general integrals and the measures which underlie them have advantages over the Riemann integral, even for functions defined on the reals. In Part IA the definition and properties of expectation were only partially explored and here we do it more fully.

If you like to see how a substantial and coherent mathematical theory is put together, you will enjoy the measure theory part of this course, and this will be essential to any further work you do in analysis. It also underpins the probability which provides motivation and application throughout the course. The course ends with the Strong Law of Large Numbers and Central Limit Theorem, both of which are of real practical importance, being the mathematical basis for the whole of statistics.

A good book to read for the early part of the course is Probability with Martingales, by D. Williams (CUP, 1991).

Principles of Statistics

Michaelmas, 24 lectures

In Part IB Statistics, an introduction to the main statistical problems and inference techniques was given, including parameter estimation, hypothesis testing and the construction of confidence sets in a variety of examples, including specific families of distributions and models. The *Principles of Statistics* course aims to give a unified perspective on these problems, and develops the main mathematical theory that underpins these basic principles of statistical inference.

The first pillar will be the inferential paradigm surrounding the likelihood function of the observations, and the associated *maximum likelihood estimator*, providing a conceptually unified and in most situations also practical solution to the problem of statistical inference. The distribution of this estimator will be shown to have a universal limiting normal distribution, permitting the use of the estimator for statistical inference. A generalisation of the Gauss–Markov theorem from the linear model can be proved for this estimator, establishing that it is in a certain sense the best among all estimators. Related to the likelihood principle, but in other respects fundamentally different, is the *Bayesian* approach to statistical inference. This approach, likewise, will be developed for general families of parametric statistical models.

The study of the notion of optimality of certain statistical procedures from a general perspective is known as *statistical decision theory*, of which the main ideas will be presented in the course. The course will also develop the main ideas of some related classical fields in statistics that are crucial in applications, such as inference methods for multivariate data, nonparametric techniques, and resampling (Monte Carlo) procedures.

Requirements: Probability IA and Statistics IB are essential.

Stochastic Financial Models

Michaelmas, 24 lectures

This is concerned with the pricing of financial assets under uncertainty. It builds towards a presentation of the celebrated Black–Scholes formula for the price of an option on stocks. The holder of a call option on a stock has the right to purchase one unit of that stock at a specified ‘strike’ price within a designated time period. The holder hopes that within the period the stock price will go above the strike price whereupon the option may be exercised with the stock being bought at the strike price and sold immediately at the higher current price to yield a profit. What is the fair price to charge for such an option? In seeking an answer to this question, the course introduces some important ideas of probability theory including martingales and Brownian motion. Deciding when the holder should exercise the option leads to the techniques of dynamic programming and optimal stopping which are applicable throughout applied probability and statistics.

The main prerequisite for this material is Part IA Probability – if you liked that course then you should enjoy this one. Probability and Measure is recommended as a companion course, but it is not strictly necessary. No previous knowledge of economics or finance is necessary. It complements, but does not rely on, Markov Chains. To get a better idea of the sort of problems the course is seeking to tackle it is worth browsing in the book *Option, Futures and Other Derivative Securities* by J. Hull (Prentice-Hall, 2nd Ed. 1993).

Applied Probability

Lent, 24 lectures

This course provides an introduction to some of the probabilistic models used to study phenomena as diverse as queueing, insurance ruin, and epidemics. The emphasis is on both the mathematical development of the models, and their application to practical problems. For example, the queueing models studied will be used to address issues that arise in the design and analysis of telecommunication networks.

The material is likely to appeal to those who enjoyed Part IA Probability; Markov Chains is useful, but the style of the course, involving a mix of theory and applications, will more closely resemble the earlier course. Probability and Measure is a loosely related Part II course.

Mathematics of Machine Learning

Lent, 16 lectures

Suppose we are given a collection of emails labelled as either being spam or legitimate (known as ham). With this data, we aim to build an algorithm to annotate new unlabelled emails as spam or ham. This is known as a classification problem and is a central problem studied in Machine Learning, a field whose broader remit, like Statistics, is about how to learn from data. The focus of this course is the classification problem, other instances of which include predicting whether or not patients have a particular disease based on their medical history, or predicting whether a user will click an internet advertisement or not (a problem of great interest to search engines).

The first part of the course will deal with the theory of empirical risk minimization (ERM), a simple but powerful general strategy for building classification algorithms. We will work with a mathematical framework where each data instance (i.e. each email and label pair in the case of our first example) is viewed as a realisation of random elements, all sharing the same distribution. Studying ERM in this context will require us to develop important probabilistic tools, so-called concentration inequalities. Part IA Probability is a necessary prerequisite and in practice it is also advisable to have attended Part IB Statistics. The results we will derive were essentially state-of-the-art until the late 1990's and early 2000's; despite being suboptimal in some cases, studying them is a useful starting point for more sophisticated results.

The second part of the course studies computational aspects of ERM, building on parts of the IB Optimisation course (relevant material will be reviewed in the course). In particular, we will introduce stochastic gradient descent, which is now the most popular optimisation strategy for machine learning problems and currently the centre of a great deal of research. The last part of the course will illustrate some of the concepts and ideas developed in the earlier parts by studying random forests, boosting and neural networks, perhaps the three most successful methods in machine learning.

Asymptotic Methods

Michaelmas 16 lectures

There are many instances, arising not only in mathematical physics, but also in analysis and number theory, where one needs an approximation to a function for which no usable convergent series expansion is available. Typically, the function is given as an integral or else as the solution of a differential equation. It turns out that excellent approximations can be obtained using certain series, called asymptotic expansions, which are normally non-convergent. Such an expansion might describe, for example, the behaviour of an integral depending on a parameter, as the parameter becomes large; alternatively, it might describe the behaviour of a solution of an ordinary differential equation, as the independent variable becomes large.

A certain amount of familiarity with the basics of complex-variable theory is essential, either through Part IB Complex Methods or Part IB Complex Analysis. This would be reinforced by the Part II Further Complex Methods course, which is desirable but not essential. An introduction to the course material is given in A. Erdelyi *Asymptotic Expansions* (Dover 1956).

Numerical Analysis

Michaelmas, 24 lectures

Many mathematical problems, e.g. differential equations, can be solved generally only by computation, using discretisation methods. In other cases, e.g. large systems of linear equations, calculation of the exact solution is impractical and, again, we need to resort to numerical methods. Numerical analysis concerns itself with the design, implementation and mathematical understanding of computational algorithms. The course will address iterative techniques for linear equations, the calculation of eigenvalues and eigenvectors, and the solution of partial differential equations by finite differences (following the treatment of ordinary differential equations in Part IB). The last section of the course deals with Fourier expansions and their generalisations.

Part IB Numerical Analysis is an obvious prerequisite but Part IB Analysis courses, Complex Methods and Linear Algebra are also highly relevant. Mathematical ability is sufficient for understanding the course, while computational experience provides only a useful advantage.

Dynamical Systems

Michaelmas, 24 lectures

Contrary to the impression that you may have gained, most differential equations can not be solved explicitly. In many cases, however, a lot can still be said about the solutions. For example, for some systems of differential equations one can show that every solution converges to an equilibrium, while for others one can prove that there is a subset of solutions which are equivalent to infinite sequences of coin tosses ('chaos'). In this course, we study differential equations which can be written in the form $\dot{x} = v(x)$ with x in some (mainly two or three-dimensional) 'state' space. We take the "dynamical systems" viewpoint, concentrating on features which are invariant under coordinate change and time rescaling. We will find that two-dimensionality imposes severe restrictions though many interesting 'bifurcations' are possible: ways that the behaviour of a system $\dot{x} = v_\mu(x)$ can change as external parameters μ are varied. We shall also study nonlinear maps, which can be thought of either as difference equations or as a way of investigating the stability of periodic solutions of differential equations. We conclude with a discussion of chaotic behaviour in maps and differential equations, including a treatment of the famous logistic map. The treatment is 'applied' in flavour, with the emphasis on describing phenomena, though key theorems will be proved when needed.

If you browse P Glendinning *Stability, instability and chaos*, CUP, 1994 you will be well prepared.

The material contained in this course is relevant to any subject involving a modern treatment of differential equations. This includes most areas of Theoretical Physics, but usually not at the undergraduate level.

Electrodynamics

Michaelmas, 16 lectures

Electrodynamics is the most successful field theory in theoretical physics and it has provided a model for all later developments. This course develops from Part IB Electromagnetism. It starts by developing electromagnetism as a classical relativistic field theory, showing how the relativistic form of Maxwell's equations can be derived from a variational principle, and presenting the covariant treatment of the energy and momentum carried by the electromagnetic field. Such a treatment is essential for later use in quantum field theory. The remainder of the course shows how Maxwell's equations describe realistic phenomena. In particular, the production of electromagnetic waves by accelerated charges is discussed, and mechanisms for scattering radiation are briefly introduced. The course ends with the treatment of electromagnetism in continuous media, which allows one to understand some of the tremendous diversity in the electric and magnetic properties of real materials. The propagation of electromagnetic waves in materials is a particular focus.

The main prerequisite is familiarity with the basic ideas (especially those involving Maxwell's equations and their relativistic formulation) of Part IB Electromagnetism and knowledge of special relativity. As often for a theoretical physics course, the *Feynman Lectures* provide good introductory reading.

Principles of Quantum Mechanics

Michaelmas, 24 lectures

This course develops the principles and ideas of quantum mechanics in a way which emphasizes the essential mathematical structure, while also laying the foundations for a proper understanding of atomic and sub-atomic phenomena. In contrast to the introductory treatment given in Part IB, which is based entirely on wavefunctions and the Schrödinger equation, observables are presented as linear operators acting on vector spaces of states. This new approach has practical as well as aesthetic advantages, leading to elegant and concise algebraic solutions of problems such as the harmonic oscillator and the quantum theory of angular momentum. Some of the other key aspects of quantum behaviour that are treated include: intrinsic spin, multi-particle systems, symmetries, and their implications. Perturbation theory techniques, which are indispensable for realistic applications, are also discussed. The course ends by examining in more detail the inherently probabilistic nature of quantum mechanics, as illustrated by Bell's inequality and related ideas.

Applications of Quantum Mechanics

Lent, 24 lectures

This course develops the ideas and methods introduced in Part IB Quantum Mechanics and Part II Principles of Quantum Mechanics and uses them to explain how we probe and understand the structure of atoms and solids. The various material objects that surround us in the everyday world exist as vast collections of particles (electrons and nuclei) making up atoms, molecules and various crystalline substances. Quantum mechanics is essential for an understanding of how this happens.

An important tool for probing the structure of matter (finding out where the particles are, how the electric charge is distributed) is the scattering of a beam of particles of appropriate energy on targets of interest. The course develops the theory of scattering in a form applicable to both atomic and crystalline targets.

There are two particularly important aspects of crystalline materials: the elastic vibrations of the atoms in the crystal matrix and the dynamics of electrons moving through the crystal. In quantum theory the elastic vibrations are understood as particle-like excitations known as phonons. In travelling through a crystal both phonons and electrons exhibit a band structure in their permitted energies. The role of phonons and electrons in condensed matter physics and the significance of this energy band structure is explained by means of simple but physically significant quantum mechanical models. Energy bands are used to understand the properties of semiconductors and some simple devices such as the *pn* junction are explained.

Some idea of the material of the course can be gained by consulting a book such as *Principles of the Theory of Solids* by J. M. Ziman, (CUP, 1972).

Statistical Physics

Lent, 24 lectures

Statistical mechanics is the art of turning the microscopic laws of physics into a description of Nature on a macroscopic scale. This requires the development of tools to understand the properties of systems which contain a very large number of particles.

The course starts by defining new concepts such as entropy, temperature and heat, all viewed from a microscopic perspective. These ideas are then used to understand a wide array of phenomena, from gases, to the vibrations of solids (phonons), to the behaviour of light (blackbody radiation). In each case, the laws governing these phenomena are derived from first principles. The course goes on to describe more subtle quantum effects that occur at low temperature such as Bose-Einstein condensation and the formation of Fermi surfaces. The course ends with a discussion of thermodynamics and phase transitions, in which the properties of a system change discontinuously.

Ideas from Part IB Quantum mechanics are essential and Part II Principles of Quantum Mechanics is highly desirable. Lecture notes from David Tong's course are available to download at <https://www.damtp.cam.ac.uk/user/tong/statphys.html>, while a shorter set of notes representing almost exactly the material presented in the classes is available at <https://www.damtp.cam.ac.uk/user/us248/Lectures/lectures.html>.

General Relativity

Lent, 24 lectures

General Relativity is a relativistic theory of gravitation which supersedes the Newtonian Theory. This course shows how the theory can be built up on the foundations of Part IA Special Relativity. The necessary ideas from differential geometry will be taught ab initio, relying on the methods courses in Parts IA and IB. As an extended example, the course includes a careful treatment of the Schwarzschild spacetime and its interpretation as a black hole.

An elegant informal treatment of much of the material is contained in chapters 1,2,7 and 8 of W. Rindler *Essential Relativity* (Springer, 1977). A slightly more formal introduction is chapters 5,6,8,9,10,14-16 of R. d'Inverno *Introducing Einstein's Relativity* (Oxford, 1992).

Integrable Systems

Lent, 16 lectures

A *soliton* was first observed in 1834 by a British experimentalist, J. Scott Russell. It was mathematically discovered in 1965 by Kruksal and Zabusky, who also introduced this name in order to emphasise the analogy with particles ('soli' for solitary and 'ton' for particles).

Solitons appear in a large number of physical circumstances, including fluid mechanics, nonlinear optics, plasma physics, elasticity, quantum field theory, relativity, biological models and nonlinear networks. This is a consequence of the fact that a soliton is the realisation of a certain physical coherence which is natural, at least asymptotically, to a variety of nonlinear phenomena. The mathematical equations modelling such phenomena are called integrable. There exist many types of integrable equations including ODEs, PDEs, singular integrodifferential equations, difference equations and cellular automata.

The mathematical structure of integrable equations is incredibly rich. Indeed soliton theory impacts on many areas of mathematics including analysis, algebraic geometry, differential geometry, group theory and topology. However, it must be emphasised that the basic concepts of the integrable theory can be introduced with only minimal mathematical tools. This course will give an introduction to soliton theory with emphasis on the occurrence of solitons in nonlinear dispersive PDEs.

Fluid Dynamics

Michaelmas, 24 lectures

How does a hummingbird hover? How does a bumblebee fly? How, for that matter, does a Boeing 747 defy the pull of gravity? Does the bath-tub vortex really rotate anti-clockwise – or is it clockwise – in the Northern Hemisphere? How can a flow which exhibits an infinite sequence of eddies in a confined space satisfy a 'minimum dissipation' theorem? How can a flow that is strictly reversible have irreversible consequences?

Such questions lie within the domain of Fluid Dynamics, a subject that contains the seeds of chaos (and indeed provides the main stimulus for much of the current intense interest in chaos). The course will address the above questions, among others, in a progression from phenomena on very small scales ('low Reynolds number problems') to phenomena on very large scales ('large Reynolds number problems'). The course thus encompasses, at one extreme, flows that arise at the biological level (e.g. the swimming of microscopic organisms) and, at the other, flows on the scale of the Earth's atmosphere and oceans, or even larger. And, in between of course, it encompasses the bath-tub! Mathematical techniques, further to those developed in IB, will be used to determine solutions to the nonlinear, time-dependent Navier–Stokes equations.

Part II Asymptotic Methods covers some material which would be useful for this course. The course has natural links with Part II Waves and less obvious links with Dynamical Systems. A number of the Computational Projects are directly relevant. Introductory reading: *Elementary Fluid Dynamics* by D.J. Acheson, chapters 1-4.

Laboratory Demonstrations in Fluid Dynamics

Non-examinable

A series of laboratory demonstrations and experiments is used to expose you to material covered by the Part IB and Part II Fluid Dynamics lecture courses. The emphasis is on understanding the physics behind the mathematics, along with the limitations of the simple analytical models. Attending this course will help you develop the physical insight necessary to derive and evaluate mathematical models, and to determine whether their predictions are reasonable. Specific topics covered include potential flow, surface waves, Reynolds experiment, Stokes flow, Kelvin's circulation theorem, spin-up, boundary layers and bubbles. Student participation is encouraged but not required.

There are no prerequisites for this course and the course is not examined.

Further details will be announced during the Part II Fluid Dynamics lectures.

Waves occur in almost all physical systems including continuum mechanics, electromagnetic theory and quantum mechanics. In this course examples will be drawn from fluid and solid mechanics, although much of the theory has application in other contexts. In the first part of the course sound waves in a gas are studied (after which you will understand why you can hear the lecturer). Small amplitude acoustic waves are described by the wave equation (see IB Methods); however at larger amplitudes nonlinear effects must be included. The change in the governing equations caused by nonlinearity leads to the formation of shocks, i.e. sonic booms. Applications of the underlying theory to both traffic flow and blood flow are mentioned.

Linear elastic waves, e.g. seismic waves, split into two types: the faster-travelling compressional waves (cf. sound waves) and the slower-travelling shear waves. The surface waves that cause most destruction in an earthquake are also studied.

Not all linear waves have a wavespeed that is independent of wavelength. In such systems it is important to distinguish the speed of wavecrests from the speed at which energy propagates; indeed, the wavecrests and energy can propagate in opposite directions. As a consequence, (a) if you throw a stone into a pond to generate a circular wave packet, you will see that the wavecrests propagate outward through the wave packet and disappear, and (b) atmospheric waves generated near ground level can appear to the eye as if they are propagating down from the heavens! Finally, the ray tracing equations are derived. These are used to describe, *inter alia*, why you can go surfing (i.e. why waves tend to approach a beach perpendicularly), why the wave pattern behind a ship (or a duck) subtends a half-angle of $19\frac{1}{2}^\circ$, and why sound can travel long distances at night.

The mathematical techniques assumed are those covered in the IA and IB Methods courses. While the course is otherwise self-contained, there is a small amount of complementary material in Asymptotic Methods, Electrodynamics, and Fluid Dynamics.

A good book to look at is *Wind waves: their generation and propagation on the ocean surface* by B. Kinsman (Prentice-Hall).

Computational Projects (CATAM)

This course is similar in nature to the Part IB course. There are a variety of projects to choose from; some are closely related to Part II courses and others are not. As in Part IB, the projects are listed in a CATAM manual which is available on the Faculty website at <https://www.maths.cam.ac.uk/undergrad/catam/II>.

The course is examined by means of work handed in close to the beginning of the Easter Term.

5 Lecture Timetable: Clashes

Each term, there are only 8 slots for lecture courses¹ This year there are a total of 37 courses to be fitted into a total of 16 slots, so there will be 11 slots with double clashes and 5 slots with triple clashes.

The general policy is to avoid triple clashing any C-courses, and to avoid clashing any applied courses, any pure courses and any applicable courses. Further, by using examination statistics to gauge the popularity of combinations of courses, the aim is to minimise the effect of the clashes (for example, by not clashing a very popular pure course with a very popular applied course). However, there are also significant constraints associated with the fact that the most popular courses in Parts II and III will fit only in the larger lecture theatres (MR2, MR3 and MR9). In addition there are constraints concerning lecturer availability. At the end of the day, it is unfortunately inevitable that some students would have wanted to attend courses that clash.

A provisional lecture timetable should be available early in September, or thereabouts, from which you will be able to see the clashes.

¹ Or 9 if 16-lecture courses were distributed asynchronously.