Introduction

When we begin our studies of mathematics at university, we are mostly concerned with things that *are*. We seek to prove that equations have solutions, that sets are equal or in bijection, that groups are isomorphic and so on. Much of this mathematics is highly constructive: we write down the needed solution, bijection or isomorphism and are content.

The much subtler part of mathematics generally concerns things which *are not*[1]. Compare the relative ease of, for example, showing that the sets $\mathbb{Z}$ and $\mathbb{Q}$ have the same cardinality with the much more difficult proof that $\mathbb{Q}$ and $\mathbb{R}$ have different cardinalities. When we want to show that two objects are not the same, we must work a good deal harder than to show that they are. For groups, the question of when two groups of the same cardinality are not isomorphic is exceedingly difficult and, indeed, is algorithmically unsolvable. We certainly cannot check every possible bijection to see if it respects the group structure!

This last statement is not always correct, of course. There are groups for which the blunt method of testing all possible functions will eventually work: the finite groups. Actually doing this in any practical context would be a fairly daft way to show two groups are not isomorphic, but it is at least a comfort to know it is possible.

A less daft thing to do than checking all possible isomorphisms is to look for group 'invariants': properties of a group that are preserved by isomorphism. If we find two groups in the gutter, we could attempt to show that they are different by checking our favourite list of invariant properties to see if the two groups have any different properties. Perhaps one is abelian, and the other is not? Perhaps one has centre, or has torsion of a particular order?

As an undergraduate, when I was first introduced to free groups, I not unnaturally asked how we could be sure that free groups of different ranks were genuinely different. The answer I was given was that the two groups had different numbers of homomorphisms to a cyclic group of order two. There are, of course, plenty of other ways to answer this particular question, but the fact remains that the easiest thing to do was to reduce the question to the behaviour of the finite quotient groups, where we can be certain that we can count things exactly.

The topic of this book is the study of how far we can push this idea: what can we learn about a group if we only know its finite quotients? Which groups will have enough finite quotients that we can use them to meaningfully study the whole group? How are we to work with such an unwieldy piece of information as "the collection of finite quotients of a given group"?

[1] No pedantry about double negatives, please.

One approach would be to try to work with each finite quotient one at a time as a sort of 'approximation' of the infinite group, and then to pass to larger and larger finite quotients as needed to witness more and more information. We do in fact do this very frequently in everyday life. By any sensible measure, the collection of dates on a timeline forms a copy of the group of integers $\mathbb{Z}$ once we have fixed a 'zero day'. We are unlikely to specify a date by actually giving the total number of days since day zero. We would first use a very small approximation to $\mathbb{Z}$, the group $\mathbb{Z}/7\mathbb{Z}$: the set of days of the week. If we could not distinguish two Tuesdays, we would pass to a larger quotient, the group $\mathbb{Z}/7\mathbb{Z} \times \mathbb{Z}/365\mathbb{Z}$, to tell us the position in the year[2]. Over a longer span of time we might add a two-digit year date, or a four-digit, to get more refined positions on the timeline. At present, in the Gregorian calendar we go no further, but in the fullness of time we may expect to add more digits to see more of the group $\mathbb{Z}$. The Mayan Long Count calendar works with a very large quotient of $\mathbb{Z}$ indeed!

Dealing so explicitly with finite quotients of groups rapidly becomes a little impractical. We would prefer to take these larger and larger finite groups and have some sort of 'limiting object' to study that contains all the information from the finite groups. This limiting object is called a *profinite group*. This book contains a study of the formal constructions and properties of profinite groups. We especially focus on those which arise as *profinite completions*: we take a 'normal' group, find the collection of all its finite quotients, and form the limiting object.

This is not the only source of profinite groups: they arise very naturally in mathematics, and especially in number theory, as Galois groups. When studying the behaviour of the Galois group of, say, the algebraic closure $\overline{\mathbb{Q}}$ of the rational numbers, one would rarely think about the entirety of $\overline{\mathbb{Q}}$. One may instead prefer to take some algebraic number of interest, take a finite field extension of $\mathbb{Q}$ which is Galois and contains the given element, and study the Galois group of the finite extension. The finite Galois group is a quotient of the full Galois group of $\overline{\mathbb{Q}}$, and larger and larger finite field extensions give better approximations of this full Galois group: the full Galois group becomes a limit of these finite groups, a profinite group. Although Galois groups and number theory are important examples and motivation to study profinite groups, in this book I will focus far more on the group theory and on profinite completions and only mention the number theory very occasionally.

I will now give a general survey of the structure of the book, and the key topics covered in each chapter. We begin with a discussion of *limits*. I have said that one goal of profinite groups is to combine the information of a system of larger and larger finite groups into a single object we can study, which certainly suggests we should seek some form of limiting object. The correct notion of limit for our purpose is that provided by *category theory*.

[2]Let us pretend leap years do not exist.

The first chapter therefore includes a discussion of category theory from the ground up. I have not assumed any prior familiarity with categories as such – although any mathematician is already familiar with plenty of examples of categories, whether they have heard them described as such or not. Categories provide the general language of objects and maps between them. Working with categories allows us to define and study limits in many settings at once, rather than having to reprove each result for sets, groups, rings, and so on.

The broad notion of 'limit' in category theory being rather *too* broad for our purposes, we then restrict to a particular kind of limit which to some extent is the star of the whole story – the *inverse limit*. This refinement gives a little more texture to the limiting process and opens up all the constructions we will need – we will constantly be referring back to the results in this chapter, in particular the theorems giving non-emptiness of various limiting sets.

At this point the other star of the story appears – *topology*. Initially our limits will be simply sets or groups, but we can equip them with natural topologies which, when we take a limit of a collection of finite groups or sets, is a *compact Hausdorff* topology. Any student of topology knows that compact Hausdorff spaces are an excellent place to work, and it is this well-behaved topology that rescues us: an incredibly unwieldy limit group becomes somewhat tamer. The behaviour of this topology is the primary reason that instead of studying limits of arbitrary groups, we study limits of finite groups – the eponymous *profinite groups*.

Chapter 1 thus gives us a foundation for our study. In Chapter 2 we will meet and study our first real profinite groups and give both examples of profinite groups and ways to construct more. In a sense, this chapter is concerned with seeing how far the world of profinite groups runs parallel to the development of group theory in general, being something of a hybrid of infinite group theory and finite group theory. We establish the familiar constructions of subgroups, quotients, products and so on, with the admixture of our compact topology leading us to restrict to *closed subgroups* and *continuous quotients*. Comfortingly neither of these constructions leaves the category we are trying to study: taking a closed subgroup or continuous quotient of a profinite group again gives a profinite group. From the theory of finite groups, the notion of a Sylow subgroup will enter the picture, allowing us in some cases to consider profinite groups "one prime at a time".

We will see key examples of profinite groups which mimic the behaviour of the familiar essential examples of groups: we have a replacement for the integers in the form of the *profinite integers* $\widehat{\mathbb{Z}}$ and we have various special and general linear groups over profinite rings. We will construct, in great generality, an appropriate notion of *free profinite groups*, from which all other profinite groups can be constructed as continuous quotients.

We also study generation of profinite groups and their actions on sets – again, if it need be emphasized further, with the inclusion of the topology so that we speak of *topological generating sets* and *continuous actions*. Even the automorphism group of a profinite group acquires a topology allowing it to act continuously on the profinite group, the *compact-open topology*. I will give a self-contained account of this topology so the reader will not need to seek out some grand old tome on topology to follow along.

The structure of profinite groups thus appears to track roughly in parallel with the abstract theory of groups, when combined with plenty of topology. In Chapter 3 we will begin to explore a close connection between the two worlds: the *profinite completion*. This operation, which is a vital part of the story of the book, takes an abstract group and builds from it a profinite group whose continuous finite quotients exactly agree with the finite quotients of the original group. The profinite group being the limit of these finite groups, one may say that the profinite completion remembers everything about the original group that could be seen in its finite quotients, and nothing else. We will prove, for example, that two (finitely generated) groups have the same collection of finite quotients if and only if their respective profinite completions are isomorphic as topological groups.

This linkage between abstract and profinite groups may be thought to have two principal directions: we may attempt to study a finitely generated group by looking at its profinite completion, or we may attempt to study the profinite completion by building finite quotients of the original group. Both directions require the original group to have a healthy supply of finite quotients if we are going to get something non-trivial: trying to take the completion of an infinite simple group, for example, would simply give the trivial group and we would learn nothing at all. Very naturally, therefore, we are drawn to study those groups for which we can guarantee a good supply of finite quotients with which to work.

The basic notion is that of *residually finite groups*: groups for which we can, at the very least, distinguish all the different group elements from each other by looking at finite quotients. These groups inject into their profinite completions as a dense subgroup, giving something of a skeleton inside the profinite group. We can then ask more detailed questions about how similar the profinite group looks to this skeleton: will their subgroup structures be similar, will their conjugacy classes interact in a sensible way, and so on. From these questions come more refined 'residual properties' studied later in the chapter, *subgroup separability* and *conjugacy separability*.

We naturally want to have examples of residually finite groups. A major source of examples, as we will see, comes from *linear groups*: the simple idea of "reducing a matrix modulo a large prime" takes one a surprising distance. For the more detailed properties like subgroup separability, however, this is insufficient, and to build our finite quotients we are drawn towards the theory of covering spaces. The need for a

well-behaved space to cover shows why, for much of this book, our examples come from *geometry*. We will thus spend a good amount of time studying free groups in their guise of fundamental groups of graphs, as well as groups of "the next dimension up": fundamental groups of surfaces.

We also ask what properties of a group can be recovered from its profinite completion. Another way to state the issue is this: if two residually finite groups have isomorphic profinite completions, what properties must they share? This topic is the study of *profinite invariants*, an area of mathematics in which much is yet to be discovered. In this part of the chapter we will introduce some profinite invariants – principally those connected in some way with abelian quotients – and see some examples of natural properties which fail to be profinite. We also give explicit examples of distinct groups sharing the same profinite completions.

In Chapter 4 we return to the study of profinite groups in their own right, or more precisely to the study of a particularly well-behaved family of profinite groups, the *pro-p groups*. In their first group theory course, the reader will have seen that finite groups whose orders are a power of a prime p are rather more tractable than arbitrary finite groups. The only simple group about which we have to worry is the cyclic group of order p, and all other p-groups may be built up step-by-step from this one group. The possibility of various sorts of 'inductive' proofs is opened up, and we will see that many of these techniques allow us to study pro-p groups, which are the inverse limits of systems of finite p-groups.

This inductive approach allows us to study generation of pro-p groups in rather more precision than for general profinite groups, and also to engage in a style of p-adic arithmetic via the techniques of Hensel's lemma. We go on to apply these methods to matrix groups over the p-adic integers, and see that these matrix groups have a rather more rigid structure than the special linear groups over the honest integers. As all profinite matrix groups are built in direct products of such p-adic matrix groups via the Chinese remainder theorem, this chapter serves as an introduction to the study of arithmetic groups, and onward to the study of p-adic analytic groups.

Restricting our attention to pro-p groups also enables us to establish an important result that, while true for general (finitely generated) profinite groups, is exceedingly difficult to prove: the 'uniqueness', in an appropriate sense, of the topology of a profinite group. That is, the compact topology which we built in Chapter 1 and have used ever since is not some artefact of the construction or dependent on any choices, but is forced upon us by the group structure itself.

At this point the book seems to change directions entirely: Chapter 5 provides the reader with a good grounding in group cohomology, or at least the parts of it that will be useful later on. It may seem somewhat odd for the longest chapter of the book to mention none of the words from the title, but cohomology is deeply engrained

in the study of profinite groups and it is essential to have a firm grounding in the abstract theory before attempting to consider the cohomology of a profinite group. I did not wish to insist that the reader should be an expert in group cohomology before beginning this book, so I have included this self-contained chapter to build from.

The treatment of cohomology theory I have chosen to give is rather category-theoretic. It is possible to study group cohomology in a fairly 'explicit' way, working directly with cocycles, and this viewpoint will be available to us, but I deemed it important for the future to introduce the proper machinery of homological algebra first. This machinery mimics the cellular structure of a cell complex, allowing us to tease apart the behaviour of our group in different 'dimensions' and quantify this behaviour as a sequence of abelian groups.

We give the general theory and construction of group cohomology, and the relations to the subgroup and quotient group structure of a given group. We go on to make a particular study of the cohomology groups in dimension two. These have a very concrete connection to the behaviour and classification of *group extensions*, which we give in full. The chapter concludes with a collection of worked examples – again those groups with geometric grounding, like surface groups, play a primary role.

In Chapter 6 the two threads of profinite groups and cohomology theory reunite, and we study the *cohomology theory of profinite groups*. At its core, this theory bears a strong similarity to the cohomology theory of abstract groups, with the word 'continuous' scattered liberally among the proofs to account for the topology. The most rigorous way to proceed would be to repeat the whole of Chapter 5 with the topology appropriately included, but this is needlessly time consuming: the proofs are so similar that I have chosen simply to incorporate the results of Chapter 5 by fiat, and proceed to talk about the more unique aspects of the profinite theory.

One key 'unique' point of this theory is the strong form of duality known as *Pontryagin duality*. This gives a strong enough symmetry between the homology and cohomology groups of a profinite group to treat them as a single unified theory, in contrast to abstract groups where one really needs both theories at different times. Another important aspect we discuss that is not present in the classical cohomology theory is the intrusion of Sylow theory: since all objects involved are limits of finite objects, we will see many scenarios in which we can simply take one prime at a time to simplify matters.

The utility of working with a single prime at a time is seen when we consider the cohomology of a pro-p group with coefficients in the cyclic group of order p. We prove the striking theorem that a pro-p group is free if and only if it is 'one-dimensional' in the appropriate sense, and extend this to quantify the one-dimensional profinite groups. Analogous results are known to be true for abstract groups, but the proofs are very much more involved than in the pro-p case. We can even establish more detailed theorems which one would very much like – but which fail – for

abstract groups: the one-dimensional cohomology controls the number of generators of a pro-p group, and the two-dimensional cohomology controls the number of relations needed to describe a pro-p group.

The obvious question, given the context of the book and the previous chapters, is "how much does the cohomology of a profinite completion look like the cohomology of the original group?". We study this question in Chapter 7, with particular focus on those groups which have the exact same cohomology as their profinite completions. This property was seemingly introduced by Serre in an exercise in his book *Cohomologie Galoisienne*, and received the throwaway name 'goodness'. Since then, the concept has grown in importance, but unfortunately has received neither a new name nor published proofs of the original exercises. I have attempted to rectify this, and offer the new name 'separable cohomology' to bring the nomenclature into line with the related notions from Chapter 3. We give a complete treatment both of the results from Serre's exercises and from certain later research papers, although I have offered new proofs of many of these.

Two key points of this chapter deserve note here: the use of separable cohomology to study residual finiteness of extensions of groups (a topic notably missing from our first chapter on residual finiteness); and the construction of sensible chain complexes to compute and work with the cohomology of a profinite group, which somewhat amends the lack of geometric classifying spaces for profinite groups.

In the preceding chapters we have considered many different constructions for building new groups from old ones: subgroups, extensions, (semi)direct products, and so on. The missing construction is the operation of 'gluing' several groups together by identifying them along a common subgroup. Such *amalgams* of groups will be studied in Chapter 8. We pay particular attention to the question of when gluing together residually finite groups should produce a group which is again residually finite. This is generally studied using the notion of *efficiency*, although I have recast the common definition into an equivalent form which I feel better expresses the geometric content of amalgams as groups which act on trees. We go on to study the equivalent notion of amalgams for profinite groups, and the extent to which profinite completions commute with our amalgamation operations.

Our studies of cohomology remain relevant in this final chapter: by means of *Mayer–Vietoris sequences*, we may connect the cohomology of an amalgam to the cohomology groups of its various pieces. This allows us to greatly expand our list of groups known to have separable cohomology beyond those studied in Chapter 7.

Many mathematical books include at this point a diagram showing the logical interrelations of the several chapters: which chapters depend on which other chapters. I have included such a diagram as Figure 1. The structure of this book is maybe not so convoluted as to really need such a diagram, but I rather enjoy them.

Each of these chapters is equipped with several exercises. As with any new piece of mathematics, it is a good habit to work through some problems as well as simply reading through the material. Many of these exercises constitute important results in their own right, and if something seems to be missing from the main text it is a good idea to check for its inclusion in the exercises. The danger with including results of substance as exercises is, of course, the frustration that comes if one ever needs to cite them. As such, I have included a full set of solutions at the end of the book, so the reader can check the proofs and cite these results as securely as if they were labelled 'Theorem'.

Certain remarks in the text, and certain sections, have been labelled 'Discussion'. These are intended to make the reader aware of various adjacent or future topics of mathematics which are relevant to the study of profinite groups and residual finiteness but whose formal proofs go too far astray, or are too difficult, to warrant inclusion in this book. These are not essential for understanding the remainder of the book and should be treated as hopefully interesting asides.

Some discussion is needed of the relation of this book to the other standard works on the subject. For the general account of profinite group theory I have drawn great assistance from the books of Wilson [142] and Ribes and Zalesskii [117], although I have in all cases tried to rework the proof of each theorem to read as well as possible and form a cohesive whole. Wilson's book includes more material on Galois groups than I have chosen to include. Ribes–Zalesskii deals with profinite groups in perhaps greater generality than I have done, by considering constructions of 'pro-$\mathcal{C}$' groups for various families of finite groups $\mathcal{C}$. I made the conscious decision to consider only the case when $\mathcal{C}$ is the class of all finite groups, or occasionally the class of finite p-groups for a prime p. This is the main focus of the vast majority of papers in the area and allows me to give what I believe is a more accessible treatment of the subject. After understanding the material covered in this book, the reader who needs a more specific pro-$\mathcal{C}$ construction will be much better equipped to consult [117] than they would be if they began the field in the greatest possible generality. Neither of these books has the amount of focus on residual finiteness or geometry that I have chosen to include in the scope of this book.

For the treatment of the classical cohomology theory I have relied on Brown [26] to fill the gaps in my own knowledge. For the profinite version I again draw some parts from [117], but also include aspects of the treatment given by Symonds and Weigel in [136] and by Serre in [123]. Again I have reworked the proofs and, at least in the case of Serre's exceedingly terse account, greatly expanded them. The end result, then, reads rather differently from any of these sources, though many of the key results are of course shared with them.

The final two chapters, concerning groups with separable cohomology and amalgams, are far more new and do not, to my knowledge, have a good treatment in any of

the present textbooks. Some material is drawn from Serre [123], from Ribes [115] and from such primary sources as are cited in the text, but I have given new and in many cases genuinely different proofs of these results. The material on residual finiteness in Chapter 3, while well known, is also rather undertreated in the standard textbooks and here again I have relied more on the listed primary sources.

I would like to thank Martin Bridson, Jack Button, Andrei Jaikin-Zapirain, Dawid Kielak, Rob Kropholler, Ismael Morales, Mark Pengitore, Alan Reid, Ric Wade, Henry Wilton and Julian Wykowski for their helpful comments and discussions about the topics in this book and/or assisting with the finding of references. Thanks are also due to my family and friends and assorted colleagues, doctoral students and the Fellows of two colleges who have had to listen to me ramble about the business of profinite groups during the writing process. I also wish to thank Prof Henry Beker for his support of my teaching post at Selwyn College, without which I would not be able to write this book.

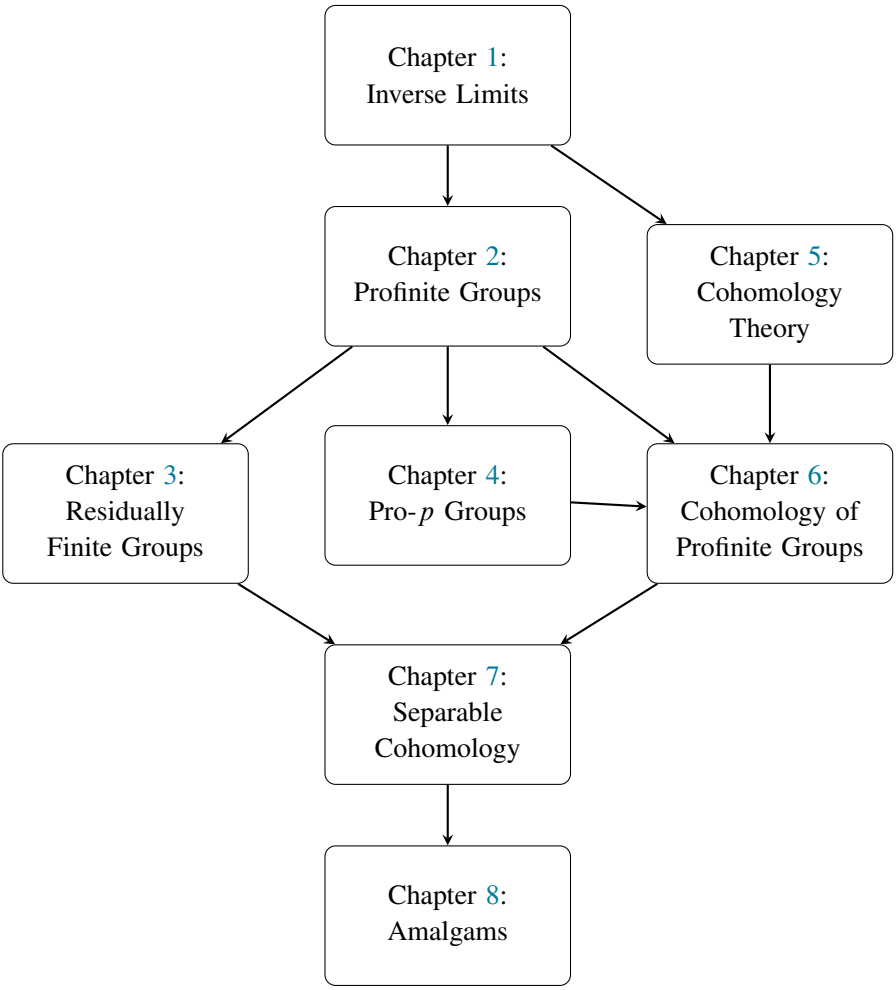

Figure 1. Logical relations of the chapters of the book.