

Preface

Libretto

What possesses us to write a book on sieve methods? Hopefully we have something to say that is worth listening to and, if we didn't write it, maybe a little something would be lost.

Many young people are drawn to mathematics and frequently what is most attractive to them are problems which are related to the most elementary objects, the integers, the squares, and the primes. The sieve method, itself an elementary idea two thousand years old, offers, even to the high school teenager, a natural approach within her own compass. Many of our colleagues can relate to this experience, as do we. Just as the student during his development acquires more sophisticated tools and tastes, so has the subject of sieve methods itself; and, whereas it is the elementary ideas that offer the initial attraction, it is the ability to infuse these with the more advanced tools that brings joy in later life. One of our goals is to transmit some of these joys to new investigators.

Eratosthenes, about 300 BC, is the first to be mentioned. His idea, illustrated in Chapter 1, is essentially an algorithm for tabulating primes. Following a very long gap, the subject was taken up by A. Legendre who gave a formula (Section 1.1) for $\pi(x)$. This was the beginning of the principle of inclusion-exclusion. Legendre's formula, however, is of very limited practical value and the ideas which were needed to turn it into a useful instrument were initiated by V. Brun in 1915. He developed his methods over the following decade before leaving the topic. We discuss his ideas in Chapter 6. Such a late placement should alert the reader that we do not follow a chronological order of sieve developments in this book.

The combinatorial complexity of Brun's most advanced works inhibited later researchers so that the subject did not receive the huge initial impetus that the power of his ideas merited.

Brun's results were improved considerably by A. Buchstab in the late 1930's using an iterative scheme, the basis of which has remained valuable to this day and which is discussed in several chapters, especially Chapter 11.

A new impetus came to the subject with the ideas of A. Selberg in the late 1940's (discussed in Chapter 7). His upper-bound sieve, which, following Selberg, we call the Λ^2 -sieve, is, in particular, very elegant in comparison with the earlier Brun sieves. It is also considerably stronger in some respects, although subsequent developments have shown that the strongest results are obtained by keeping both the Λ^2 -sieve and the combinatorial sieve in mind.

More recent progress has involved many new ideas and many names. Not only has the theory of the sieve apparatus grown but it has developed enhancements which permit the implementation of sophisticated results from harmonic analysis,

from exponential sums, from arithmetic geometry, and from automorphic forms. The reachable targets have expanded, not only from almost-primes to primes, but also to questions of greater variety, for example to solutions of diophantine equations. We hope the contents of the book will speak for themselves.

Let's take a quick run-through of the chapters. Although we follow roughly the order of development of the subject there are significant detours. Thus, for example, Bombieri who appears in Chapter 3 is thought to be somewhat younger than Eratosthenes who appears in Chapter 4. The reader we hope has sufficient mathematical maturity not to be troubled by this.

In Chapter 1, following the introduction to exclusion-inclusion and the formula of Eratosthenes-Legendre, we give a number of examples of arithmetic problems which can be attacked by the sieve. Although there are many we do not cover, those presented here suffice to show the versatility of sieve ideas.

In Chapter 2 we present, still following the historical order, ideas (due to Dirichlet and Tchebyshev) which are of an elementary nature and which will be used in sieve arguments. These are methods for evaluating sums of arithmetic functions of convolution type and are based on changing the order of summation in the spirit of the Dirichlet hyperbola method. A second purpose is the evaluation of sums and products over prime numbers to which we shall need frequent reference. Further results of this nature are postponed to Appendix A.

In Chapter 3 we develop Bombieri's sieve which, despite the name, we consider to be the transition between the preceding ideas of Tchebyshev and the sieve ideas to follow. Consequently, although historically out of place, it seems to fit here naturally. An understanding of these ideas should help the reader to appreciate the magic of positivity in the sieves that are to follow. More importantly, the results of this section offer the best way to view the parity phenomenon which lies at the heart of the limitations of classical sieve theory.

We are now ready to start with the sieve. However, before entering the mainstream, we take the opportunity in Chapter 4, for historical reasons, to show the extent to which the original sieve of Eratosthenes and Legendre can be developed by today's techniques. Surprisingly, these straightforward ideas can even lead to asymptotic formulae under favourable conditions.

In Chapter 5 we introduce many of the basic sieve principles, terminology and notation. To some extent, we have already out of necessity been doing this, even in Chapter 1. However, at the risk of a little repetition, we want to gather these in a single location for ease of reference. From the start, a serious impediment to an understanding of the sieve has been the difficulty in coming up with a satisfactory notation. The literature is littered with awkward choices. It is especially difficult to be consistent throughout the presentation of a full-length book. We think the reader is better served by the adoption of a notation which sacrifices uniqueness in order to be more suggestive. Consequently, one finds slightly modified names for similar objects. In such cases, the context helps to clarify the issue. The material in the final few sections of Chapter 5 is of a more advanced nature but still deals with general sieve principles so is appropriate to include before we get to specific sieve procedures. Moreover, it is beneficial to see it inserted now, while the notation is fresh, making its exposition more transparent.

In Chapter 6 we finally come to the real beginning of our story. Brun gave up on the goal of finding asymptotic formulae and was willing to settle for upper

and lower bounds. This allowed him, by positivity, to simply throw away many of the pieces of the problem which he could not handle. The fact that the few pieces remaining were enough to give useful results still today seems a miracle. Even the simplest version of Brun's sieve, which we present in the first section, produces remarkable results. Although less than perfect, such results seem hopeless to derive without the sieve by any modern means. In the following few sections we refine the original constructions of Brun, not as he did, but following along his lines and equipped with a good deal of hindsight. In particular, we take advantage of an important observation due to Buchstab which, by iteration, produces a sequence of successively stronger sieves. This idea will lead us later to the construction of sieves which are optimal in the most important cases. Before we come to these strong but complicated results, we give in Section 6.5 a simpler treatment, nevertheless, already obtaining theorems of fundamental nature. The resulting upper and lower bounds have the right order of magnitude and are sufficient for many, possibly most, of the sieve applications in the general literature. In the final two sections we illustrate some of these applications. In particular, the application in the last section gives a glimpse of the techniques we shall use to go beyond the natural limitations of the sieve. Along the lines of the earlier Section 6.8, which is needed to prepare for this, one learns that it is useful to understand how the sieve mechanism works and not merely the final statements of the theorems.

Returning to the theorem in Section 6.5, an important feature of the result is the fact that the upper and lower bounds for the sifting function approach each other, and at more than exponential speed, as the sifting range gets smaller. This feature is of such importance throughout the subject that it merits the name by which it has become known, "Fundamental Lemma", a term which was coined by Halberstam and Richert.

Chapter 7 is devoted to Selberg's sieve. His upper-bound sieve, prompted by his earlier work on the zeta-function, introduced ideas fundamentally different from those of Brun. We discuss various features of the Selberg sieve at length. It is distinguished not only by the power of the results but also by the elegance of the arguments. The estimates are particularly strong when the sieve dimension is large. This power is exploited brilliantly in the work of Goldston, Pintz and Yıldırım on small gaps between consecutive primes. Their argument is one case where the Selberg sieve seems to be essential. In the final sections of the chapter we show how these ideas work.

Given any upper-bound sieve, one can obtain a lower-bound sieve (and vice-versa) by Buchstab's recurrence formula. This possibility is particularly convenient in the case of the Selberg upper bound which, unlike the situation with Brun, has no natural lower bound counterpart. Ankeny and Onishi were the first to perform this procedure. The resulting estimates, upper and lower, can be further improved by iterations of the operation. Actually, it requires infinitely many such iterations to obtain the best results in all ranges. We do not present these in the book beyond the first iteration. In the case of the linear sieve, a combination of Selberg's upper-bound sieve with the construction of a combinatorial sieve of Brun's type led Jurkat and Richert to the essentially best possible upper and lower bounds for the sifting function in all ranges. These are the same results as could have been obtained by an infinite number of iterations. We shall also obtain the same results as Jurkat and

Richert (in Chapter 11), but by a more direct construction, using the combinatorial sieve.

An important problem in the theory of the sieve is concerned with the “sifting limit”, that is, the maximal possible range for which one can get a positive lower bound in a general setting. If κ denotes the dimension of the sieve problem it has been conjectured that this sifting limit parameter should be 2κ (actually we shall see that it is sometimes better). The results that we shall develop later show that this indeed is true for $\kappa = 1$ and $\kappa = \frac{1}{2}$. For κ very large the results in this respect were rather far from the truth until Selberg introduced a new combinatorial device into his upper-bound sieve, giving bounds for the sifting limits which approach 2κ as $\kappa \rightarrow \infty$. We present his arguments (slightly modified) in Section 7.6.

There are interesting problems which require the removal of elements from many residue classes per modulus. The combinatorial sieve proves to be quite effective if the number of unwanted classes is small or at least bounded. On the other hand, Selberg’s upper-bound sieve shows to advantage here in case the number of classes is large or even increases with the modulus. For example, it is capable of bounding correctly the numbers of squares by removing half of the classes for each prime modulus. Here, the concept of sieve dimension is meaningless and the Selberg sieve benefits from being insensitive to the sieve dimension. In Chapter 8 we give a number of applications of Selberg’s sieve in such circumstances.

There are, however, very different techniques for studying these “large sieve” problems. In Chapter 9 we develop some of these techniques, discussing in Section 9.5 the extent to which they are related. Linnik originated the large sieve theory. Over time, various alternative approaches and general formulations have been developed, reaching the point that the subject is no longer recognizable as a sieve. Nevertheless, the main statement retains the name “large sieve inequality”. The most frequently used, among its many forms, is an inequality for trigonometrical polynomials at well-spaced points, which we establish in Theorem 9.1. The main idea behind results of this type is the duality principle for linear operators on a vector space.

An important example of well-spaced points is provided by the Farey fractions. This example leads to estimates for character sums, both additive and multiplicative. The results are so strong that they compete with the consequences of the Riemann Hypothesis. Indeed, the large sieve inequalities for multiplicative characters lead to bounds for the error term in the Prime Number Theorem in arithmetic progressions (Bombieri–Vinogradov theorem) which have served as an unconditional substitute for the Riemann Hypothesis in many applications. This will be evident in our frequent use of this theorem throughout. The Farey points are just the roots of linear congruences. Another example of well-spaced points is given by the roots of quadratic congruences and the large sieve produces powerful consequences; see Propositions 20.5 and 20.7. A beautiful aspect of the large sieve inequality is that it holds for character sums with general complex coefficients. This offers applications inaccessible even to the Riemann Hypothesis. In the later sections of Chapter 9 we derive several estimates for bilinear forms in arithmetic progressions which, due to this generality, will be of great use to us later.

By this time, we feel it is appropriate to say a little more about the sieve weights. Although introduced simply as tools to serve our purpose, these surrogates for the Möbius function have an inner structure which is of interest on its own. Revealing

such properties can also provide a greater understanding of the sieve mechanism. In Chapter 10 we probe this “molecular” structure.

We have already seen in Chapter 6 the motivation for and construction of the beta-sieve weights and their application to the Fundamental Lemma. Although extremely powerful when one is not sifting very far, the results there suffer in quality when the sifting range is large, which is the case in the most advanced applications. This is not the fault of the choice of the weights but rather because we used there some crude bounds which could be quickly executed and are sufficient for the Fundamental Lemma. In Chapter 11 we reconsider the beta-sieve and work a lot harder to determine the optimal value of β and, using that, the optimal upper and lower bounds for the sifting function in all ranges. The analysis is very delicate. It depends on the theory of differential-difference equations which we postpone to Appendix B and, in particular, on a careful treatment of a sensitive convergence question. We conclude the chapter with some applications of the results for the fractional dimensions $0 < \kappa < 1$. The most important cases of $\kappa = 1$ and $\kappa = \frac{1}{2}$ are given individual additional treatment in subsequent chapters.

In Chapters 12 and 13 we concentrate attention on the linear sieve, that is, the sieve of dimension $\kappa = 1$. Because the upper and lower bounds are best possible with respect to the main terms, as we illustrate in Section 12.3 using a construction of Selberg, it makes sense to invest the effort to obtain sharpened secondary terms which didn’t merit attention in the general development in Chapter 11. In particular, we expend considerable energy to get the best possible error term in the main terms by using differential-difference equations especially designed for that purpose. More importantly, we pay renewed attention to the remainder. We manage to give the remainder in the shape of a bilinear form. This opens the possibility to apply various techniques of harmonic analysis, such as given in Chapter 13, to obtain estimates superior to those following from the traditional trivial treatment. Consequently, the results are stronger. The flexibility of ranges in our bilinear remainder allows us to choose the parameters to fit the results of the harmonic analysis in the best way. This new input from harmonic analysis to sieve theory makes the latter less elementary but the rewards are great.

Chapter 13 is entirely devoted to the illustration of this input from harmonic analysis for a variety of sifting sequences. First we consider the sequence of integers in an arithmetic progression, use Fourier expansion of the remainder terms, apply bounds for Kloosterman type sums, and conclude an upper bound for $\pi(x; q, a)$ which is quite strong, uniformly in large q . We then consider the sequence of integers in a short interval and use similar analysis, but with exponential sums suited to this sequence, getting similar improvements. Finally, we show how this analysis can be combined with Tchebyshev’s method rather than the sieve to show the existence, in the sequence, of integers having a large prime power factor. Here we have only considered two of the simplest sequences. But harmonic analysis is very versatile. A lot more will follow.

The beta-sieve with $\kappa = \frac{1}{2}$ is called the semi-linear sieve and several of its applications are featured in Chapter 14. In this case the sifting limit is $\beta = 1$ which is the best that one could hope for in any sieve. This means that one can perform sieving in the maximal range conceivable, right up to the level of distribution. The lower and upper bounds given by the sifted sums are optimal. The example which demonstrates this optimality is attractive from an arithmetic point of view. Indeed,

we show the upper bound becomes an asymptotic for some interesting sequences, such as the sequence of values of a binary quadratic form. The lower bound is also very interesting, even when the exponent of distribution is as small as $\frac{1}{2}$, in which case it just misses in general to give the desired output. Moreover, in many cases it can, with additional arguments, be made to do so. For example, when the sifting sequence consists of shifted primes we do have exponent of distribution $\frac{1}{2}$, due to the Bombieri–Vinogradov theorem, and we conclude that shifted primes are representable, for example by sums of two squares, infinitely often.

Another interesting instance which involves a double application of the semi-linear sieve arises in the study of primes represented by polynomials over varieties. We give in Theorem 14.10 an example wherein one can derive bounds of the correct order of magnitude provided the level of distribution is sufficiently good. The level needed is not yet known to hold but is weaker than the Elliott–Halberstam conjecture. In addition to its applications to prime number problems the semi-linear sieve proves to be a powerful tool for the study of diophantine problems. We show how this works for cubic surfaces. The arguments exemplify once again that one can accomplish more with an imaginative transformation of sieve results as compared to a direct application of the theorems.

The versatility of the sieve shows itself not only in its adaptation to different mathematical techniques but also in its applicability to differing arithmetic goals. In Chapter 15 we consider a few problems which at first glance may seem quite outside the compass of the sieve.

In the next Chapter 16 we switch from practical aspects to a question of great theoretical importance, the parity problem of the linear sieve. As mentioned earlier, this phenomenon is most clearly visible when examined in the context of the Bombieri sieve and so we follow closely his treatment. This requires considerable generalization of the results of Chapter 3 so as to provide a wider class of test functions. Such general results, when confronted with Selberg-type examples, reveal that one cannot produce primes within the framework of the classical sieve. Actually, the parity phenomenon reveals even further limitations.

Chapter 17 deals with combinatorial identities for primes and related arithmetical functions. The subject received its impetus from Vinogradov who, like Brun, took his inspiration from the Eratosthenes sieve. Although not, strictly speaking, a part of sieve methods the ideas are closely related and the results play an important supporting role in modern sieve theory. The objective of these identities is to create from sums over primes, by replacement, sums of special linear forms and bilinear forms. These identities are truly of a combinatorial nature, as opposed to identities such as Poisson summation, obtained by analytic transformations. It is amazing how such simple transformations, obtained by merely adding and subtracting terms, can be so powerful that they lead to asymptotic formulas for sums over primes in many sequences. The point is that, by decomposing in this fashion, one may apply different types of analytic arguments to different pieces, none of which would be appropriate for all. It is interesting to compare at this point Brun’s decomposition with that of Vinogradov. The latter has the advantage, when it works, of producing asymptotic formulas. Brun’s decomposition, which is also into special linear forms and bilinear forms, has been cleverly chosen so that all the bilinear forms appear with the same sign. Therefore, they can be discarded by positivity and this offers

the possibility of successful bounds (but not asymptotic formulae) even in cases when some of the parts are inaccessible to analysis of any kind.

Brun's decomposition could be regarded as a combinatorial identity, à la Vinogradov, were we not to discard the relevant bilinear forms. In fact, we can and do reclaim some of these terms, getting results beyond the classical framework. Therefore, one can consider modern sieve theory as standing between the two. Vinogradov's original decomposition has over the years been replaced by much simpler variants and in this chapter we give a panorama of identities of this type. Similar in principle, they nevertheless have different features which can make one or the other more convenient in specific cases.

In Chapter 18 we develop some further combinatorial identities, designed for application to specific problems in Chapters 19 and 21. These identities have the added feature that they contain pieces which we shall be unable to properly evaluate, but which fortunately have zero measure in the decomposition so that, for these, a crude sieve upper bound will suffice. The fact that we can put aside these zero measure parts makes the other parts just barely fit to the analysis. Without this feature we would fail. In fact we are, in this situation, very close to the combinatorial identity end of the spectrum but we succeed only by moving the tiniest little bit in the direction of the sieve. We should like to point out that the Fundamental Lemma is essential in this sort of "asymptotic identity". We state the main results of this chapter as formulae for sums over primes in terms of special linear forms, bilinear forms and small terms. To complete the job in any given case one does not need to look back into the machinery but only to estimate the forms. One may consider these identities as an axiomatization of the asymptotic sieve for primes. The linear form part is classical, while the bilinear form part is the new axiom by means of which one breaks the parity barrier. In Theorem 18.6 one sees clearly what is missing in the Bombieri sieve in the case $k = 1$, in other words, the reason it failed for primes. The treatment of the bilinear forms in the asymptotic sieve for primes, no matter what methods are used, ends up with sums over primes or the Möbius function, but without restrictions, and so the theory of L -functions applies. Therefore, the asymptotic sieve for primes produces primes in special sequences because it reduces the question to primes in regular sequences.

In Chapter 19 we apply one of the asymptotic identities of the previous chapter to establish the equidistribution of roots of a quadratic polynomial to prime moduli. What still needs to be done is to verify the axioms, and that is the bulk of the problem. To this end we apply harmonic analysis, this time the spectral theory of automorphic forms which is stronger and more appropriate to this problem than is classical Fourier analysis. This application shows clearly that modern sieve theory does not live on its own. It draws strength from external sources and its ability to adapt to do so is the essence of its current success and the key to its future growth.

In Chapter 20 we give another application of the asymptotic sieve for primes, this time to primes $p = \ell^2 + m^2$ with ℓ in an arbitrary but relatively dense set, such as, for example, the primes. Along the way we need to establish a large sieve inequality for "quadratic Farey points". We shall need this type of inequality again.

In Chapter 21 we present some results concerning the multiplicative structure of polynomials in two variables. These polynomial sequences are rather sparse and sieve theory has only begun to scratch the surface. In the case of the polynomials $X^2 + Y^4$ and $X^3 + 2Y^3$ one is able to produce primes. In our third example $X^2 + Y^6$

only the first difficult step has been successfully treated, namely the divisor function τ_3 . Even for these few successes, the technical details are far too lengthy to include in this book. The sieve aspects for $X^2 + Y^4$ have already been treated in Chapter 18 and we briefly discuss here those for $X^3 + 2Y^3$. We also point out some interesting results established in the course of the proofs.

One may say that the theory of the sieve can be divided into two basic components. The first of these is the search for sieve weights which give results as strong as possible for the main term without causing havoc for the remainder. The main ideas used are partly combinatorial and partly analytic. Any set of sieve weights reduces the problem of finding, say primes, to the problem of finding a very good level of distribution for the congruence sums (and bilinear forms). It is in relation to this second component that an infusion of ideas from many parts of mathematics has proved to be particularly fruitful. In Chapter 22 we demonstrate a number of such achievements. A great challenge is provided by sparse sequences and we concentrate on these, for example, sums of two biquadrates. The problem becomes more acute for a general polynomial in several variables and for varieties of high dimension. Here, it is not even obvious in which order to count the points, less so even to guess which varieties should possess primes and, if so, how many of them. We highlight very recent work of Bourgain, Gamburd and Sarnak which captures much of the essence of these issues and solves some of the problems.

In Chapter 23 we apply the linear sieve to estimate from below the number of primes in a short interval. As in many other places, we do not attempt to present the most advanced results. The point we wish to make is that one can go further by combining the sieve with analytic tools than did the earlier results on this problem which applied zeta-function techniques alone. The extra advantage coming from the sieve is the added flexibility it brings to the theory of Dirichlet polynomials. This flexibility is sufficiently great that we could entirely dispense with the zero-density estimates. As did Hoheisel, we still need a good zero-free region, which is implicitly responsible for allowing us to break the parity barrier of the linear sieve.

In the ensuing Chapter 24 we address the problem of the least prime in an arithmetic progression. This is analogous to the Hoheisel problem but, as is well known, it is much more difficult and requires sophisticated new ideas. Our goal is to prove Linnik's theorem using ideas from sieve theory and as little else as we can get away with. Especially, we succeeded in avoiding any use of the repulsion phenomenon of Deuring-Heilbronn and the log-free zero-density theorems for L -functions which are centrepieces of the previous proofs. That we could do so is largely due to combinatorial arrangements in the sieve decomposition. As in all proofs, the arguments split into two cases. In our sieve treatment the case when the exceptional character $\chi \bmod q$ exists is particularly attractive. Rather than considering the sequence of integers in the progression $a \bmod q$ we apply the sieve to that sequence weighted by $1 * \chi$. In this illusory circumstance our sequence is rather lacunary. One would think that our task would be harder but in fact it is easier. The point is that, by weighting in this fashion, we have implicitly performed a preliminary sieve (of relatively small level) which has killed so many prime divisors that what is left to do in the second step is sieve by a very sparse set of primes. Hence, that becomes a problem of sieve dimension zero. So, no surprise that we succeeded. The arguments involving the exceptional character hold within a quite

general framework and are of a pleasing sieve-theoretic flavor so we examine them further in the final section of the chapter.

Despite our best efforts, in the most famous problems it is still impossible to produce primes. Right from Brun's early work, there have been results producing "almost-primes" and an interest in reducing as far as possible the number of prime factors. Any lower bound for $S(\mathcal{A}, z)$ of the right order gives plenty of these. However, as first noticed by Kuhn, one can do better in reducing the number of factors by attaching to the sequence certain weights of a combinatorial nature which are negative at unwanted elements. There has been considerable development in the art of choosing these weights and the most powerful choices are due to G. Greaves. We use in Chapter 25 a simpler earlier construction due to Richert, which is still quite powerful. We strengthen it further by implementing within it the new shape of the linear sieve having bilinear remainder. Consequently, using harmonic analysis, we succeeded in many cases to get results which surpass not only the earlier ones but, more importantly, even those which are conjectured to be the limit for such a weighted sieve in the traditional setting.

The book ends with two appendices. The first one delivers standard but much-needed asymptotics for sums of multiplicative functions and related results. In some of these asymptotics the main terms involve functions which are most conveniently described as continuous solutions to certain differential-difference equations. There is a large theory concerning such equations and such equations occur in many places in sieve theory and elsewhere. In the second appendix we present what we need.

We had a number of goals in writing this book. Foremost among them, we would like to encourage young people to study the subject. In doing so it would be very helpful, in case they do not already have the relevant background, to have at hand some books on analytic number theory, such as [9], [28], [109], [129], [132]. An inexperienced reader may find it helpful to also consult some classic texts in sieve theory, such as the first few chapters of [79] or of [76]. For experienced readers, we also recommend the article of Selberg [145] as a companion to the more theoretical aspects of our book.

We do intend the book also for more experienced readers, especially from other parts of number theory and mathematics, whether to master the subject or to learn those parts they need for application. We made an effort to make the book a handy reference for theorems and techniques. Experience shows, however, that frequently the theorems available are not precisely in the form that is wanted for a particular application. So, it is important to be in command of the ideas behind the theorems in order to perform the necessary adjustments, sometimes small and sometimes not so small.

A word of warning. A lot of ground is covered and the shortest route is not always chosen. The book is intentionally not written in a linear order. Some topics require tools developed later, some topics are discussed before earlier discoveries and some topics are more difficult than others which are treated later. Don't get discouraged! Although we hope that very few readers will choose to take a minimal route, an exceedingly bare-bones introduction to the subject can be attempted by means of Sections 1.1–1.3, Sections 5.1–5.8, Sections 6.1–6.5 and 6.9, and finally, Sections 7.1, 7.4–7.6 and 7.8–7.12.

On the other hand, we did not intend to make the book encyclopaedic. The subject of sieve methods has undergone tremendous growth in the past few decades,

especially when it comes to applications, many of them highly interesting. Already on their own, Erdős and his collaborators are responsible for dozens of these. Reluctant choices needed to be made, including omission of some personal favourites of the authors. Inevitably, there are instruments which we do not allow to perform in our opera. Nevertheless, we hope that there will be readers who find plenty of interesting points and will devote their time to the further development of the subject. Not least, we hope we have succeeded in conveying to many an appreciation of the music that is the sieve.

* * *

Acknowledgments. Thanks are due to Barbara Miller, Lucile Lo, Ida Bulat, Jemima Merisca and Allison Lee for typing various parts of the manuscript and to Luann Cole for directing the book through the production process. We are grateful to Sergei Gelfand for his unflagging interest over the (approximately five year) period we enjoyed working on the book.

JF was supported in part by NSERC Grant A5123 and, during 2003–2005, by a Killam Research Fellowship. HI was supported in part by NSF Grants DMS-03-01168 and DMS-08-02246 and the University of Toronto during numerous visits. Both authors are grateful to the Banff International Research Station for a Research in Teams Award in August 2007, during which, in beautiful surroundings, some of our favourite tunes were composed.