Prerequisites and Notation

The basic requirements for most of this text are standard introductory graduate courses in algebra, analysis (including Lebesgue integration and complex analysis), and probability. Of course, knowledge and familiarity with basic number theory (for instance, the distribution of primes up to the Bombieri–Vinogradov Theorem) are helpful, but we review in Appendix C all the results that we use. Similarly, Appendix B summarizes the notation and facts from probability theory that are the most important for us.

We will use the following notation:

(1) For subsets Y_1 and Y_2 of an arbitrary set X, we denote by $Y_1 - Y_2$ the difference set, that is, the set of elements $x \in Y_1$ such that $x \notin Y_2$.
(2) A locally compact topological space is always assumed to be separated (i.e., Hausdorff), as in Bourbaki [15].
(3) For a set X, $|X| \in [0, +\infty]$ denotes its cardinal, with $|X| = \infty$ if X is infinite. There is no distinction in this text between the various infinite cardinals.
(4) If X is a set and f, g two complex-valued functions on X, then we write synonymously $f = O(g)$ or $f \ll g$ to say that there exists a constant $C \geqslant 0$ (sometimes called an "implied constant") such that $|f(x)| \leqslant Cg(x)$ for all $x \in X$. Note that this implies that in fact $g \geqslant 0$. We also write $f \asymp g$ to indicate that $f \ll g$ and $g \ll f$.
(5) If X is a topological space, $x_0 \in X$ and f and g are functions defined on a neighborhood of x_0, with $g(x) \neq 0$ for x in a neighborhood of x_0, then we say that $f(x) = o(g(x))$ as $x \to x_0$ if $f(x)/g(x) \to 0$ as $x \to x_0$, and that $f(x) \sim g(x)$ as $x \to x_0$ if $f(x)/g(x) \to 1$.
(6) We write $a \mid b$ for the divisibility relation "a divides b"; we denote by (a, b) the gcd of two integers a and b, and by $[a, b]$ their lcm.

(7) Usually, the variable p will always refer to prime numbers. In particular, a series $\sum_p(\cdots)$ refers to a series over primes (summed in increasing order, in case it is not known to be absolutely convergent), and similarly for a product over primes.

(8) We denote by $\mathbf{F}_p$ the finite field $\mathbf{Z}/p\mathbf{Z}$, for p prime, and more generally by $\mathbf{F}_q$ a finite field with q elements, where $q = p^n$, $n \geqslant 1$, is a power of p. We will recall the properties of finite fields when we require them.

(9) For a complex number z, we write $e(z) = e^{2i\pi z}$. If $q \geqslant 1$ and $x \in \mathbf{Z}/q\mathbf{Z}$, then $e(x/q)$ is well defined by taking any representative of x in $\mathbf{Z}$ to compute the exponential.

(10) If $q \geqslant 1$ and $x \in \mathbf{Z}$ (or $x \in \mathbf{Z}/q\mathbf{Z}$) is an integer that is coprime to q (or a residue class invertible modulo q), we sometimes denote by $\bar{q}$ the inverse class such that $x\bar{x} = 1$ in $\mathbf{Z}/q\mathbf{Z}$. This will always be done in such a way that the modulus q is clear from context, in the case where x is an integer.

(11) Given a probability space $(\Omega, \Sigma, \mathbf{P})$, we denote by $\mathbf{E}(\cdot)$ (resp. $\mathbf{V}(\cdot)$) the expectation (resp. the variance) computed with respect to $\mathbf{P}$. It will often happen that we have a sequence $(\Omega_N, \Sigma_N, \mathbf{P}_N)$ of probability spaces; we will then denote by $\mathbf{E}_N$ or $\mathbf{V}_N$ the respective expectation and variance with respect to $\mathbf{P}_N$.

(12) Given a measure space (Ω, Σ, μ) (not necessarily a probability space), a set Y with a σ-algebra Σ' and a measurable map $f\colon \Omega \longrightarrow Y$, we denote by $f_*(\mu)$ (or sometimes $f(\mu)$) the image measure on Y; in the case of a probability space, so that f is seen as a random variable on Ω, this is the probability law of f seen as a "random Y-valued element." If the set Y is given without specifying a σ-algebra, we will view it usually as given with the σ-algebra generated by sets $Z \subset Y$ such that $f^{-1}(Z)$ belongs to Σ.

(13) As a typographical convention, we will use sans-serif fonts like X to denote an *arithmetic* random variable and more standard fonts (like X) for "abstract" random variables. When using the same letter, this will usually mean that somehow the "purely random" X is the "model" of the arithmetic quantity X.