CHAPTER 2

What is an L-function?

In this chapter we begin by defining, axiomatically, the class of L-functions we will consider in this book. In broad terms, our L-functions should be those that play a role in the Langlands program, i.e., those that are ultimately tied to an automorphic representation. Our axiomatic formulation involves a detailed specification of the potential functional equation and Euler product of an L-function. This approach is in contrast to most other axiomatic approaches to L-functions, such as that of Selberg (see Section 2.4.3), in which the axioms are intentionally made as vague as possible. Since it is conjectured that the set of L-function we define is actually the same as the "Selberg class," one could argue that Selberg's minimalist approach is better. The reason we prefer our approach is that all L-functions actually come from some object which is of its own independent interest, and many L-functions seem to arise from several different objects (e.g., from modular forms, from curves, from surfaces, from automorphic representations, etc.). The particular origin of the L-function imparts many constraints on the L-function, constraints which manifest themselves in properties of the functional equation and Euler product. Since a detailed understanding of the huge spectrum of sources for L-functions is beyond the reach of most mathematicians, our hope is that by first focusing on L-functions as analytic objects we can use that perspective as a step to understanding the myriad of objects which are related to L-functions. We also hope that this approach will help people who are experts on just one approach to the study of L-functions to use these connections to learn about the relationship with other objects, and, in turn, learn something new about objects which are already familiar.

We note that the term "L-function" is used inconsistently in the literature, and some objects which have been called "L-functions" do not, or have not yet been proven to, meet our definition. And some people use "zeta-function" as a synonym for "L-function". In informal discussions these inconsistencies may cause slight confusion, but no more than other cases in mathematics where the same term is used in several contexts. The specific definition does matter when properties of an L-function are used in a proof. In that situation it is important to be specific about exactly what properties of the function are being used in each step.

There are at least three points of view from which one can define an L-function. First, one could define them as always being attached to an automorphic representation. Such an attachment confers certain properties to the L-function, many of which we include in our axioms. We find this approach lacking since many functions which deserve to be called L-functions, such as most Hasse-Weil L-functions, have not been shown to be automorphic. Second, one could define an L-function to be an object that satisfies the Selberg axioms. We find this approach less than ideal because the axioms are so general that it is an unsolved problem to show that they imply the properties we expect an L-function to have. Our approach is a middle path: there is an explicit list of axioms and they correspond to the properties L-functions associated to automorphic forms are expected to have. It is a wide-open problem that will not be solved any time soon that these three points of view are describing the same set of objects.

We begin by stating and describing the axioms. We give names to the properties an L-function should satisfy and state the properties as restrictively as possible. We state the axioms somewhat free of motivation and explanation, two things we provide in the subsequent section. After some comments on the axioms, we discuss what we call the "structural properties" of the L-function. For example, we discuss what the distribution of the coefficients of the L-function tells us about the L-function.

Continuing, we discuss the "analytic properties" of an L-function, including such topics as its functional equation, its poles, its zeros, and its order of growth. Having discussed L-functions from an analytic perspective, we connect them to arithmetic in the following two sections: first, we define an "arithmetic L-function" and, second, we discuss the special and critical values of such an object.

We conclude with three sections in which we discuss (1) moments of L-functions and connect them to subconvexity, (2) operations we can carry out on Euler products and (3) strong multiplicity one, a description of what data we need to specify an L-function exactly.

2.1. L-function axioms

⟨sec:Lfunctionaxioms⟩ An L-function is a function $L(s)$ of the complex variable $s = \sigma + it$, with $\sigma, t \in \mathbb{R}$, satisfying the three sets of axioms listed below. In the statements of the axioms we introduce terminology which later we will show to be well-defined. Also, we warn the reader that we use the so-called "analytic normalization" in our axioms. The "arithmetic normalization" which some people prefer (although it is not defined for all L-functions) is discussed in Section 2.6. All undefined notation is described in later sections.

Axiom 1 **(Analytic properties):** $L(s)$ is given by a Dirichlet series

$$L(s) = \sum_{n=1}^{\infty} \frac{a_n}{n^s}, \tag{2.1.1}$$?eqn:DS?

where $a_n \in \mathbb{C}$.

a) *Convergence:* $L(s)$ converges absolutely for $\sigma > 1$.
b) *Analytic continuation:* $L(s)$ continues to a meromorphic function having only finitely many poles, with all poles [[in $\sigma > 0$]] lying on the $\sigma = 1$ line.

Axiom 2 **(Functional equation):** There is a positive integer N called the *level* of the L-function, a positive integer d called the *degree* of the L-function, a pair of non-negative integers (d_1, d_2) called the *signature* of the L-function, where $d = d_1 + 2d_2$, and complex numbers $\{\mu_j\}$ and $\{\nu_j\}$ called the *spectral parameters* of the L-function, such that the *completed L-function*

$$\Lambda(s) = N^{s/2} \prod_{j=1}^{d_1} \Gamma_{\mathbb{R}}(s + \mu_j) \prod_{j=1}^{d_2} \Gamma_{\mathbb{C}}(s + \nu_j) \cdot L(s) \tag{2.1.2}$$ eqn:Lambda

has the following properties:

a) *Bounded in vertical strips:* Away from the poles of the L-function, $\Lambda(s)$ is bounded in vertical strips $\sigma_1 \le \sigma \le \sigma_2$.

b) *Functional equation:* There exists $\varepsilon \in \mathbb{C}$, called the *sign* of the functional equation, such that

$$\Lambda(s) = \varepsilon\overline{\Lambda}(1-s). \tag{2.1.3}$$ eqn:FE

c) *Partial Selberg bound:* For every j, we have $\mathrm{Re}(\mu_j) > -\frac{1}{2}$ and $\mathrm{Re}(\nu_j) > 0$, and furthermore the sets $\{\mu_j\}$ and $\{\nu_j\}$ satisfy the unitary pairing condition. ›reciseselbergeigenvalue⟩?

Axiom 3 (Euler product): There is a product formula

$$L(s) = \prod_{p \text{ prime}} F_p(p^{-s})^{-1}, \tag{2.1.4}$$?eqn:EP?

absolutely convergent for $\sigma > 1$.

a) *Polynomial:* For every prime p, the F_p is a polynomial of degree at most d, with $F_p(0) = 1$.

b) *Central character:* There exists a Dirichlet character χ mod N, called the *central character* of the L-function, such that

$$F_p(z) = 1 - a_p z + \cdots + (-1)^d \chi(p) p^{-it_0} z^d, \tag{2.1.5}$$ eqn:Fpchi

where $d\,t_0 = \mathrm{Im}\left(\sum \mu_j + 2\sum \nu_j\right)$. [[check that –DF]]

c) *Parity:* The spectral parameters determine the parity of the central character:

$$\chi(-1) = (-1)^{\mathrm{Re}(\sum \mu_j + \sum(2\nu_j+1))}. \tag{2.1.6}$$ eqn:parityaxiom

d) *Partial Ramanujan bound:* Write F_p in factored form as

$$F_p(z) = (1-\alpha_{1,p}z)\cdots(1-\alpha_{d_p,p}z) \tag{2.1.7}$$ {?}

with $\alpha_{j,p} \neq 0$. There exists $\theta < \frac{1}{2}$ such that for every j we have $|\alpha_{j,p}| \le p^\theta$, and furthermore the set $\{\alpha_{j,p}\}$ satisfies the unitary pairing condition.

At first glance, these axioms appear dense with notation and light on motivation, and may contain unfamiliar terminology. In the following sections, we fill in these gaps.

2.1.1. Some notation.

We describe some of the notation used in the axioms.

In Axiom 2, we use $\Gamma_\mathbb{R}$ and $\Gamma_\mathbb{C}$ which are defined as: $\Gamma_\mathbb{R}(s) = \pi^{-s/2}\Gamma(s/2)$ and $\Gamma_\mathbb{C}(s) = 2(2\pi)^{-s}\Gamma(s)$, where Γ is the Euler Γ-function; see Appendix 11.3 for background.

Also in Axiom 2, we write $\overline{\Lambda}(z) = \overline{\Lambda(\overline{z})}$. In general, if $f(z)$ is an analytic function we write $\overline{f}(z)$ to mean $\overline{f(\overline{z})}$, the Schwartz reflection of $f(z)$.

Axioms 2 and 3 refer to the "unitary pairing condition," which is described in Section **??**.

2.2. Comments on the axioms, and some terminology

In this section we make some comments on the axioms. We focus on the most immediate consequences of the axioms, and address the fact that the set of axioms is not minimal.

2.2.1. Comments on Axiom 1 (Analytic properties). We choose to start with a Dirichlet series as the first property of an L-function, because that is how they arose historically, and also from an analytic or computational perspective the Dirichlet series is more natural than the Euler product. One could have started with an Euler product (which of course multiplies out to give a Dirichlet series), which would be a natural starting point if, for example, one wished to focus on automorphic L-functions or those attached to elliptic curves. We have chosen the axioms, and the order they are presented, because in this book we take the L-functionsas analytic objects which exist without reference to another underlying object.

Axiom 1c) In contrast to the axioms presented by Selberg, see Section 2.4.3, we allow poles to occur anywhere on the $\sigma = 1$ line.

2.2.2. Comments on Axiom 2 (Functional equation). We use the notation $\Lambda(s)$ for the completed L-function, which we sometimes write as

$$\Lambda(s) = L_\infty(s)L(s), \tag{2.2.1}$$ {?}

where

$$L_\infty(s) = N^{s/2}\prod_{j=1}^{d_1}\Gamma_\mathbb{R}(s+\mu_j)\prod_{j=1}^{d_2}\Gamma_\mathbb{C}(s+\nu_j). \tag{2.2.2}$$ {?}

In some cases (particularly when discussing automorphic L-functions), people use "L-function" for what we call the "completed L-function," so their $L(s)$ is our $\Lambda(s)$. In that context, our $L(s)$ would be called a "partial L-function." In some cases the sign of the functional equation is incorporated into the (completed) L-function, the factor being the square root of our ε.

The sign of the functional equation satisfies $|\varepsilon| = 1$. To see this, apply the functional equation twice to get $\Lambda(s) = \varepsilon\overline{\varepsilon}\Lambda(s)$. The sign is sometimes called the *root number*.

The duplication formula for the Γ-function implies $\Gamma_\mathbb{C}(s) = \Gamma_\mathbb{R}(s)\Gamma_\mathbb{R}(s+1)$. This may appear to introduce an ambiguity in how the functional equation is written, but in Proposition 2.3.1 we show that all the parameters in the functional equation are well-defined.

Axiom 2c) Our axiom concerns the partial Selberg bound. The (full) Selberg bound is:

Conjecture 2.2.1 (The Selberg bound). *Let*

$$L_\infty(s) = N^{s/2}\prod_{j=1}^{d_1}\Gamma_\mathbb{R}(s+\mu_j)\prod_{j=1}^{d_2}\Gamma_\mathbb{C}(s+\nu_j) \tag{2.2.3}$$ {?}

be the Γ*-factor for an* L*-function. Then* $\mathrm{Re}(\mu_j) \in \{0,1\}$ *and* $\mathrm{Re}(\nu_j) \in \{\frac{1}{2}, 1, \frac{3}{2}, 2, ...\}$.

An alternative name to the Selberg bound, is the *temperedness condition*. The Selberg bound is also known as the Selberg Eigenvalue Conjecture.

The Selberg bound is not known to be a theorem for most automorphic L-functions. The partial Selberg bound, Axiom 2c), does hold for the standard L-function of an [[adjectives]] representation on $\mathrm{GL}(n)$, as we describe in Section **??**. For such L-functions the Selberg bound is conjectured to hold.

It is clear that the bounds $\mathrm{Re}(\mu_j) > -\frac{1}{2}$ and $\mathrm{Re}(\nu_j) > 0$ are a weak version of the Selberg bound. However, just having an inquality on the spectral parameters is not sufficient to apply ideas from [[automorphic representations????]] on $\mathrm{GL}(n)$ to an L-function which does not arise from an underlying object. We require an additional condition which we call the *unitary pairing condition*. This essentially says that if the Selberg bound fails for one Γ-factor, it must also fail for another Γ-factor, in a way that form balanced pairs around the values allowed by the Selberg conjecture. For example, the following Γ-factor satisfies the unitary pairing condition:

$$\begin{aligned}\Gamma_{\mathbb{R}}(s-0.2)\Gamma_{\mathbb{R}}(s+0.2)\Gamma_{\mathbb{R}}(s)^3\Gamma_{\mathbb{R}}(s+0.9)\Gamma_{\mathbb{R}}(s+1.1)&\\ \times\Gamma_{\mathbb{C}}(s+0.7)\Gamma_{\mathbb{C}}(s+1.3)^2\Gamma_{\mathbb{C}}(s+1.7)\Gamma_{\mathbb{C}}(s+7),&\end{aligned}\qquad (2.2.4)\ \{?\}$$

as does this one

$$\begin{aligned}\Gamma_{\mathbb{R}}(s-0.2+3i)\Gamma_{\mathbb{R}}(s+0.2+3i)\Gamma_{\mathbb{R}}(s+1)\Gamma_{\mathbb{R}}(s+1-8i)&\\ \times\Gamma_{\mathbb{C}}(s+0.7)\Gamma_{\mathbb{C}}(s+1.3)\Gamma_{\mathbb{C}}(s+1.3-7i)\Gamma_{\mathbb{C}}(s+1.7-7i).&\end{aligned}\qquad (2.2.5)\ \{?\}$$

We now describe the unitary pairing condition, or more accurately, the *unitary pairing condition at infinity*. In the definition we use the following notation: if $x \in \mathbb{R}$ and $\xi \in \mathbb{C}$ then $(x,\xi)^* = (x, -\overline{\xi})$.

DEFINITION 2.2.2. The multisets $\{\mu_j\}$ and $\{\nu_j\}$ meet the *unitary pairing condition at infinity* if it is possible to write $\mu_j = \delta_j + \alpha_j$ and $\nu_j = \eta_j + \beta_j$, where $\delta_j \in \{0,1\}$ and $\eta_j \in \{\frac{1}{2}, 1, \frac{3}{2}, \ldots\}$, with $|\mathrm{Re}(\alpha_j)|, |\mathrm{Re}(\beta_j)| < \frac{1}{2}$, such that the multisets $\{(\delta_j, \alpha_j)\}$ and $\{(\eta_j, \beta_j)\}$ are closed under the operation $S \to S^*$.

We have written the functional equation in the *analytic normalization*, meaning that the functional equation relates s to $1-s$. This is natural when one is considering the analytic properties of the L-function. There are other cases where it is more natural to consider the *arithmetic normalization*, in which the functional equation relates s to $w+1-s$ for some non-negative integer w, known as the *motivic weight* of the L-function. See Section 2.6 for a discussion.

⟨sec:commentsEuler⟩ **2.2.3. Comments on Axiom 3 (Euler product).** There are some obvious redundancies in the Euler product axioms, for example Axiom 3b) directly implies Axiom 3a). We include Axiom 3a) because it states explicitly that F_p is a polynomial of degree at most the degree of the L-function, two axioms that should be stated explicitly. [[can we give a better explanation? I can't think of one. –DF]]

The convergence of the infinite series $L(s)$ allows us to conclude that the infinite product converges and is equal to $L(s)$. Specifically, the fact that any rearrangement of an absolutely convergent series converges to the same value, tells us that if the Dirichlet series $L(s) = \sum_{n=1}^{\infty} a_n n^{-s}$ converges absolutely at some point s, then the Euler product converges absolutely and is equal to $L(s)$.

The terms in the Euler product, $F_p(p^{-s})^{-1}$, are known as the *local factors* in the Euler product.

Axiom 3b) We say that the prime p is *bad* if $p|N$, and *good* if $p \nmid N$. By (2.1.5), p is good if and only if d_p, the degree of F_p, equals d, the degree of the L-function.

Axiom 3c) By the unitary pairing condition, the exponent $\mathrm{Re}\left(\sum \mu_j + \sum(\nu_j + 1)\right)$ is an integer. So by (2.1.6) we have $\chi(-1) = 1$ or -1, that is, the central character is even or odd.

Axiom 3d) The $\alpha_{j,p}$, which are the reciprocals of the roots of F_p, are called the *Satake parameters* of the L-function at p.

Note that if $L(s)$ arises from an arithmetic object, then for each good prime that object will have its own Satake parameters, which in general will be different than the Satake parameters of the L-function. Indeed, that object will give rise to an infinite list of L-functions, each of which has different Satake parameters.

There is a close analogy between the (partial) Selberg bound and the (partial) Ramanujan bound. Indeed, we have:

CONJECTURE 2.2.3 (The Ramanujan bound). *Write the local factor of an L-function of level N in factored form as*

$$F_p(z) = (1 - \alpha_{1,p}z) \cdots (1 - \alpha_{d_p,p}z) \tag{2.2.6}$$ {?}

with $\alpha_{j,p} \neq 0$. If $p \nmid N$ then $|\alpha_{j,p}| = 1$ for all j. If $p \mid N$ then $|\alpha_{j,p}| = p^{-m_j/2}$ for some $m_j \in \{0, 1, 2, ...\}$, and $\sum m_j \leq d - d_p$.

There is a slick way to write the Ramanujan bound for all p: $|\alpha_{j,p}| = p^{-m_j/2}$ for some $m_j \in \{0, 1, 2, ...\}$, and $\sum m_j \leq d - d_p$. This is the condition above at a bad prime, but it also gives the condition at a good prime, because in that case $d_p = d$, so $m_j = 0$.

One sees a direct analogy between the Ramanujan and the Selberg bound, where the symmetry is with respect to a circle instead of a line. Just as in the archimedean case, the Ramanujan conjecture is also known as the *temperedness condition*. Virtually every statement about the Selberg bound translates directly to a statement about the Ramanujan bound: the Ramanujan bound is not known to be a theorem for most automorphic L-functions. The partial Ramanujan bound, Axiom 3c), holds for the standard L-function of an [[adjectives]] representation on GL(n). For such L-functions the Ramanujan bound is conjectured to hold. Details about these assertions, and a discussion of the connection between the Selberg and Ramanujan bounds, can be found in Section **??**.

Just as in the archimedean case, we see that the partial Ramanujan condition $|\alpha_{j,p}| \leq p^\theta$ is a weak version of the Ramanujan bound. And again, an inequality on the Satake parameters is not sufficient when we want to employ the machinery of automorphic representations on $GL(n)$, so we need a unitary pairing condition. At a good prime the unitary pairing condition is easy to state.

DEFINITION 2.2.4. Suppose p is a good prime. The multiset $\{\alpha_1, \dots, \alpha_d\}$ meets the *unitary pairing condition at p* with partial Ramanujan bound $\theta < \frac{1}{2}$, if $|\alpha_j| \leq p^\theta$ and the multiset is closed under the operation $x \to 1/\overline{x}$. Equivalently, the polynomial $F(z) = \prod_j (1 - \alpha_j z)$ has all its roots in $|z| \geq p^{-\theta}$ and satisfies the self-reciprocal condition

$$F(z) = \xi z^d \overline{F}(z^{-1}), \tag{2.2.7}$$ eqn:self-reciprocal

where $\xi = (-1)^d \prod_j \alpha_j$.

The term *self-reciprocal* refers to the fact that, up to multiplication by a constant, the coefficients of the polynomial are the same if read either order.

If $|\alpha_j| = 1$ then the unitary pairing condition at p says nothing, because $\alpha_j = 1/\overline{\alpha_j}$. But those Satake parameters which are not on the unit circle occur in pairs: if $\alpha_j = re^{i\theta}$ with $r \neq 1$, then $r^{-1}e^{i\theta}$ is also a Satake parameter. Those two points are located symmetrically with respect to the unit circle.

The general case of the unitary pairing condition at p, which includes the good prime version above, closely follows the archimedean case. Specifically, the $\Gamma_{\mathbb{R}}$ factors are like the good primes, and the $\Gamma_{\mathbb{C}}$ factors are similar to the bad primes. The representation-theoretic explanation for this similarity is discussed in Section **??**. [[note to the author of that section: it has to do with whether the irreducible components of the representation are 1- or 2-dimensional.]]

Recall the notation $(x, \xi)^* = (x, -\overline{\xi})$.

DEFINITION 2.2.5. The multiset $\{\alpha_1, \ldots, \alpha_M\}$ meets the *unitary pairing condition at* p of degree d and partial Ramanujan bound $\theta < \frac{1}{2}$, if it is possible to write $\alpha_j = p^{-\eta_j - \beta_j}$ where $\eta_j \in \{0, \frac{1}{2}, 1, \frac{3}{2}, \ldots\}$, with $\sum \eta_j \le d - M$ and $|\mathrm{Re}(\beta_j)| \le \theta$, such that the multiset $S = \{(\eta_j, \beta_j)\}$ is closed under the operation $S \to S^*$.

[[Somewhere say that equality means Iwahori-Spherical]]

[[put somewhere: $\sum \mu_j + \sum (2\nu_j + 1)$ determines the infinity contribution to the sign of the Dirichlet character. (How about at p?) (It is trivial to compute that sum, even if the Selberg bound is not true.)]]

Note that if $\{\alpha_1, \ldots, \alpha_M\}$ are the Satake parameters at a good prime, then $M = d$ and the above condition implies that $\eta_j = 0$.

In the archimedean case we wrote $|\mathrm{Re}(\beta_j)| < \frac{1}{2}$ instead of $|\mathrm{Re}(\beta_j)| \le \theta$ for some $\theta < \frac{1}{2}$. But those conditions are actually equivalent because there are only finitely many spectral parameters.

2.3. Strong multiplicity one

?⟨sec:smoL⟩? If is not obvious that the quanitites mentioned in our L-functions axioms are well defined. Perhaps the duplication formula for the Γ-function introduces some ambiguity, or maybe there is more than one possible level for an L-function? Below we show that there is no ambiguity. Such results fall into the class of *strong multiplicity one* theorems. The idea is that partial information about an L-function uniquely determines the L-function. The first version we give below is fairly weak, meaning that one must assume a lot of partial information. Later versions will require fewer hypotheses.

[[motivation: axioms are well defined]]

[[motivation: we don't need to specify the bad factors – if the L-function exists then the bad factords are uniquely determined.]]

[[maybe include this here:

Note that Axiom 2 c) implies that there is no ambiguity in the definition of $\Lambda(s)$ due to the duplication formula $\Gamma_{\mathbb{C}}(s) = \Gamma_{\mathbb{R}}(s)\Gamma_{\mathbb{R}}(s+1)$.

]]

[[we need to introduce the partial Ramanujan/Selberg bound before the statement of this proposition]]

⟨prop:smoL⟩ PROPOSITION 2.3.1. *Suppose* $L_j(s) = \prod_p F_{p,j}(p^{-s})^{-1}$, *for* $j = 1, 2$, *are* L-*functions, meaning that they satisfy the three sets of axioms. If* $F_{p,1} = F_{p,2}$ *for all but finitely many* p, *then* $F_{p,1} = F_{p,2}$ *for all* p, *and* L_1 *and* L_2 *have the same functional equation data.*

A useful special case is: the functional equation data and the bad factors of an L-function are determined by the good local factors, although the proof is not constructive. The proposition shows in particular that the functional equation data