

Chapter 8

Additional Topics

The foundation in analysis provided by the first seven chapters is sufficient background for the exploration of some advanced and historically important topics. The writing in this chapter is similar to that in the concluding project sections of each individual chapter. Exercises are included within the exposition and are designed to make each section a narrative investigation into a significant achievement in the field of analysis.

8.1 The Generalized Riemann Integral

Chapter 7 concluded with Henri Lebesgue's elegant result that a bounded function is Riemann-integrable if and only if its points of discontinuity form a set of measure zero. To eliminate the dependence of integrability on continuity, Lebesgue proposed a new method of integration that has become the standard integral in mathematics. In the Epilogue to Chapter 7, we briefly outlined some of the strengths and weaknesses of the Lebesgue integral, concluding with a look back to the Fundamental Theorem of Calculus (Theorem 7.5.1). (Lebesgue's measure-zero criterion is not a prerequisite for understanding the material in this section, but the discussion in Section 7.7 provides some useful context for what follows.)

If F is a differentiable function on $[a, b]$, then in a perfect world we might hope to prove that

$$(1) \quad \int_a^b F' = F(b) - F(a).$$

Notice that although this is the conclusion of part (i) of Theorem 7.5.1, there we needed the additional requirement that F' be Riemann-integrable. To drive this point home, Section 7.6 concluded with an example of a function that has

a derivative that the Riemann integral cannot handle. The Lebesgue integral alluded to earlier is a significant improvement. It can integrate our example from Section 7.6, but ultimately it too suffers from the same setback. Not every derivative is integrable, no matter which integral is used.

What follows is a short introduction to the generalized Riemann integral, discovered independently around 1960 by Jaroslav Kurzweil and Ralph Henstock. As mentioned in Section 7.7, this lesser-known modification of the Riemann integral can actually integrate a larger class of functions than Lebesgue's ubiquitous integral and yields a surprisingly simple proof of equation (1) above with no additional hypotheses.

The Riemann Integral as a Limit

Let

$$P = \{x_0, x_1, x_2, \dots, x_n\}$$

be a partition of $[a, b]$. A *tagged partition* is one where in addition to P we have chosen points c_k in each of the subintervals $[x_{k-1}, x_k]$. This sets the stage for the concept of a Riemann sum. Given a function $f : [a, b] \rightarrow \mathbf{R}$, and a tagged partition $(P, \{c_k\}_{k=1}^n)$, the *Riemann sum* generated by this partition is given by

$$R(f, P) = \sum_{k=1}^n f(c_k)(x_k - x_{k-1}).$$

Looking back at the definition of the upper sum

$$U(f, P) = \sum_{k=1}^n M_k(x_k - x_{k-1}) \quad \text{where} \quad M_k = \sup\{f(x) : x \in [x_{k-1}, x_k]\},$$

and the lower sum

$$L(f, P) = \sum_{k=1}^n m_k(x_k - x_{k-1}) \quad \text{where} \quad m_k = \inf\{f(x) : x \in [x_{k-1}, x_k]\},$$

it should be clear that

$$L(f, P) \leq R(f, P) \leq U(f, P)$$

for any bounded function f . In Definition 7.2.7, we characterized integrability by insisting that the infimum of the upper sums equal the supremum of the lower sums. Any Riemann sum is going to fall between a particular upper and lower sum. If the upper and lower sums are converging to some common value, then the Riemann sums are also eventually close to this value as well. The next theorem shows that it is possible to characterize Riemann integrability in a way equivalent to Definition 7.2.7 using an ϵ - δ -type definition applied to Riemann sums.

Definition 8.1.1. Let $\delta > 0$. A partition P is δ -fine if every subinterval $[x_{k-1}, x_k]$ satisfies $x_k - x_{k-1} < \delta$. In other words, every subinterval has width less than δ .

Theorem 8.1.2 (Limit Criterion for Riemann Integrability). A bounded function $f : [a, b] \rightarrow \mathbf{R}$ is Riemann-integrable with

$$\int_a^b f = A$$

if and only if, for every $\epsilon > 0$, there exists a $\delta > 0$ such that, for any tagged partition $(P, \{c_k\})$ that is δ -fine, it follows that

$$|R(f, P) - A| < \epsilon.$$

Before attempting the proof, we should point out that, in some treatments, the criterion in Theorem 8.1.2 is actually taken as the *definition* of Riemann integrability. In fact, this is how Riemann originally defined the concept. The spirit of this theorem is close to what is taught in most introductory calculus courses. To approximate the area under the curve, Riemann sums are constructed. The hope is that as the partitions become finer, the corresponding approximations get closer to the value of the integral. The content of Theorem 8.1.2 is that if the function is integrable, then these approximations do indeed converge to the value of the integral, regardless of how the tags are chosen. Conversely, if the approximating Riemann sums for finer and finer partitions collect around some value A , then the function is integrable and integrates to A .

Proof. (\Rightarrow) For the forward direction, we begin with the assumption that f is integrable on $[a, b]$. Given an $\epsilon > 0$, we must produce a $\delta > 0$ such that if $(P, \{c_k\})$ is any tagged partition that is δ -fine, then $|R(f, P) - \int_a^b f| < \epsilon$.

Because f is integrable, we know there exists a partition P_ϵ such that

$$U(f, P_\epsilon) - L(f, P_\epsilon) < \frac{\epsilon}{3}.$$

Let $M > 0$ be a bound on $|f|$, and let n be the number of subintervals of P_ϵ (so that P_ϵ really consists of $n + 1$ points in $[a, b]$). We will argue that choosing

$$\delta = \epsilon/9nM$$

has the desired property.

Here is the idea. Let $(P, \{c_k\})$ be an arbitrary tagged partition of $[a, b]$ that is δ -fine, and let $P' = P \cup P_\epsilon$. The key is to establish the string of inequalities

$$L(f, P') - \frac{\epsilon}{3} < L(f, P) \leq U(f, P) < U(f, P') + \frac{\epsilon}{3}.$$

Exercise 8.1.1. (a) Explain why both the Riemann sum $R(f, P)$ and $\int_a^b f$ fall between $L(f, P)$ and $U(f, P)$.

(b) Explain why $U(f, P') - L(f, P') < \epsilon/3$.

By the previous exercise, if we can show $U(f, P) < U(f, P') + \epsilon/3$ (and similarly $L(f, P') - \epsilon/3 < L(f, P)$), then it will follow that

$$\left| R(f, P) - \int_a^b f \right| < \epsilon$$

and the proof will be done. Thus, we turn our attention toward estimating the distance between $U(f, P)$ and $U(f, P')$.

Exercise 8.1.2. Explain why $U(f, P) - U(f, P') \geq 0$.

A typical term in either $U(f, P)$ or $U(f, P')$ has the form $M_k(x_k - x_{k-1})$, where M_k is the supremum of f over $[x_{k-1}, x_k]$. A good number of these terms appear in both upper sums and so cancel out.

Exercise 8.1.3. (a) In terms of n , what is the largest number of terms of the form $M_k(x_k - x_{k-1})$ that could appear in one of $U(f, P)$ or $U(f, P')$ but not the other?

(b) Finish the proof in this direction by arguing that

$$U(f, P) - U(f, P') < \epsilon/3.$$

(\Leftarrow) For this direction, we assume that the ϵ - δ criterion in Theorem 8.1.2 holds and argue that f is integrable. Integrability, as we have defined it, depends on our ability to choose partitions for which the upper sums are close to the lower sums. We have remarked that given any partition P , it is always the case that

$$L(f, P) \leq R(f, P) \leq U(f, P)$$

no matter which tags are chosen to compute $R(f, P)$.

Exercise 8.1.4. (a) Show that if f is continuous, then it is possible to pick tags $\{c_k\}_{k=1}^n$ so that

$$R(f, P) = U(f, P).$$

Similarly, there are tags for which $R(f, P) = L(f, P)$ as well.

(b) If f is not continuous, it may not be possible to find tags for which $R(f, P) = U(f, P)$. Show, however, that given an arbitrary $\epsilon > 0$, it is possible to pick tags for P so that

$$U(f, P) - R(f, P) < \epsilon.$$

The analogous statement holds for lower sums.

Exercise 8.1.5. Use the results of the previous exercise to finish the proof of Theorem 8.1.2. \square

Gauges and $\delta(x)$ -fine Partitions

The key to the generalized Riemann integral is to allow the δ in Theorem 8.1.2 to be a *function of x* .

Definition 8.1.3. A function $\delta : [a, b] \rightarrow \mathbf{R}$ is called a *gauge* on $[a, b]$ if $\delta(x) > 0$ for all $x \in [a, b]$.

Definition 8.1.4. Given a particular gauge $\delta(x)$, a tagged partition $(P, \{c_k\}_{k=1}^n)$ is *$\delta(x)$ -fine* if every subinterval $[x_{k-1}, x_k]$ satisfies $x_k - x_{k-1} < \delta(c_k)$. In other words, each subinterval $[x_{k-1}, x_k]$ has width less than $\delta(c_k)$.

It is important to see that if $\delta(x)$ is a constant function, then Definition 8.1.4 says precisely the same thing as Definition 8.1.1. In the case where $\delta(x)$ is not a constant, Definition 8.1.4 describes a way of measuring the fineness of partitions that is quite different.

Exercise 8.1.6. Consider the interval $[0, 1]$.

- (a) If $\delta(x) = 1/9$, find a $\delta(x)$ -fine tagged partition of $[0, 1]$. Does the choice of tags matter in this case?
- (b) Let

$$\delta(x) = \begin{cases} 1/4 & \text{if } x = 0 \\ x/3 & \text{if } 0 < x \leq 1. \end{cases}$$

Construct a $\delta(x)$ -fine tagged partition of $[0, 1]$.

The tinkering required in Exercise 8.1.6 (b) may cast doubt on whether an arbitrary gauge always admits a $\delta(x)$ -fine partition. However, it is not too difficult to show that this is indeed the case.

Theorem 8.1.5. *Given a gauge $\delta(x)$ on an interval $[a, b]$, there exists a tagged partition $(P, \{c_k\}_{k=1}^n)$ that is $\delta(x)$ -fine.*

Proof. Let $I_0 = [a, b]$. It may be possible to find a tag such that the trivial partition $P = \{a, b\}$ works. Specifically, if $b - a < \delta(x)$ for some $x \in [a, b]$, then we can set c_1 equal to such an x and notice that $(P, \{c_1\})$ is $\delta(x)$ -fine. If no such x exists, then bisect $[a, b]$ into two equal halves.

Exercise 8.1.7. Finish the proof of Theorem 8.1.5. □

Generalized Riemann Integrability

Keeping in mind that Theorem 8.1.2 offers an equivalent way to define Riemann integrability, we now propose a new method for defining the value of the integral.

Definition 8.1.6. A function f on $[a, b]$ has *generalized Riemann integral* A if, for every $\epsilon > 0$, there exists a gauge $\delta(x)$ on $[a, b]$ such that for each tagged partition $(P, \{c_k\}_{k=1}^n)$ that is $\delta(x)$ -fine, it is true that

$$|R(f, P) - A| < \epsilon.$$

In this case, we write $A = \int_a^b f$.

Theorem 8.1.7. *If a function has a generalized Riemann integral, then the value of the integral is unique.*

Proof. Assume that a function f has generalized Riemann integral A_1 and that it also has generalized Riemann integral A_2 . We must prove $A_1 = A_2$. □

Exercise 8.1.8. Finish the argument.

The implications of Definition 8.1.6 on the resulting class of integrable functions are far reaching. This is somewhat surprising given that the criteria for integrability in Definition 8.1.6 and Theorem 8.1.2 differ in such a small way. One observation that should be immediately evident is the following.

Exercise 8.1.9. Explain why every function that is Riemann-integrable with $\int_a^b f = A$ must also have generalized Riemann integral A .

The converse statement is not true, and that is the important point. One example that we have of a non-Riemann-integrable function is Dirichlet's function

$$g(x) = \begin{cases} 1 & \text{if } x \in \mathbf{Q} \\ 0 & \text{if } x \notin \mathbf{Q} \end{cases}$$

which has discontinuities at every point of \mathbf{R} .

Theorem 8.1.8. *Dirichlet's function $g(x)$ is generalized Riemann-integrable on $[0, 1]$ with $\int_0^1 g = 0$.*

Proof. Let $\epsilon > 0$. By Definition 8.1.6, we must construct a gauge $\delta(x)$ on $[0, 1]$ such that whenever $(P, \{c_k\}_{k=1}^n)$ is a $\delta(x)$ -fine tagged partition, it follows that

$$0 \leq \sum_{k=1}^n g(c_k)(x_k - x_{k-1}) < \epsilon.$$

The gauge represents a restriction on the size of $\Delta x_k = x_k - x_{k-1}$ in the sense that $\Delta x_k < \delta(c_k)$. The Riemann sum consists of products of the form $g(c_k)\Delta x_k$. Thus, for irrational tags, there is nothing to worry about because $g(c_k) = 0$ in this case. Our task is to make sure that any time a tag c_k is rational, it comes from a suitably thin subinterval.

Let $\{r_1, r_2, r_3, \dots\}$ be an enumeration of the countable set of rational numbers contained in $[0, 1]$. For each r_k , set $\delta(r_k) = \epsilon/2^{k+1}$. For x irrational, set $\delta(x) = 1$.

Exercise 8.1.10. Show that if $(P, \{c_k\}_{k=1}^n)$ is a $\delta(x)$ -fine tagged partition, then $R(g, P) < \epsilon$. \square

Dirichlet's function fails to be Riemann-integrable because, given any (un-tagged) partition, it is possible to make $R(g, P) = 1$ or $R(g, P) = 0$ by choosing the tags to be either all rational or all irrational. For the generalized Riemann integral, choosing all rational tags results in a tagged partition that is not $\delta(x)$ -fine (when $\delta(x)$ is small on rational points) and so does not have to be considered. In general, allowing for nonconstant gauges allows us to be more discriminating about which tagged partitions qualify as $\delta(x)$ -fine. The result, as we have just seen, is that it may be easier to achieve the inequality

$$|R(f, P) - A| < \epsilon$$

for the often smaller and more carefully selected set of tagged partitions that remain.

The Fundamental Theorem of Calculus

We conclude this brief introduction to the generalized Riemann integral with a proof of the Fundamental Theorem of Calculus. As was alluded to earlier, the most notable distinction between the following theorem and part (i) of Theorem 7.5.1 is that here we do not need to assume that the derivative function is integrable. Using the generalized Riemann integral, every derivative is integrable, and the integral can be evaluated using the antiderivative in the familiar way. It is also interesting to note that in Theorem 7.5.1 the Mean Value Theorem played the crucial role in the argument, but it is not needed here.

Theorem 8.1.9. *Assume $F : [a, b] \rightarrow \mathbf{R}$ is differentiable at each point in $[a, b]$ and set $f(x) = F'(x)$. Then, f has the generalized Riemann integral*

$$\int_a^b f = F(b) - F(a).$$

Proof. Let $P = \{x_0, x_1, x_2, \dots, x_n\}$ be a partition of $[a, b]$. Both this proof and the proof of Theorem 7.5.1 make use of the following fact.

Exercise 8.1.11. Show that

$$F(b) - F(a) = \sum_{k=1}^n [F(x_k) - F(x_{k-1})].$$

If $\{c_k\}_{k=1}^n$ is a set of tags for P , then we can estimate the difference between the Riemann sum $R(f, P)$ and $F(b) - F(a)$ by

$$\begin{aligned} |F(b) - F(a) - R(f, P)| &= \left| \sum_{k=1}^n [F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})] \right| \\ &\leq \sum_{k=1}^n |F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})|. \end{aligned}$$

Let $\epsilon > 0$. To prove the theorem, we must construct a gauge $\delta(c)$ such that

$$(2) \quad |F(b) - F(a) - R(f, P)| < \epsilon$$

for all $(P, \{c_k\})$ that are $\delta(c)$ -fine. (Using the variable c in the gauge function is more convenient than x in this case.)

Exercise 8.1.12. For each $c \in [a, b]$, explain why there exists a $\delta(c) > 0$ (a $\delta > 0$ depending on c) such that

$$\left| \frac{F(x) - F(c)}{x - c} - f(c) \right| < \epsilon$$

for all $0 < |x - c| < \delta(c)$.

This $\delta(c)$ is the desired gauge on $[a, b]$. Let $(P, \{c_k\}_{k=1}^n)$ be a $\delta(c)$ -fine partition of $[a, b]$. It just remains to show that equation (2) is satisfied for this tagged partition.

Exercise 8.1.13. (a) For a particular $c_k \in [x_{k-1}, x_k]$ of P , show that

$$|F(x_k) - F(c_k) - f(c_k)(x_k - c_k)| < \epsilon(x_k - c_k)$$

and

$$|F(c_k) - F(x_{k-1}) - f(c_k)(c_k - x_{k-1})| < \epsilon(c_k - x_{k-1}).$$

(b) Now, argue that

$$|F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})| < \epsilon(x_k - x_{k-1}),$$

and use this fact to complete the proof of the theorem. \square

If we consider the function

$$F(x) = \begin{cases} x^{3/2} \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

then it is not too difficult to show that F is differentiable everywhere, including $x = 0$, with

$$F'(x) = \begin{cases} (3/2)\sqrt{x} \sin(1/x) - (1/\sqrt{x}) \cos(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

What is notable here is that the derivative is unbounded near the origin. The theory of the ordinary Riemann integral begins with the assumption that we only consider bounded functions on closed intervals, but there is no such restriction for the generalized Riemann integral. Theorem 8.1.9 proves that F' has a generalized integral. Now, *improper* Riemann integrals have been created to extend Riemann integration to some unbounded functions, but it is another interesting fact about the generalized Riemann integral that any function having an improper integral must already be integrable in the sense described in Definition 8.1.6.

As a parting gesture, let's show how Theorem 8.1.9 yields a short verification of the substitution technique from calculus.

Theorem 8.1.10 (Change-of-variable Formula). *Let $g : [a, b] \rightarrow \mathbf{R}$ be differentiable at each point of $[a, b]$, and assume F is differentiable on the set $g([a, b])$. If $f(x) = F'(x)$ for all $x \in g([a, b])$, then*

$$\int_a^b (f \circ g) \cdot g' = \int_{g(a)}^{g(b)} f.$$

Proof. The hypothesis of the theorem guarantees that the function $(F \circ g)(x)$ is differentiable for all $x \in [a, b]$.

Exercise 8.1.14. (a) Why are we sure that f and $(F \circ g)'$ have generalized Riemann integrals?

(b) Use Theorem 8.1.9 to finish the proof. □

The impressive properties of the generalized Riemann integral do not end here. The central source for the material in this section is Robert Bartle's award winning article "Return to the Riemann Integral," which appeared in the *American Mathematical Monthly*, October, 1996. The article goes on to discuss convergence theorems for this new integral in the spirit of Theorem 7.4.4, and outlines the argument that the collection of integrable functions is strictly larger when the Lebesgue integral is replaced by the generalized Riemann integral. In light of this, the author boldly declares that "*the time has come to discard the Lebesgue integral as the primary integral.*" (Italics in the original.)

That this revolution has not come to pass may simply be due to a case of overwhelming inertia, but a contributing factor is very likely the geometrically satisfying intuition of Lebesgue's theory. At the heart of Lebesgue's approach to integration is the desire to generalize the concepts of length and area. Although one can certainly use a properly developed integral to give a rigorous definition for the length—or measure—of a general set, there is a compelling argument that this puts the ideas in the wrong pedagogical order. Rather than using a sophisticated integral to generalize a primitive notion such as length, Lebesgue found an effective way to talk about the length of a very wide class of sets, and used that to build his definition of the integral. The very elegant result of his endeavor is likely to be the industry standard for a long time to come.

8.2 Metric Spaces and the Baire Category Theorem

A natural question to ask is whether the theorems we have proved about sequences, series, and functions in \mathbf{R} have analogues in the plane \mathbf{R}^2 or in even higher dimensions. Looking back over the proofs, one crucial observation is that most of the arguments depend on just a few basic properties of the absolute value function. Interpreting the statement “ $|x - y|$ ” to mean the “distance from x to y in \mathbf{R} ,” our aim is to experiment with other ways of measuring distance on other sets such as \mathbf{R}^2 and $C[0, 1]$, the space of continuous functions on $[0, 1]$.

Definition 8.2.1. Given a set X , a function $d : X \times X \rightarrow \mathbf{R}$ is a *metric* on X if for all $x, y \in X$:

- (i) $d(x, y) \geq 0$ with $d(x, y) = 0$ if and only if $x = y$,
- (ii) $d(x, y) = d(y, x)$, and
- (iii) for all $z \in X$, $d(x, y) \leq d(x, z) + d(z, y)$.

A *metric space* is a set X together with a metric d .

Property (iii) in the previous definition is the “triangle inequality.” The next two exercises illustrate the point that the same set X can be home to several different metrics. When referring to a metric space, we must specify the set *and* the particular distance function d .

Exercise 8.2.1. Decide which of the following are metrics on $X = \mathbf{R}^2$. For each, we let $x = (x_1, x_2)$ and $y = (y_1, y_2)$ be points in the plane.

- (a) $d(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$.
- (b) $d(x, y) = \max\{|x_1 - y_1|, |x_2 - y_2|\}$.
- (c) $d(x, y) = |x_1x_2 + y_1y_2|$.

The metric in part (a) of the previous exercise is the familiar Euclidean distance between two points in the plane. This is often referred to as the “usual” or “standard” metric on \mathbf{R}^2 . The usual metric on \mathbf{R} is our old friend $d(x, y) = |x - y|$.

Exercise 8.2.2. Let $C[0, 1]$ be the collection of continuous functions on the closed interval $[0, 1]$. Decide which of the following are metrics on $C[0, 1]$.

- (a) $d(f, g) = \sup\{|f(x) - g(x)| : x \in [0, 1]\}$.
- (b) $d(f, g) = |f(1) - g(1)|$.
- (c) $d(f, g) = \int_0^1 |f - g|$.

The following distance function is called the *discrete metric* and can be defined on any set X . For any $x, y \in X$, let

$$\rho(x, y) = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

Exercise 8.2.3. Verify that the discrete metric is actually a metric.

Basic Definitions

Definition 8.2.2. Let (X, d) be a metric space. A sequence $(x_n) \subseteq X$ *converges* to an element $x \in X$ if for all $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that $d(x_n, x) < \epsilon$ whenever $n \geq N$.

Definition 8.2.3. A sequence (x_n) in a metric space (X, d) is a *Cauchy sequence* if for all $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that $d(x_m, x_n) < \epsilon$ whenever $m, n \geq N$.

Exercise 8.2.4. Show that a convergent sequence is Cauchy.

The Cauchy Criterion, as it is called in \mathbf{R} , was an “if and only if” statement. In the general metric space setting, however, the converse statement does not always hold. Recall that, in \mathbf{R} , the assertion that “Cauchy sequences converge” was shown to be equivalent to the Axiom of Completeness. In order to transport the Axiom of Completeness into a metric space, we would need to have an ordering on our space so that we could discuss such things as upper bounds. It is an interesting observation that not every set can be ordered in a satisfying way (the points in \mathbf{R}^2 for example). Even without an ordering, we are still going to want completeness. For metric spaces, the convergence of Cauchy sequences is taken to be the definition of completeness.

Definition 8.2.4. A metric space (X, d) is *complete* if every Cauchy sequence in X converges to an element of X .

Exercise 8.2.5. (a) Consider \mathbf{R}^2 with the discrete metric $\rho(x, y)$ examined in Exercise 8.2.3. What do Cauchy sequences look like in this space? Is \mathbf{R}^2 complete with respect to this metric?

(b) Show that $C[0, 1]$ is complete with respect to the metric in Exercise 8.2.2 (a).

(c) Define $C^1[0, 1]$ to be the collection of differentiable functions on $[0, 1]$ whose derivatives are also continuous. Is $C^1[0, 1]$ complete with respect to the metric defined in Exercise 8.2.2 (a)?

Because completeness is a prerequisite for doing anything significant in the way of analysis, the metric in Exercise 8.2.2 (a) is the most natural metric to consider when working with $C[0, 1]$. The notation

$$\|f - g\|_\infty = d(f, g) = \sup\{|f(x) - g(x)| : x \in [0, 1]\}$$

is standard, and setting $g = 0$ gives the so-called “sup norm”

$$\|f\|_\infty = d(f, 0) = \sup\{|f(x)| : x \in [0, 1]\}.$$

In all upcoming discussions, it is assumed that the space $C[0, 1]$ is endowed with this metric unless otherwise specified.

Definition 8.2.5. Let (X, d_1) and (Y, d_2) be metric spaces. A function $f : X \rightarrow Y$ is *continuous at* $x \in X$ if for all $\epsilon > 0$ there exists a $\delta > 0$ such that $d_2(f(x), f(y)) < \epsilon$ whenever $d_1(x, y) < \delta$.

Exercise 8.2.6. Which of these functions from $C[0, 1]$ to \mathbf{R} (with the usual metric) are continuous?

(a) $g(f) = \int_0^1 fk$, where k is some fixed function in $C[0, 1]$.

(b) $g(f) = f(1/2)$.

(c) $g(f) = f(1/2)$, but this time with respect to the metric on $C[0, 1]$ from Exercise 8.2.2 (c).

Topology on Metric Spaces

Definition 8.2.6. Given $\epsilon > 0$ and an element x in the metric space (X, d) , the ϵ -neighborhood of x is the set $V_\epsilon(x) = \{y \in X : d(x, y) < \epsilon\}$.

Exercise 8.2.7. Describe the ϵ -neighborhoods in \mathbf{R}^2 for each of the different metrics described in Exercise 8.2.1. How about for the discrete metric?

With the definition of an ϵ -neighborhood, we can now define *open sets*, *limit points*, and *closed sets* exactly as we did before. A set $O \subseteq X$ is *open* if for every $x \in O$ we can find a neighborhood $V_\epsilon(x) \subseteq O$. A point x is a *limit point* of a set A if every $V_\epsilon(x)$ intersects A in some point other than x . A set C is *closed* if it contains its limit points.

Exercise 8.2.8. Let (X, d) be a metric space.

(a) Verify that a typical ϵ -neighborhood $V_\epsilon(x)$ is an open set. Is the set

$$C_\epsilon(x) = \{y \in X : d(x, y) \leq \epsilon\}$$

a closed set?

(b) Show that a set $E \subseteq X$ is open if and only if its complement is closed.

Exercise 8.2.9. (a) Show that the set $Y = \{f \in C[0, 1] : \|f\|_\infty \leq 1\}$ is closed in $C[0, 1]$.

(b) Is the set $T = \{f \in C[0, 1] : f(0) = 0\}$ open, closed, or neither in $C[0, 1]$?

We define compactness in metric spaces just as we did for \mathbf{R} .

Definition 8.2.7. A subset K of a metric space (X, d) is *compact* if every sequence in K has a convergent subsequence that converges to a limit in K .

An extremely useful characterization of compactness in \mathbf{R} is the proposition that a set is compact if and only if it is closed and bounded. For abstract metric spaces, this proposition only holds in the forward direction.

Exercise 8.2.10. (a) Supply a definition for *bounded* subsets of a metric space (X, d) .

(b) Show that if K is a compact subset of the metric space (X, d) , then K is closed and bounded.

(c) Show that $Y \subseteq C[0, 1]$ from Exercise 8.2.9 (a) is closed and bounded but not compact.

A good hint for part (c) of the previous exercise can be found in Exercise 6.2.14 from Chapter 6. This exercise defines the concept of an *equicontinuous* family of functions, which is a key ingredient in the Arzela–Ascoli Theorem (Exercise 6.2.15). The Arzela–Ascoli Theorem states that any bounded, *equicontinuous* collection of functions in $C[0, 1]$ must have a uniformly convergent subsequence. One way to summarize this famous result—which we did not have the language for in Chapter 6—is as a statement describing a particular class of compact subsets in $C[0, 1]$. Looking at the definition of compactness, and remembering that the uniform limit of continuous functions is continuous, the Arzela–Ascoli Theorem states that any closed, bounded, equicontinuous collection of functions is a compact subset of $C[0, 1]$.

Definition 8.2.8. Given a subset E of a metric space (X, d) , the *closure* \overline{E} is the union of E together with its limit points. The *interior* of E is denoted by E° and is defined as

$$E^\circ = \{x \in E : \text{there exists } V_\epsilon(x) \subseteq E\}.$$

Closure and interior are dual concepts. Results about these concepts come in pairs and exhibit an elegant and useful symmetry.

Exercise 8.2.11. (a) Show that E is closed if and only if $\overline{E} = E$. Show that E is open if and only if $E^\circ = E$.

(b) Show that $\overline{E^c} = (E^\circ)^c$, and similarly that $(E^\circ)^c = \overline{E^c}$.

A good hint for this exercise is to review the proofs from Chapter 3, where closure at least is discussed. Thinking of all of these concepts as they relate to \mathbf{R} or \mathbf{R}^2 with the usual metric is not a bad idea. However, it is important to remember also that rigorous proofs must be constructed purely from the relevant definitions.

Exercise 8.2.12. (a) Show

$$\overline{V_\epsilon(x)} \subseteq \{y \in X : d(x, y) \leq \epsilon\},$$

in an arbitrary metric space (X, d) .

- (b) To keep things from sounding too familiar, find an example of a specific metric space where

$$\overline{V_\epsilon(x)} \neq \{y \in X : d(x, y) \leq \epsilon\}.$$

We are on our way toward the Baire Category Theorem. The next definitions provide the final bit of vocabulary needed to state the result.

Definition 8.2.9. A set $A \subseteq X$ is *dense* in the metric space (X, d) if $\overline{A} = X$. A subset E of a metric space (X, d) is *nowhere-dense* in X if \overline{E}° is empty.

Exercise 8.2.13. If E is a subset of a metric space (X, d) , show that E is nowhere-dense in X if and only if \overline{E}^c is dense in X .

The Baire Category Theorem

In Section 3.5, we proved Baire's Theorem, which states that it is impossible to write the real numbers \mathbf{R} as the countable union of nowhere-dense sets. Previous to this, we knew that \mathbf{R} was too big to be written as the countable union of single points (\mathbf{R} is uncountable), but Baire's Theorem improves on this by asserting that the only way to make \mathbf{R} from a countable union of arbitrary sets is for the closure of at least one of these sets to contain an interval. The keystone to the proof of Baire's Theorem is the completeness of \mathbf{R} . The idea now is to replace \mathbf{R} with an arbitrary complete metric space and prove the theorem in this more general setting. This leads to a statement that can be used to discuss the size and structure of other spaces such as \mathbf{R}^2 and $C[0, 1]$. At the end of Chapter 3, we mentioned one particularly fascinating implication of this result for $C[0, 1]$, which is that—despite the substantial difficulty required to produce an example of one—most continuous functions are nowhere-differentiable. It would be a good idea at this point to reread Sections 3.6 and 5.5. We are now equipped to carry out the details promised in these discussions.

Theorem 8.2.10. *Let (X, d) be a complete metric space, and let $\{O_n\}$ be a countable collection of dense, open subsets of X . Then, $\bigcap_{n=1}^\infty O_n$ is not empty.*

Proof. When we proved this theorem on \mathbf{R} , completeness manifested itself in the form of the Nested Interval Property. We could derive something akin to NIP in the metric space setting, but instead let's take an approach that uses the convergence of Cauchy sequences (because this is how we have defined completeness).

Pick $x_1 \in O_1$. Because O_1 is open, there exists an $\epsilon_1 > 0$ such that $V_{\epsilon_1}(x_1) \subseteq O_1$.

Exercise 8.2.14. (a) Give the details for why we know there exists a point $x_2 \in V_{\epsilon_1}(x_1) \cap O_2$ and an $\epsilon_2 > 0$ satisfying $\epsilon_2 < \epsilon_1/2$ with $V_{\epsilon_2}(x_2)$ contained in O_2 and

$$\overline{V_{\epsilon_2}(x_2)} \subseteq V_{\epsilon_1}(x_1).$$

- (b) Proceed along this line and use the completeness of (X, d) to produce a single point $x \in O_n$ for every $n \in \mathbf{N}$. \square

Theorem 8.2.11 (Baire Category Theorem). *A complete metric space is not the union of a countable collection of nowhere-dense sets.*

Exercise 8.2.15. Complete the proof of the theorem.

This result is called the Baire Category Theorem because it creates two categories of size for subsets in a metric space. A set of “first category” is one that *can* be written as a countable union of nowhere-dense sets. These are the small, intuitively thin subsets of a metric space. We now see that if our metric space is complete, then it is necessarily of “second category,” meaning it cannot be written as a countable union of nowhere-dense sets. Given a subset A of a complete metric space X , showing that A is of first category is a mathematically precise way of demonstrating that A constitutes a very minor portion of the set X . The term “meager” is often used to mean a set of first category.

With the stage set, we now outline the argument that continuous functions that are differentiable at even one point of $[0, 1]$ form a meager subset of the metric space $C[0, 1]$.

Theorem 8.2.12. *The set*

$$D = \{f \in C[0, 1] : f'(x) \text{ exists for some } x \in [0, 1]\}$$

is a set of first category in $C[0, 1]$.

Proof. For each pair of natural numbers m, n , define

$$A_{m,n} = \left\{ f \in C[0, 1] : \text{there exists } x \in [0, 1] \text{ where} \right. \\ \left. \left| \frac{f(x) - f(t)}{x - t} \right| \leq n \text{ whenever } 0 < |x - t| < \frac{1}{m} \right\}.$$

This definition takes some time to digest. Think of $1/m$ as defining a δ -neighborhood around the point x , and view n as an upper bound on the magnitude of the slopes of lines through the two points $(x, f(x))$ and $(t, f(t))$. The set $A_{m,n}$ contains any function in $C[0, 1]$ for which it is possible to find at least one point x where the slopes through $(x, f(x))$ and points on the function nearby—within $1/m$ to be precise—are bounded by n .

Exercise 8.2.16. Show that if $f \in C[0, 1]$ is differentiable at a point $x \in [0, 1]$, then $f \in A_{m,n}$ for some pair $m, n \in \mathbf{N}$.

The collection of subsets $\{A_{m,n} : m, n \in \mathbf{N}\}$ is countable, and we have just seen that the union of these sets contains our set D . Because it is not difficult to see that a subset of a set of first category is first category, the final hurdle in the argument is to prove that each $A_{m,n}$ is nowhere-dense in $C[0, 1]$.

Fix m and n . The first order of business is to prove that $A_{m,n}$ is a closed set. To this end, let (f_k) be a sequence in $A_{m,n}$ and assume $f_k \rightarrow f$ in $C[0, 1]$. We need to show $f \in A_{m,n}$.

Because $f_k \in A_{m,n}$, then for each $k \in \mathbf{N}$ there exists a point $x_k \in [0, 1]$ where

$$\left| \frac{f_k(x_k) - f_k(t)}{x_k - t} \right| \leq n \quad \text{for all } 0 < |x_k - t| < 1/m.$$

Exercise 8.2.17. (a) The sequence (x_k) does not necessarily converge, but explain why there exists a subsequence (x_{k_l}) that is convergent. Let $x = \lim(x_{k_l})$.

(b) Prove that $f_{k_l}(x_{k_l}) \rightarrow f(x)$.

(c) Now finish the proof that $A_{m,n}$ is closed.

Because $A_{m,n}$ is closed, $\overline{A_{m,n}} = A_{m,n}$. In order to prove that $A_{m,n}$ is nowhere-dense, we just have to show that it contains no ϵ -neighborhoods, so pick an arbitrary $f \in A_{m,n}$, let $\epsilon > 0$, and consider the ϵ -neighborhood $V_\epsilon(f)$ in $C[0, 1]$. To show that this set is not contained in $A_{m,n}$, we must produce a function $g \in C[0, 1]$ that satisfies $\|f - g\|_\infty < \epsilon$ and has the property that there is no point $x \in [0, 1]$ where

$$\left| \frac{g(x) - g(t)}{x - t} \right| \leq n \quad \text{for all } 0 < |x - t| < 1/m.$$

Exercise 8.2.18. A continuous function is called *polygonal* if its graph consists of a finite number of line segments.

(a) Show that there exists a polygonal function $p \in C[0, 1]$ satisfying $\|f - p\|_\infty < \epsilon/2$.

(b) Show that if h is any function in $C[0, 1]$ that is bounded by 1, then the function

$$g(x) = p(x) + \frac{\epsilon}{2}h(x)$$

satisfies $g \in V_\epsilon(f)$.

(c) Construct a polygonal function $h(x)$ in $C[0, 1]$ that is bounded by 1 and leads to the conclusion $g \notin A_{m,n}$, where g is defined as in (b). Explain how this completes the argument for Theorem 8.2.12. \square

8.3 Euler's Sum

In Section 6.1 we saw Euler's first and most famous derivation of the formula

$$1 + \frac{1}{4} + \frac{1}{9} + \frac{1}{16} + \frac{1}{25} + \cdots = \frac{\pi^2}{6}.$$

At the crux of this argument are two representations for the function $\sin(x)$. The first is the standard Taylor series representation

$$(1) \quad \sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \cdots,$$

and the second is an infinite product representation

$$(2) \quad \sin(x) = x \left(1 - \frac{x}{\pi}\right) \left(1 + \frac{x}{\pi}\right) \left(1 - \frac{x}{2\pi}\right) \left(1 + \frac{x}{2\pi}\right) \cdots.$$

Although we have since made rigorous sense of the first equation (Example 6.6.4), proving the validity of equation (2) is still beyond our means.

The news is not all bad, however. In the time since Euler first made this discovery, dozens of different proofs for this result have been published, starting with several by Euler himself and continuing right up to the present. The machinery required in these arguments runs the gamut from multi-variable calculus to Fourier series to complex integration, but one in particular due to Boon Rim Choe relies mainly on Taylor series expansions and properties of uniformly convergent series. Choe's argument was published in 1987 but actually has much in common with one of Euler's original attempts. The proof outlined in this section follows Choe's argument with some simplifications due to Peter Duren.¹

Wallis's Product

Even though we don't currently have the tools to prove the infinite product formula for $\sin(x)$ in equation (2), we can prove a special case.

Exercise 8.3.1. Supply the details to show that when $x = \pi/2$ the product formula in (2) is equivalent to

$$(3) \quad \frac{\pi}{2} = \lim_{n \rightarrow \infty} \left(\frac{2 \cdot 2}{1 \cdot 3}\right) \left(\frac{4 \cdot 4}{3 \cdot 5}\right) \left(\frac{6 \cdot 6}{5 \cdot 7}\right) \cdots \left(\frac{2n \cdot 2n}{(2n-1)(2n+1)}\right),$$

where the infinite product in (2) is interpreted to be a limit of partial products. (Although it is not necessary for what follows, it might be useful to review the treatment of infinite products in Exercises 2.4.10 and 2.7.10.)

The goal of the next few exercises is to supply a proper proof for equation (3). This curious formula involving π was first discovered by John Wallis (1616–1703) and will provide some key ingredients for our proof of Euler's sum. It resurfaces again in Section 8.4 where the factorial function is defined.

Set

$$b_n = \int_0^{\frac{\pi}{2}} \sin^n(x) dx, \quad \text{for } n = 0, 1, 2, \dots$$

The first few terms are easy enough to calculate; in particular,

$$b_0 = \int_0^{\frac{\pi}{2}} 1 dx = \frac{\pi}{2} \quad \text{and} \quad b_1 = \int_0^{\frac{\pi}{2}} \sin(x) dx = 1.$$

¹[13], p. 92–95

Exercise 8.3.2. Assume $h(x)$ and $k(x)$ have continuous derivatives on $[a, b]$ and derive the integration-by-parts formula

$$\int_a^b h(t)k'(t)dt = h(b)k(b) - h(a)k(a) - \int_a^b h'(t)k(t)dt.$$

Exercise 8.3.3. (a) Using the simple identity $\sin^n(x) = \sin^{n-1}(x)\sin(x)$ and the previous exercise, derive the recurrence relation

$$b_n = \frac{n-1}{n}b_{n-2} \quad \text{for all } n \geq 2.$$

(b) Use this relation to generate the first three even terms and the first three odd terms of the sequence (b_n) .

(c) Write a general expression for b_{2n} and b_{2n+1} .

Because $0 \leq \sin^{n+1}(x) \leq \sin^n(x)$ on $[0, \pi/2]$, it follows that $b_{n+1} \leq b_n$ and (b_n) is decreasing. It turns out that $(b_n) \rightarrow 0$ but that isn't the limit we are interested in at the moment.

Exercise 8.3.4. Show

$$\lim_{n \rightarrow \infty} \frac{b_{2n}}{b_{2n+1}} = 1,$$

and use this fact to finish the proof of Wallis's product formula in (3).

There are some standard techniques for working with the notation of equation (3). For instance,

$$2 \cdot 4 \cdot 6 \cdots (2n) = 2^n n!$$

and

$$1 \cdot 3 \cdot 5 \cdots (2n+1) = \frac{(2n+1)!}{2 \cdot 4 \cdot 6 \cdots (2n)} = \frac{(2n+1)!}{2^n n!}.$$

Exercise 8.3.5. Derive the following alternative form of Wallis's product formula:

$$\sqrt{\pi} = \lim_{n \rightarrow \infty} \frac{2^{2n}(n!)^2}{(2n)!\sqrt{n}}.$$

Taylor Series

The next step in the argument is to generate the Taylor series for $\arcsin(x)$. This is not really possible to do directly from Taylor's formula for the coefficients, but keeping in mind that

$$(\arcsin(x))' = \frac{1}{\sqrt{1-x^2}},$$

we can get where we want to go by first finding the expansion for $1/\sqrt{1-x}$.

Exercise 8.3.6. Show that $1/\sqrt{1-x}$ has Taylor expansion $\sum_{n=0}^{\infty} c_n x^n$, where $c_0 = 1$ and

$$c_n = \frac{(2n)!}{2^{2n}(n!)^2} = \frac{1 \cdot 3 \cdot 5 \cdots (2n-1)}{2 \cdot 4 \cdot 6 \cdots 2n}$$

for $n \geq 1$.

The coefficients c_n should look familiar from our work on Wallis's product. Exercise 8.3.5 can be rephrased as

$$\sqrt{\pi} = \lim_{n \rightarrow \infty} \frac{1}{c_n \sqrt{n}}.$$

Exercise 8.3.7. Show that $\lim c_n = 0$ but $\sum_{n=0}^{\infty} c_n$ diverges.

The divergence of $\sum_{n=0}^{\infty} c_n$ makes sense when we consider the Taylor series for $1/\sqrt{1-x}$. We want to determine the values of x for which

$$(4) \quad \frac{1}{\sqrt{1-x}} = \sum_{n=0}^{\infty} c_n x^n,$$

and $x = 1$ is not in the domain of the left side. We do aim to prove (4) for all $x \in (-1, 1)$ but the usual word of warning is in order. Having computed the coefficients c_n , it is not enough to simply argue that the series on the right side converges when $|x| < 1$. To properly establish (4) we are going to show that the error function

$$E_N(x) = \frac{1}{\sqrt{1-x}} - \sum_{n=0}^N c_n x^n$$

tends to zero as $N \rightarrow \infty$. Back in Section 6.6, the primary tool we used for this task was Lagrange's Remainder Theorem (Theorem 6.6.3), but it is not up to this particular challenge

Exercise 8.3.8. Using the expression for $E_N(x)$ from Lagrange's Remainder Theorem, show that equation (4) is valid for all $|x| < 1/2$. What goes wrong when we try to use this method to prove (4) for $x \in (1/2, 1)$?

The Integral Form of the Remainder

The moral of the previous exercise is that we need a different method for estimating $E_N(x)$. The Lagrange form of the remainder grows out of the Mean Value Theorems and yields a formula for the error function in terms of the derivative $f^{(N+1)}$. Now that we are in possession of a proper definition of the integral, we can derive another useful formula for $E_N(x)$.

Theorem 8.3.1 (Integral Remainder Theorem). *Let f be differentiable $N + 1$ times on $(-R, R)$ and assume $f^{(N+1)}$ is continuous. Define $a_n = f^{(n)}(0)/n!$ for $n = 0, 1, \dots, N$, and let*

$$S_N(x) = a_0 + a_1 x + a_2 x^2 + \cdots + a_N x^N.$$

For all $x \in (-R, R)$, the error function $E_N(x) = f(x) - S_N(x)$ satisfies

$$E_N(x) = \frac{1}{N!} \int_0^x f^{(N+1)}(t)(x-t)^N dt.$$

Proof. The case $x = 0$ is easy to check, so let's take $x \neq 0$ in $(-R, R)$ and keep in mind that x is a fixed constant in what follows. To avoid a few technical distractions, let's just consider the case $x > 0$.

Exercise 8.3.9. (a) Show

$$f(x) = f(0) + \int_0^x f'(t) dt.$$

(b) Now use a previous result from this section to show

$$f(x) = f(0) + f'(0)x + \int_0^x f''(t)(x-t) dt.$$

(c) Continue in this fashion to complete the proof of the theorem. □

To gain a better understanding of this formulation for $E_N(x)$ and simultaneously make some headway on our exploration of equation (4), let's return to the special case $f(x) = 1/\sqrt{1-x}$.

Exercise 8.3.10. (a) Make a rough sketch of $1/\sqrt{1-x}$ and $S_2(x)$ over the interval $(-1, 1)$, and compute $E_2(x)$ for $x = 1/2, 3/4$, and $8/9$.

(b) For a general x satisfying $-1 < x < 1$, show

$$E_2(x) = \frac{15}{16} \int_0^x \left(\frac{x-t}{1-t} \right)^2 \frac{1}{(1-t)^{3/2}} dt.$$

(c) Explain why the inequality

$$\left| \frac{x-t}{1-t} \right| \leq |x|$$

is valid, and use this to find an overestimate for $|E_2(x)|$ that no longer involves an integral. Note that this estimate will necessarily depend on x . Confirm that things are going well by checking that this overestimate is in fact larger than $|E_2(x)|$ at the three computed values from part (a).

(d) Finally, show $E_N(x) \rightarrow 0$ as $N \rightarrow \infty$ for an arbitrary $x \in (-1, 1)$.

Having established that the Taylor series in (4) does indeed converge for all $|x| < 1$, it is now clear sailing to produce a Taylor series representation for $\arcsin(x)$. The first step is to substitute x^2 for x in (4) to get

$$\frac{1}{\sqrt{1-x^2}} = \sum_{n=0}^{\infty} c_n x^{2n} \quad \text{for all } |x| < 1.$$

The next step is to take the term-by-term anti-derivative of this series. Any time we start manipulating infinite series as though they were finite in nature we need to pause and make sure we are on solid footing.

Exercise 8.3.11. Assuming that the derivative of $\arcsin(x)$ is indeed $1/\sqrt{1-x^2}$, supply the justification that allows us to conclude

$$(5) \quad \arcsin(x) = \sum_{n=0}^{\infty} \frac{c_n}{2n+1} x^{2n+1} \quad \text{for all } |x| < 1.$$

Exercise 8.3.12. Our work thus far shows that the Taylor series in (5) is valid for all $|x| < 1$, but note that $\arcsin(x)$ is continuous for all $|x| \leq 1$. Carefully explain why the series in (5) converges uniformly to $\arcsin(x)$ on the closed interval $[-1, 1]$.

Summing $\sum_{n=1}^{\infty} 1/n^2$

Every proof of Euler's sum contains a moment of genuine ingenuity at some point, and this is where our proof takes an unanticipated turn.

Let's make the substitution $x = \sin(\theta)$ in (5) where we restrict our attention to $-\pi/2 \leq \theta \leq \pi/2$. The result is

$$\theta = \arcsin(\sin(\theta)) = \sum_{n=0}^{\infty} \frac{c_n}{2n+1} \sin^{2n+1}(\theta)$$

which converges uniformly on $[-\pi/2, \pi/2]$.

Exercise 8.3.13. (a) Show

$$\int_0^{\pi/2} \theta d\theta = \sum_{n=0}^{\infty} \frac{c_n}{2n+1} b_{2n+1},$$

being careful to justify each step in the argument. The term b_{2n+1} refers back to our earlier work on Wallis's product.

(b) Deduce

$$\frac{\pi^2}{8} = \sum_{n=0}^{\infty} \frac{1}{(2n+1)^2},$$

and use this to finish the proof that $\pi^2/6 = \sum_{n=1}^{\infty} 1/n^2$.

The Riemann-Zeta Function

Euler's determination of the value of $\sum 1/n^2$ brought him international recognition and represented a significant milestone in what would be a lifelong exploration of series of the form $\sum 1/n^s$. Euler's original argument for summing $\sum 1/n^2$ discussed in Section 6.1 involved equating the coefficient of x^2 in two different series expansions for $\sin(x)/x$. By equating the coefficients of higher powers of x he was also able to sum $\sum 1/n^s$ for $s = 4, 6, 8, 10$ and 12 . (Try it for $s = 4$.) Eventually, Euler worked out a general formula for any even natural number, and in the process he shifted his focus to thinking about $\sum 1/n^s$ as a *function* of the variable s . The iconic notation

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s} \quad \text{for all } s > 1,$$

and the name—the Riemann-zeta function—would come one hundred years later, but it was Euler who first unearthed many deep properties of this function. Significant among these is a connection to the prime numbers, evident in the Eulerian formula

$$(6) \quad \sum_{n=1}^{\infty} \frac{1}{n^s} = \left(\frac{1}{1-2^{-s}} \right) \left(\frac{1}{1-3^{-s}} \right) \left(\frac{1}{1-5^{-s}} \right) \left(\frac{1}{1-7^{-s}} \right) \cdots,$$

where the product is taken over all the primes. The mathematics underlying the Riemann-zeta function gets complicated very quickly, but this particular formula is actually quite accessible. Notice that for each prime p ,

$$\frac{1}{1-p^{-s}} = 1 + \frac{1}{p^s} + \frac{1}{p^{2s}} + \frac{1}{p^{3s}} + \frac{1}{p^{4s}} + \cdots.$$

Multiplying out the product on the right in (6) in this fashion and using the fact that every $n \in \mathbf{N}$ is a unique product of primes leads naturally to the given relationship.

Euler returned to study $\zeta(s)$ many times in his career, in part it seems to tend to the unfinished business of evaluating $\sum 1/n^s$ for the odd integers. Amid his many successes, this was a challenge that eluded Euler, as it has eluded every mathematician since.

8.4 Inventing the Factorial Function

The goal of this section is to produce a function $f(x)$, defined on all of \mathbf{R} , with the property that $f(n) = n!$ for each $n \in \mathbf{N}$. With no other restriction on f , this is as easy as it is uninteresting—simply define f piecewise in such a way that it passes through the points $(1,1)$, $(2,2)$, $(3,6)$, $(4,24)$, and so on. Letting

$$f(x) = \begin{cases} n! & \text{if } n \leq x < n+1, n \in \mathbf{N} \\ 1 & \text{if } x < 1 \end{cases}$$

does the trick.

To make this problem meaningful we need to be much more discriminating about what properties we require f to have. Should f be continuous? Differentiable? Twice differentiable? We shall see about this. This problem actually has its origins in a series of 1729 letters between Christian Goldbach (of “Goldbach’s Conjecture” fame, although that is a different story) and Leonard Euler. The term “function” in Euler’s day implicitly referred to a mapping defined by an analytic expression comprised of the elementary functions and operations of calculus. Logarithms, exponentials, polynomials, and power series were examples of 18th century functions; the piecewise concoction proposed above was not.

Thus, a better statement of our goal—although still a little imprecise—is to find a function defined by a single, organic formula which extends the definition of $n!$ in a meaningful way to non-natural numbers.

Exercise 8.4.1. For $n \in \mathbf{N}$, let

$$n\# = n + (n - 1) + (n - 2) + \cdots + 2 + 1.$$

- (a) Without looking ahead, decide if there is a natural way to define $0\#$. How about $(-2)\#$? Conjecture a reasonable value for $\frac{7}{2}\#$.
- (b) Now prove $n\# = \frac{1}{2}n(n + 1)$ for all $n \in \mathbf{N}$, and revisit part (a).

The formula in part (b) of the previous exercise not only simplifies the calculation of $n\#$ for large values of n , but also yields a properly defined function on \mathbf{R} when the discrete variable n is replaced with the continuous variable x . Indeed, Euler would be perfectly comfortable with the expression $x\# = \frac{1}{2}x(x + 1)$.

We are seeking something similar for $n!$. What is the right definition for $x!$ when $x \in \mathbf{R}$?

The Exponential Function

The idea of extending the definition of a function defined on \mathbf{N} to all of \mathbf{R} may at first sound like a somewhat whimsical enterprise, but it is perfectly analogous to the way we come to understand a function like 2^x . Similar to $n!$, 2^n for $n \in \mathbf{N}$ is unambiguous and meaningful the minute we understand multiplication, but something like $2^{-\pi}$ is another matter. Because it is instructive, and because we are going to presently need functions of the form t^x , let’s take a moment to define exponential functions in a rigorous way.

Typically the way a function like 2^x gets defined on \mathbf{R} is through a series of domain expansions. Starting with 2^n , we first expand the domain to \mathbf{Z} using reciprocals, then to \mathbf{Q} using roots, and finally to \mathbf{R} using continuity. Although we could follow this strategy, we are going to take a different approach that has the advantage of yielding the important properties we need more efficiently.

Step one is to properly define the natural exponential function e^x . Back in Chapter 6, we assumed e^x was already defined and showed how it could be

represented by its Taylor series. Here we flip this process around. The problem on the table is to rigorously construct a proper definition for e^x , and the theory of power series gives us a bedrock foundation on which to build.

Define

$$(1) \quad E(x) = \sum_{n=0}^{\infty} \frac{x^n}{n!} = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$$

Exercise 8.4.2. Verify that the series converges absolutely for all $x \in \mathbf{R}$, that $E(x)$ is differentiable on \mathbf{R} , and $E'(x) = E(x)$.

Exercise 8.4.3. (a) Use the results of Exercise 2.8.7 and the binomial formula to show that $E(x+y) = E(x)E(y)$ for all $x, y \in \mathbf{R}$.

(b) Show that $E(0) = 1$, $E(-x) = 1/E(x)$, and $E(x) > 0$ for all $x \in \mathbf{R}$.

The takeaway here is that the power series $E(x)$ satisfies all the properties we associate with the exponential function, and we can therefore give ourselves permission to go back to the more familiar notation e^x in place of $E(x)$. What happens if we have a momentary relapse and interpret e^x as the real number $e \approx 2.71828\dots$ raised to the power x rather than $E(x)$? Not to worry—the two interpretations coincide, whenever the former is defined in the usual way.

Exercise 8.4.4. Define $e = E(1)$. Show $E(n) = e^n$ and $E(m/n) = (\sqrt[n]{e})^m$ for all $m, n \in \mathbf{Z}$.

One final property of e^x we need is its behavior as $x \rightarrow \pm\infty$.

Definition 8.4.1. Given $f : [a, \infty] \rightarrow \mathbf{R}$, we say that $\lim_{x \rightarrow \infty} f(x) = L$ if, for all $\epsilon > 0$, there exists $M > a$ such that whenever $x \geq M$ it follows that $|f(x) - L| < \epsilon$.

Exercise 8.4.5. Show $\lim_{x \rightarrow \infty} x^n e^{-x} = 0$ for all $n = 0, 1, 2, \dots$

To get started notice that when $x \geq 0$, all the terms in (1) are positive.

Other Bases

Having set e^x on solid mathematical footing, we can now do the same for t^x where $t > 0$. This requires use of the natural logarithm.

Exercise 8.4.6. (a) Explain why we know e^x has an inverse function—let's call it $\log x$ —defined on the strictly positive real numbers and satisfying

- (i) $\log(e^y) = y$ for all $y \in \mathbf{R}$ and
- (ii) $e^{\log x} = x$, for all $x > 0$.

(b) Prove $(\log x)' = 1/x$. (See Exercise 5.2.12.)

(c) Fix $y > 0$ and differentiate $\log(xy)$ with respect to x . Conclude that

$$\log(xy) = \log x + \log y \quad \text{for all } x, y > 0.$$

- (d) For $t > 0$ and $n \in \mathbf{N}$, t^n has the usual interpretation as $t \cdot t \cdots t$ (n times). Show that

$$(2) \quad t^n = e^{n \log t} \quad \text{for all } n \in \mathbf{N}.$$

Part (d) of the previous exercise is the pivotal formula because the expression on the right of the equal sign is meaningful if we replace n with $x \in \mathbf{R}$. This is our cue to use the identity in (2) as a template for the definition of t^x on all of \mathbf{R} .

Definition 8.4.2. Given $t > 0$, define the exponential function t^x to be

$$t^x = e^{x \log t} \quad \text{for all } x \in \mathbf{R}.$$

Exercise 8.4.7. (a) Show $t^{m/n} = (\sqrt[n]{t})^m$ for all $m, n \in \mathbf{N}$.

(b) Show $\log(t^x) = x \log t$, for all $t > 0$ and $x \in \mathbf{R}$.

(c) Show t^x is differentiable on \mathbf{R} and find the derivative.

Finding the right definition for $x!$ is harder than defining t^x , but the strategy is essentially the same. We are seeking a formula of the form $n! = g(n)$ where g yields a meaningful formula when n is replaced by x . What might such a function $g(x) = x!$ look like when graphed over \mathbf{R} ? For $x \geq 0$ it must grow extremely rapidly to keep up with $n!$, but how about on $x < 0$? Using a functional equation for $x!$ we can create a reasonable artist's rendering of the function we are looking for.

The Functional Equation

A defining property of the factorial on \mathbf{N} is that $1! = 1$ and $n! = n(n-1)!$ for all $n \geq 2$. Thus it seems reasonable to require the same from our currently mythic function $x!$ defined on \mathbf{R} . Whatever $x!$ means it should satisfy

$$x! = x(x-1)! \quad \text{for all } x \in \mathbf{R}.$$

Setting $n = 1$ in this equation, for example, yields $1 = 0!$.

Exercise 8.4.8. Inspired by the fact that $0! = 1$ and $1! = 1$, let $h(x)$ satisfy

(i) $h(x) = 1$ for all $0 \leq x \leq 1$, and

(ii) $h(x) = xh(x-1)$ for all $x \in \mathbf{R}$.

(a) Find a formula for $h(x)$ on $[1, 2]$, $[2, 3]$, and $[n, n+1]$ for arbitrary $n \in \mathbf{N}$.

(b) Now do the same for $[-1, 0]$, $[-2, -1]$, and $[-n, -n+1]$.

(c) Sketch h over the domain $[-4, 4]$.

Notice that $h(x)$ satisfies $h(n) = n!$ and it is at least continuous for $x \geq 0$, but its piecewise definition and its many non-differentiable corners disqualify it from being our sought after factorial function. One legitimate conclusion that arises out of this exercise is that $x!$, when we find it, will exhibit the same asymptotic behavior as h at $x = -1, -2, -3, \dots$, and thus won't be defined on the negative integers.

Improper Riemann Integrals

For reasons that will become clear, we need to make rigorous sense of an expression like

$$\int_0^{\infty} e^{-t} dt.$$

Most likely familiar from calculus, integrals over unbounded regions like $[0, \infty)$ are called *improper Riemann integrals* and are defined by taking the limit of “proper” integrals.

Definition 8.4.3. Assume f is defined on $[a, \infty)$ and integrable on every interval of the form $[a, b]$. Then define $\int_a^{\infty} f$ to be

$$\lim_{b \rightarrow \infty} \int_a^b f,$$

provided the limit exists. In this case we say the improper integral $\int_a^{\infty} f$ *converges*.

Exercise 8.4.9. (a) Show that the improper integral $\int_a^{\infty} f$ converges if and only if, for all $\epsilon > 0$ there exists $M > a$ such that whenever $d > c \geq M$ it follows that

$$\left| \int_c^d f \right| < \epsilon.$$

(In one direction it will be useful to consider the sequence $a_n = \int_a^{a+n} f$.)

(b) Show that if $0 \leq f \leq g$ and $\int_a^{\infty} g$ converges then $\int_a^{\infty} f$ converges.

(c) Part (a) is a Cauchy criterion, and part (b) is a comparison test. State and prove an absolute convergence test for improper integrals.

Exercise 8.4.10. (a) Use the properties of e^t previously discussed to show

$$\int_0^{\infty} e^{-t} dt = 1.$$

(b) Show

$$(3) \quad \frac{1}{\alpha} = \int_0^{\infty} e^{-\alpha t} dt, \quad \text{for all } \alpha > 0.$$

Just for a moment, let's take our analysis gloves off and ask what we think might happen if we differentiate formula (3) with respect to α .

On the left-hand side we certainly get

$$\left[\frac{1}{\alpha}\right]' = -\frac{1}{\alpha^2}.$$

On the right-hand side of (3), let's brazenly crash through the integral sign and take the derivative of the integrand $e^{-\alpha t}$ with respect to α (thinking of t as a constant.) The result is

$$[e^{-\alpha t}]' = e^{-\alpha t} \cdot (-t).$$

The question, then, is whether this is a valid manipulation. Is it true that

$$(4) \quad \frac{1}{\alpha^2} = \int_0^{\infty} te^{-\alpha t} dt?$$

Well, let's compute the integral and find out.

Exercise 8.4.11. (a) Evaluate $\int_0^b te^{-\alpha t} dt$ using the integration-by-parts formula from Exercise 7.5.6. The result will be an expression in α and b .

(b) Now compute $\int_0^{\infty} te^{-\alpha t} dt$ and verify equation (4).

Apparently, our bold differentiation of equation (3) into equation (4) worked out. Now it's time to put our analysis gloves back on and see why this is so.

Differentiating Under the Integral

Let $f(x, t)$ be a function of two variables, defined for all $a \leq x \leq b$ and $c \leq t \leq d$. The domain of f is then a *rectangle* D in \mathbf{R}^2 .

What does it mean to say f is continuous at a point (x_0, t_0) in D ? Section 8.2 on metric spaces gives a more thorough explanation, but the only real difference from the single variable setting is that we have to replace our sense of *distance* between points (x_0, t_0) and (x, t) with the familiar Euclidean distance formula

$$\|(x, t) - (x_0, t_0)\| = \sqrt{(x - x_0)^2 + (t - t_0)^2}.$$

Definition 8.4.4. A function $f : D \rightarrow \mathbf{R}$ is continuous at (x_0, t_0) if for all $\epsilon > 0$, there exists $\delta > 0$ such that whenever $\|(x, t) - (x_0, t_0)\| < \delta$, it follows that

$$|f(x, t) - f(x_0, t_0)| < \epsilon.$$

Exercise 8.4.12. Assume the function $f(x, t)$ is continuous on the rectangle $D = \{(x, t) : a \leq x \leq b, c \leq t \leq d\}$. Explain why the function

$$F(x) = \int_c^d f(x, t) dt$$

is properly defined for all $x \in [a, b]$.

It should not be too surprising that Theorem 4.4.7 has an analogue in the \mathbf{R}^2 setting. The set D is compact in \mathbf{R}^2 , and a continuous function on D is uniformly continuous in the sense that the δ in Definition 8.4.4 can be chosen independently of the point (x_0, t_0) .

Theorem 8.4.5. *If $f(x, t)$ is continuous on D , then $F(x) = \int_c^d f(x, t) dt$ is uniformly continuous on $[a, b]$.*

Exercise 8.4.13. Prove Theorem 8.4.5.

Taking inspiration from equations (3) and (4), let's add the assumption that for each fixed value of t in $[c, d]$, the function $f(x, t)$ is a differentiable function of x ; that is,

$$f_x(x, t) = \lim_{z \rightarrow x} \frac{f(z, t) - f(x, t)}{z - x}$$

exists for all $(x, t) \in D$. In addition, let's assume that the derivative function $f_x(x, t)$ is continuous.

Theorem 8.4.6. *If $f(x, t)$ and $f_x(x, t)$ are continuous on D , then the function $F(x) = \int_c^d f(x, t) dt$ is differentiable and*

$$F'(x) = \int_c^d f_x(x, t) dt.$$

Proof. Fix x in $[a, b]$ and let $\epsilon > 0$ be arbitrary. Our task is to find a $\delta > 0$ such that

$$(5) \quad \left| \frac{F(z) - F(x)}{z - x} - \int_c^d f_x(x, t) dt \right| < \epsilon$$

whenever $0 < |z - x| < \delta$.

Exercise 8.4.14. Finish the proof of Theorem 8.4.6 □

Improper Integrals, Revisited

Theorem 8.4.6 is a formal justification for differentiating under the integral sign, but we need to extend this result to the case where the integral is improper. Looking back one more time to our motivating example in equation (3), we see that what we have is a function $f(x, t)$ where the domain of the variable t is the unbounded interval $c \leq t < \infty$.

Let's fix x from some set $A \subseteq \mathbf{R}$. For such an x , we define

$$(6) \quad F(x) = \int_c^\infty f(x, t) dt = \lim_{d \rightarrow \infty} \int_c^d f(x, t) dt,$$

provided the limit exists.

Notice that the formula in (6) is a *pointwise* statement. Given an $x \in A$ and $\epsilon > 0$, we can find an M (perhaps dependent on x) where

$$\left| F(x) - \int_c^d f(x, t) dt \right| < \epsilon$$

whenever $d \geq M$. As we have seen on numerous occasions, the elixir required to ensure that good behavior in the finite setting extends to the infinite setting is uniformity.

Definition 8.4.7. Given $f(x, t)$ defined on $D = \{(x, t) : x \in A, c \leq t\}$, assume $F(x) = \int_c^\infty f(x, t) dt$ exists for all $x \in A$. We say the improper integral *converges uniformly* to $F(x)$ on A if for all $\epsilon > 0$, there exists $M > c$ such that

$$\left| F(x) - \int_c^d f(x, t) dt \right| < \epsilon$$

for all $d \geq M$ and all $x \in A$.

Exercise 8.4.15. (a) Show that the improper integral $\int_0^\infty e^{-xt} dt$ converges uniformly to $1/x$ on the set $[1/2, \infty)$.

(b) Is the convergence uniform on $(0, \infty)$?

Exercise 8.4.16. Prove the following analogue of the Weierstrass M-Test for improper integrals: If $f(x, t)$ satisfies $|f(x, t)| \leq g(t)$ for all $x \in A$ and $\int_a^\infty g(t) dt$ converges, then $\int_a^\infty f(x, t) dt$ converges uniformly on A .

An immediate consequence of Definition 8.4.7 is that if the improper integral converges uniformly then the sequence of functions defined by

$$F_n(x) = \int_c^{c+n} f(x, t) dt$$

converges uniformly to $F(x)$ on $[a, b]$. This observation gives us access to the host of useful results we developed in Chapter 6.

Theorem 8.4.8. If $f(x, t)$ is continuous on $D = \{(x, t) : a \leq x \leq b, c \leq t\}$, then

$$F(x) = \int_c^\infty f(x, t) dt$$

is uniformly continuous on $[a, b]$, provided the integral converges uniformly.

Exercise 8.4.17. Prove Theorem 8.4.8.

Theorem 8.4.9. Assume the function $f(x, t)$ is continuous on $D = \{(x, t) : a \leq x \leq b, c \leq t\}$ and $F(x) = \int_c^\infty f(x, t) dt$ exists for each $x \in [a, b]$. If the derivative function $f_x(x, t)$ exists and is continuous, then

$$(7) \quad F'(x) = \int_c^\infty f_x(x, t) dt,$$

provided the integral in (7) converges uniformly.

Exercise 8.4.18. Prove Theorem 8.4.9.

The Factorial Function

It's time to return our attention to equation (3) from earlier in this section:

$$\frac{1}{\alpha} = \int_0^{\infty} e^{-\alpha t} dt, \quad \text{for all } \alpha > 0.$$

Exercise 8.4.19. (a) Although we verified it directly, show how to use the theorems in this section to give a second justification for the formula

$$\frac{1}{\alpha^2} = \int_0^{\infty} t e^{-\alpha t} dt, \quad \text{for all } \alpha > 0.$$

(b) Now derive the formula

$$(8) \quad \frac{n!}{\alpha^{n+1}} = \int_0^{\infty} t^n e^{-\alpha t} dt, \quad \text{for all } \alpha > 0.$$

If we set $\alpha = 1$ in equation (8) we get

$$n! = \int_0^{\infty} t^n e^{-t} dt.$$

The appearance of $n!$ on the left side of this equation is an exciting development, especially because where n appears on the right it can be meaningfully replaced by a real variable x , at least when $x \geq 0$. This is the equation we have been looking for!

Definition 8.4.10. For $x \geq 0$, define the factorial function

$$x! = \int_0^{\infty} t^x e^{-t} dt.$$

Exercise 8.4.20. (a) Show that $x!$ is an infinitely differentiable function on $(0, \infty)$ and produce a formula for the n^{th} derivative. In particular show that $(x!)'' > 0$.

(b) Use the integration-by-parts formula employed earlier to show that $x!$ satisfies the functional equation

$$(x+1)! = (x+1)x!.$$

The previous exercise is our first piece of evidence that we have found the right definition for $x!$. There is more to come.

A consequence of $(x!)'' > 0$ is that $x!$ is a *convex* function. In calculus this is usually referred to as “concave up” and means that the line segment connecting two points on the graph of $x!$ always sits above the curve. Said another way, there are no inflection points in $x!$ and the slope of the curve steadily increases as the graph passes through the points $(n, n!)$ for $n = 0, 1, 2, \dots$. We did not

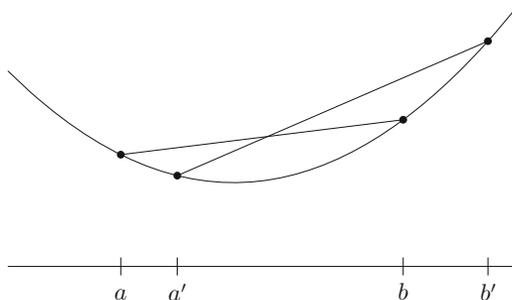


Figure 8.1: INCREASING CHORD SLOPES ON A CONVEX FUNCTION.

mention this property at the time, but reflecting on our earlier analogy between 2^x and $x!$, convexity is a natural condition to desire in our factorial function.

In fact, not only is $x!$ convex but $\log(x!)$ is also convex. This is a stronger statement. (Consider, for instance, the graphs of $x^2 + 1$ and $\log(x^2 + 1)$.) The proof is a little technical and we won't go through it, but the fact that $\log(x!)$ is convex on $x \geq 0$ is quite significant. Here's why.

Theorem 8.4.11 (Bohr–Mollerup Theorem). *There is a unique positive function f defined on $x \geq 0$ satisfying*

- (i) $f(0) = 1$
- (ii) $f(x + 1) = (x + 1)f(x)$, and
- (iii) $\log(f(x))$ is convex.

Because $x!$ satisfies properties (i), (ii), and (iii), it follows that $f(x) = x!$.

Proof. We need one more geometrically plausible fact about convex functions. If $[a, b]$ and $[a', b']$ are two intervals in the domain of a convex function ϕ , and $a \leq a'$ and $b \leq b'$, then the slopes of the chords over these intervals satisfy

$$\frac{\phi(b) - \phi(a)}{b - a} \leq \frac{\phi(b') - \phi(a')}{b' - a'}.$$

(See Figure 8.1).

Because f satisfies properties (i) and (ii) we know $f(n) = n!$ for all $n \in \mathbf{N}$. Now fix $n \in \mathbf{N}$ and $x \in (0, 1]$.

Exercise 8.4.21. (a) Use the convexity of $\log(f(x))$ and the three intervals $[n - 1, n]$, $[n, n + x]$, and $[n, n + 1]$ to show

$$x \log(n) \leq \log(f(n + x)) - \log(n!) \leq x \log(n + 1).$$

(b) Show $\log(f(n + x)) = \log(f(x)) + \log((x + 1)(x + 2) \cdots (x + n))$.

(c) Now establish that

$$0 \leq \log(f(x)) - \log\left(\frac{n^x n!}{(x+1)(x+2)\cdots(x+n)}\right) \leq x \log\left(1 + \frac{1}{n}\right).$$

(d) Conclude that

$$f(x) = \lim_{n \rightarrow \infty} \frac{n^x n!}{(x+1)(x+2)\cdots(x+n)}, \quad \text{for all } x \in (0, 1].$$

(e) Finally, show that the conclusion in (d) holds for all $x \geq 0$.

Because we have arrived at an explicit formula for $f(x)$, the function $f(x)$ must be unique. By virtue of the fact that $x!$ satisfies conditions (i), (ii), and (iii) of the theorem, we can conclude that $x!$ is this unique function; i.e., $f(x) = x!$. Thus, not only have we proved the theorem, but we have also discovered an alternate representation for the factorial function called the *Gauss product formula*:

$$(9) \quad x! = \int_0^\infty t^x e^{-t} dt = \lim_{n \rightarrow \infty} \frac{n^x n!}{(x+1)(x+2)\cdots(x+n)},$$

for all $x \geq 0$. □

What happens if $x < 0$? The integral in Definition 8.4.10 becomes improper for a second reason when $x < 0$ because t^x is unbounded and undefined at $t = 0$. If $-1 < x < 0$, it is not hard to show that the integral still converges. On the other hand, the functional equation in Exercise 8.4.20(b) provides a natural way to extend the definition of $x!$ to all of \mathbf{R} . Just as in Exercise 8.4.8, the resulting function is never zero, alternating between positive and negative components with vertical asymptotes at $x = -1, -2, -3, \dots$

The Gamma Function

The focus of our discussion has been on the ingredients that go into the definition of $x!$ —improper integrals, proper definitions of exponential functions, differentiating under the integral sign—but the end result is a function worthy of its own separate chapter. Since its discovery by Euler, the factorial function has become ubiquitous in numerous branches of analysis.

One of the early modifications that occurred was a shift in the domain of $x!$ and a change in the notation. Adrien Marie Legendre introduced the Greek letter Γ (gamma) and set

$$\Gamma(x) = (x-1)! = \int_0^\infty t^{x-1} e^{-t} dt,$$

so that $\Gamma(n+1) = n!$ and $x\Gamma(x) = \Gamma(x+1)$. This convention eventually became the standard, and so it is the gamma function that routinely appears in formulas from number theory, probability, geometry, and beyond.

Philip Davis's article on the history of the gamma function (see [11]) is an excellent place to get a sense of the important role the gamma function has played in the development of analysis.² Davis's essay seems to be at least part of the inspiration for a wonderful series of articles by David Fowler that explore the properties of $x!$ in an original and accessible way.³ Here is one of the anecdotes Fowler offers, which serves as an enticing clue for how intricately the gamma/factorial function is connected to the larger mathematical landscape.

Recall that when $x!$ is extended to all of \mathbf{R} via the functional equation $x! = x(x-1)!$ we get asymptotes at every negative integer. Thus, there is a compelling reason to consider the reciprocal function $1/x!$ which we can take to be zero for $x = -1, -2, -3, \dots$

Exercise 8.4.22. (a) Where does $g(x) = \frac{x}{x!(-x)!}$ equal zero? What other familiar function has the same set of roots?

(b) The function e^{-x^2} provides the raw material for the all-important Gaussian bell curve from probability, where it is known that $\int_{-\infty}^{\infty} e^{-x^2} dx = \sqrt{\pi}$. Use this fact (and some standard integration techniques) to evaluate $(1/2)!$.

(c) Now use (a) and (b) to conjecture a striking relationship between the factorial function and a well-known function from trigonometry.

Exercise 8.4.23. As a parting shot, use the value for $(1/2)!$ and the Gauss product formula in equation (9) to derive the famous product formula for π discovered by John Wallis in the 1650s:

$$\frac{\pi}{2} = \lim_{n \rightarrow \infty} \left(\frac{2 \cdot 2}{1 \cdot 3} \right) \left(\frac{4 \cdot 4}{3 \cdot 5} \right) \left(\frac{6 \cdot 6}{5 \cdot 7} \right) \cdots \left(\frac{2n \cdot 2n}{(2n-1)(2n+1)} \right).$$

8.5 Fourier Series

In his famous treatise, *Théorie Analytique de la Chaleur* (The Analytical Theory of Heat), 1822, Joseph Fourier (1768–1830) boldly asserts, “Thus there is no function $f(x)$, or part of a function, which cannot be expressed by a trigonometric series.”⁴

It is difficult to exaggerate the mathematical richness of this idea. It has been convincingly argued by mathematical historians that the ensuing investigation into the validity of Fourier's conjecture was the fundamental catalyst for the pursuit of rigor that characterizes 19th century mathematics. Power series had been in wide use in the 150 years leading up to Fourier's work, largely because they behaved so well under the operations of calculus. A function expressed as a power series is continuous, differentiable an infinite number of times, and

²Exercise 8.4.1, as well as the insight of comparing the development of $x!$ to 2^x , are borrowed from this piece.

³Exercise 8.4.8 is borrowed from Fowler's treatment in [15].

⁴Quoted passages in this section are taken from [9].

can be integrated and differentiated as though it were a polynomial. In the presence of such agreeable behavior, there was no compelling reason for mathematicians to formulate a more precise understanding of “limit” or “convergence” because there were no arguments to resolve. Fourier’s successful implementation of trigonometric series to the study of heat flow changed all of this. To understand what the fuss was really about, we need to look more closely at what Fourier was asserting, focusing individually on the terms “function,” “express,” and “trigonometric series.”

Trigonometric Series

The basic principle behind any series representations is to express a given function $f(x)$ as a sum of simpler functions. For power series, the component functions are $\{1, x, x^2, x^3, \dots\}$, so that the series takes the form

$$f(x) = \sum_{n=0}^{\infty} a_n x^n = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \dots$$

A *trigonometric series* is a very different type of infinite series where the functions

$$\{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \cos(3x), \sin(3x), \dots\}$$

serve as the components. Thus, a trigonometric series has the form

$$\begin{aligned} f(x) &= a_0 + a_1 \cos(x) + b_1 \sin(x) + a_2 \cos(2x) + b_2 \sin(2x) + a_3 \cos(3x) + \dots \\ &= a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx). \end{aligned}$$

The idea of representing a function in this way was not completely new when Fourier first publicly proposed it in 1807. About 50 years earlier, Jean Le Rond d’Alembert (1717–1783) published the partial differential equation

$$(1) \quad \frac{\partial^2 u}{\partial x^2} = \frac{\partial^2 u}{\partial t^2}$$

as a means of describing the motion of a vibrating string. In this model, the function $u(x, t)$ represents the displacement of the string at time $t \geq 0$ and at some point x , which we will take to be in the interval $[0, \pi]$. Because the string is understood to be attached at each end of this interval, we have

$$(2) \quad u(0, t) = 0 \quad \text{and} \quad u(\pi, t) = 0$$

for all values of $t \geq 0$. Now, at $t = 0$, the string is displaced some initial amount, and at the moment it is released we assume

$$(3) \quad \frac{\partial u}{\partial t}(x, 0) = 0,$$

meaning that, although the string immediately starts to move, it is given no initial velocity at any point. Finding a function $u(x, t)$ that satisfies equations (1), (2), and (3) is not too difficult.

Exercise 8.5.1. (a) Verify that

$$u(x, t) = b_n \sin(nx) \cos(nt)$$

satisfies equations (1), (2), and (3) for any choice of $n \in \mathbf{N}$ and $b_n \in \mathbf{R}$. What goes wrong if $n \notin \mathbf{N}$?

- (b) Explain why any finite sum of functions of the form given in part (a) would also satisfy (1), (2), and (3). (Incidentally, it is possible to hear the different solutions in (a) for values of n up to 4 or 5 by isolating the harmonics on a well-made stringed instrument.)

Now, we come to the truly interesting issue. We have just seen that any function of the form

$$(4) \quad u(x, t) = \sum_{n=1}^N b_n \sin(nx) \cos(nt)$$

solves d'Alembert's *wave equation*, as it is called, but the particular solution we want depends on how the string is originally "plucked." At time $t = 0$, we will assume that the string is given some initial displacement $f(x) = u(x, 0)$. Setting $t = 0$ in our family of solutions in (4), the hope is that the initial displacement function $f(x)$ can be expressed as

$$(5) \quad f(x) = \sum_{n=1}^N b_n \sin(nx).$$

What this means is that *if* there exist suitable coefficients b_1, b_2, \dots, b_N so that $f(x)$ can be written as a sum of sine functions as in (5), then the vibrating-string problem is completely solved by the function $u(x, t)$ given in (4). The obvious question to ask, then, is just what types of functions can be constructed as linear combinations of the functions $\{\sin(x), \sin(2x), \sin(3x), \dots\}$. How general can $f(x)$ be? Daniel Bernoulli (1700–1782) is usually credited with proposing the idea that by taking an *infinite* sum in equation (5), it may be possible to represent *any* initial position $f(x)$ over the interval $[0, \pi]$.

Fourier was studying the propagation of heat when trigonometric series resurfaced in his work in a very similar way. For Fourier, $f(x)$ represented an initial temperature applied to the boundary of some heat-conducting material. The differential equations describing heat flow are slightly different from d'Alembert's wave equation, but they still involve the second derivatives that make expressing $f(x)$ as a sum of trigonometric functions the crucial step in finding a solution.

Periodic Functions

In the early stages of his work, Fourier focused his attention on even functions (i.e., functions satisfying $f(x) = f(-x)$) and sought out ways to represent them

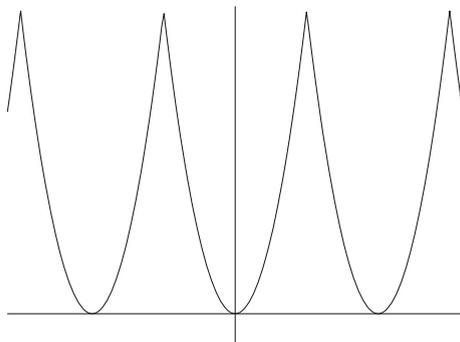


Figure 8.2: $f(x) = x^2$ OVER $(-\pi, \pi]$, EXTENDED TO BE 2π -PERIODIC.

as series of the form $\sum a_n \cos(nx)$. Eventually, he arrived at the more general formulation of the problem, which is to find suitable coefficients (a_n) and (b_n) to express a function $f(x)$ as

$$(6) \quad f(x) = a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx).$$

As we begin to explore how arbitrary $f(x)$ can be, it is important to notice that every component of the series in equation (6) is periodic with period 2π . Turning our attention to the term “function,” it now follows that any function we hope to represent by a trigonometric series will necessarily be periodic as well. We will give primary attention to the interval $(-\pi, \pi]$. What this means is that, given a function such as $f(x) = x^2$, we will restrict our attention to f over the domain $(-\pi, \pi]$ and then extend f periodically to all of \mathbf{R} via the rule $f(x) = f(x + 2k\pi)$ for all $k \in \mathbf{Z}$ (Fig. 8.2).

This convention of focusing on just the part of $f(x)$ over the interval $(-\pi, \pi]$ hardly seems controversial, but it did generate some confusion in Fourier’s time. In Sections 1.2 and 4.1, we alluded to the fact that in the early 1800s the term “function” was used to mean something more like “formula.” It was generally believed that a function’s behavior over the interval $(-\pi, \pi]$ determined its behavior everywhere else, a point of view that follows naturally from an overly zealous faith in Taylor series. The modern definition of function given in Definition 1.2.3 is attributed to Dirichlet from the 1830s, although the idea had been suggested earlier by others. In *Théorie Analytique de la Chaleur*, Fourier clarifies his own use of the term by stating that a “function $f(x)$ represents a succession of values or ordinates, each of which is arbitrary. . . We do not suppose these ordinates to be subject to a common law; they succeed each other in any matter whatever, and each of them is given as if it were a single quantity.”

In the end, we will need to make a few assumptions about the nature of our functions, but the requirements we will need are quite mild, especially when compared with restrictions such as “infinitely differentiable,” which are necessary—but not sufficient—for the existence of a Taylor series representation.

Types of Convergence

This brings us to a discussion of the word “expressed.” The assumptions we must ultimately place on our function depend on the kind of convergence we aim to demonstrate. How are we to understand the equal sign in equation (6)? Our usual course of action with infinite series is first to define the partial sum

$$(7) \quad S_N(x) = a_0 + \sum_{n=1}^N a_n \cos(nx) + b_n \sin(nx).$$

To “express $f(x)$ as a trigonometric series” then means finding coefficients $(a_n)_{n=0}^{\infty}$ and $(b_n)_{n=1}^{\infty}$ so that

$$(8) \quad f(x) = \lim_{N \rightarrow \infty} S_N(x).$$

The question remains as to what kind of limit this is. Fourier probably imagined something akin to a *pointwise* limit because the concept of uniform convergence had not yet been formulated. In addition to pointwise convergence and uniform convergence, there are still other ways to interpret the limit in equation (8). Although it won’t be discussed here, it turns out that proving

$$\int_{-\pi}^{\pi} |S_N(x) - f(x)|^2 dx \rightarrow 0$$

is a natural way to understand equation (8) for a particular class of functions. This is referred to as L^2 convergence. An alternate type of convergence that we will discuss, called *Cesaro mean convergence*, relies on demonstrating that the *averages* of the partial sums converge, in our case uniformly, to $f(x)$.

Fourier Coefficients

In the discussion that follows, we are going to need a few calculus facts.

Exercise 8.5.2. Using trigonometric identities when necessary, verify the following integrals.

(a) For all $n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos(nx) dx = 0 \quad \text{and} \quad \int_{-\pi}^{\pi} \sin(nx) dx = 0.$$

(b) For all $n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos^2(nx) dx = \pi \quad \text{and} \quad \int_{-\pi}^{\pi} \sin^2(nx) dx = \pi.$$

(c) For all $m, n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos(mx) \sin(nx) dx = 0.$$

For $m \neq n$,

$$\int_{-\pi}^{\pi} \cos(mx) \cos(nx) dx = 0 \quad \text{and} \quad \int_{-\pi}^{\pi} \sin(mx) \sin(nx) dx = 0.$$

The consequences of these results are much more interesting than their proofs. The intuition from inner-product spaces is useful. Interpreting the integral as a kind of dot product, this exercise can be summarized by saying that the functions

$$\{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \cos(3x), \dots\}$$

are all *orthogonal* to each other. The content of what follows is that they in fact form a *basis* for a large class of functions.

The first order of business is to deduce some reasonable candidates for the coefficients (a_n) and (b_n) in equation (6). Given a function $f(x)$, the trick is to *assume* we are in possession of a representation described in (6) and then manipulate this equation in a way that leads to formulas for (a_n) and (b_n) . This is exactly how we proceeded with Taylor series expansions in Section 6.6. Taylor's formula for the coefficients was produced by repeatedly differentiating each side of the desired representation equation. Here, we integrate.

To compute a_0 , integrate each side of equation (6) from $-\pi$ to π , brazenly take the integral inside the infinite sum, and use Exercise 8.5.2 to get

$$\begin{aligned} \int_{-\pi}^{\pi} f(x) dx &= \int_{-\pi}^{\pi} \left[a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx) \right] dx \\ &= \int_{-\pi}^{\pi} a_0 dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} [a_n \cos(nx) + b_n \sin(nx)] dx \\ &= a_0(2\pi) + \sum_{n=1}^{\infty} a_n 0 + b_n 0 = a_0(2\pi). \end{aligned}$$

Thus,

$$(9) \quad a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) dx.$$

The switching of the sum and the integral sign in the second step of the previous calculation should rightly raise some eyebrows, but keep in mind that we are really working backward from a hypothetical representation for $f(x)$ to get a *proposal* for what a_0 should be. The point is not to justify the derivation of the formula but rather to show that using this value for a_0 ultimately gives us the representation we want. That hard work lies ahead.

Now, consider a fixed $m \geq 1$. To compute a_m , we first multiply each side of equation (6) by $\cos(mx)$ and again integrate over the interval $[-\pi, \pi]$.

Exercise 8.5.3. Derive the formulas

$$(10) \quad a_m = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos(mx) dx \quad \text{and} \quad b_m = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(mx) dx$$

for all $m \geq 1$.

Let's take a short break and empirically test our recipes for (a_m) and (b_m) on a few simple functions.

Example 8.5.1. Let

$$f(x) = \begin{cases} 1 & \text{if } 0 < x < \pi \\ 0 & \text{if } x = 0 \text{ or } x = \pi \\ -1 & \text{if } -\pi < x < 0. \end{cases}$$

The fact that f is an odd function (i.e., $f(-x) = -f(x)$) means we can avoid doing any integrals for the moment and just appeal to a symmetry argument to conclude

$$a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) dx = 0 \quad \text{and} \quad a_n = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos(nx) dx = 0$$

for all $n \geq 1$. We can also simplify the integral for b_n by writing

$$\begin{aligned} b_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(nx) dx = \frac{2}{\pi} \int_0^{\pi} \sin(nx) dx \\ &= \frac{2}{\pi} \left(\frac{-1}{n} \cos(nx) \Big|_0^{\pi} \right) \\ &= \begin{cases} 4/n\pi & \text{if } n \text{ is odd} \\ 0 & \text{if } n \text{ is even.} \end{cases} \end{aligned}$$

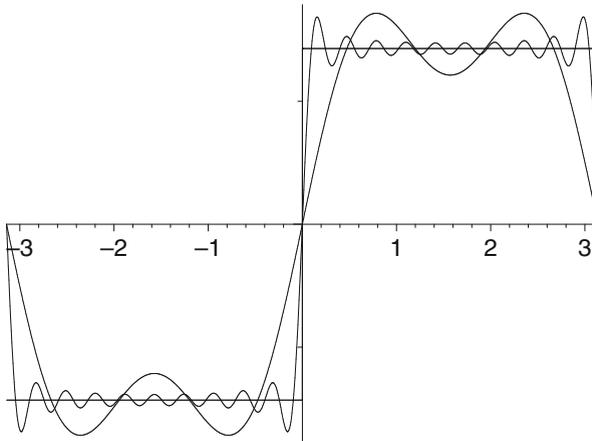


Figure 8.3: f , S_4 , AND S_{20} ON $[-\pi, \pi]$.

Proceeding on blind faith, we plug these results into equation (6) to get the representation

$$f(x) = \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{1}{2n+1} \sin((2n+1)x).$$

A graph of a few of the partial sums of this series (Fig. 8.3) should generate some optimism about the legitimacy of what is happening.

Exercise 8.5.4. (a) Referring to the previous example, explain why we can be sure that the convergence of the partial sums to $f(x)$ is *not* uniform on any interval containing 0.

(b) Repeat the computations of Example 8.5.1 for the function $g(x) = |x|$ and examine graphs for some partial sums. This time, make use of the fact that g is even ($g(x) = g(-x)$) to simplify the calculations. By just looking at the coefficients, how do we know this series converges uniformly to something?

(c) Use graphs to collect some empirical evidence regarding the question of term-by-term differentiation in our two examples to this point. Is it possible to conclude convergence or divergence of either differentiated series by looking at the resulting coefficients? Theorem 6.4.3 is about the legitimacy of term-by-term differentiation. Can it be applied to either of these examples?

The Riemann–Lebesgue Lemma

In the examples we have seen to this point, the sequences of Fourier coefficients (a_n) and (b_n) all tend to 0 as $n \rightarrow \infty$. This is always the case. Understanding why this happens is crucial to our upcoming convergence proof.

We start with a simple observation. The reason

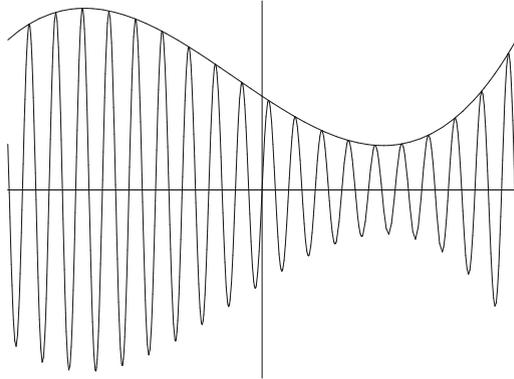
$$\int_{-\pi}^{\pi} \sin(x) dx = 0$$

is that the positive and negative portions of the sine curve cancel each other out. The same is true of

$$\int_{-\pi}^{\pi} \sin(nx) dx = 0.$$

Now, when n is large, the period of the oscillations of $\sin(nx)$ becomes very short— $2\pi/n$ to be precise. If $h(x)$ is a continuous function, then the values of h do not vary too much as $\sin(nx)$ ranges over each short period. The result is that the successive positive and negative oscillations of the product $h(x)\sin(nx)$ (Fig. 8.4) are nearly the same size so that the cancellation leads to a small value for

$$\int_{-\pi}^{\pi} h(x) \sin(nx) dx.$$

Figure 8.4: $h(x)$ AND $h(x) \sin(nx)$ FOR LARGE n .

Theorem 8.5.2 (Riemann–Lebesgue Lemma). Assume $h(x)$ is continuous on $(-\pi, \pi]$. Then,

$$\int_{-\pi}^{\pi} h(x) \sin(nx) dx \rightarrow 0 \quad \text{and} \quad \int_{-\pi}^{\pi} h(x) \cos(nx) dx \rightarrow 0$$

as $n \rightarrow \infty$.

Proof. Remember that, like all of our functions from here on, we are mentally extending h to be 2π -periodic. Thus, while our attention is generally focused on the interval $(-\pi, \pi]$, the assumption of continuity is intended to mean that the periodically extended h is continuous on all of \mathbf{R} . Note that in addition to continuity on $(-\pi, \pi]$, this amounts to insisting that $\lim_{x \rightarrow -\pi^+} h(x) = h(\pi)$.

Exercise 8.5.5. Explain why h is uniformly continuous on \mathbf{R} .

Given $\epsilon > 0$, choose $\delta > 0$ such that $|x - y| < \delta$ implies $|h(x) - h(y)| < \epsilon/2$. The period of $\sin(nx)$ is $2\pi/n$, so choose N large enough so that $\pi/n < \delta$ whenever $n \geq N$. Now, consider a particular interval $[a, b]$ of length $2\pi/n$ over which $\sin(nx)$ moves through one complete oscillation.

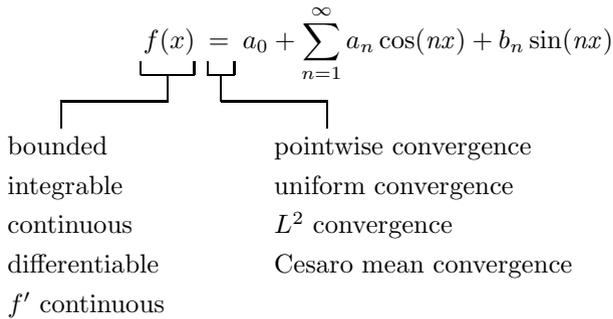
Exercise 8.5.6. Show that $\left| \int_a^b h(x) \sin(nx) dx \right| < \epsilon/n$, and use this fact to complete the proof. \square

Applications of Fourier series are not restricted to continuous functions (Example 8.5.1). Even though our particular proof makes use of continuity, the Riemann–Lebesgue lemma holds under much weaker hypotheses. It is true, however, that any proof of this fact ultimately takes advantage of the cancellation of positive and negative components. Recall from Chapter 2 that this type of cancellation is the mechanism that distinguishes conditional convergence from absolute convergence. In the end, what we discover is that, unlike power series,

Fourier series can converge conditionally. This makes them less robust, perhaps, but more versatile and capable of more interesting behavior.

A Pointwise Convergence Proof

Let's return once more to Fourier's claim that every "function" can be "expressed" as a trigonometric series. Our recipe for the Fourier coefficients in equations (9) and (10) implicitly requires that our function be integrable. This is the major motivation for Riemann's modification of Cauchy's definition of the integral. Because integrability is a prerequisite for producing a Fourier series, we would like the class of integrable functions to be as large as possible. The natural question to ask now is whether Riemann integrability is enough or whether we need to make some additional assumptions about f in order to guarantee that the Fourier series converges back to f . The answer depends on the type of convergence we hope to establish.



There is no tidy way to summarize the situation. For pointwise convergence, integrability is not enough. At present, "integrable" for us means Riemann-integrable, which we have only rigorously defined for bounded functions. In 1966, Lennart Carleson proved (via an extremely complicated argument) that the Fourier series for such a function converges pointwise at every point in the domain excluding possibly a set of *measure zero*. This term surfaced in our discussion of the Cantor set (Section 3.1) and is defined rigorously in Section 7.6. Sets of measure zero are small in one sense, but they can be uncountable, and there are examples of continuous functions with Fourier series that diverge at uncountably many points. Lebesgue's modification of Riemann's integral in 1901 proved to be a much more natural setting for Fourier analysis. Carleson's proof is really about Lebesgue-integrable functions which are allowed to be unbounded but for which $\int_{-\pi}^{\pi} |f|^2$ is finite. One of the cleanest theorems in this area states that, for this class of square Lebesgue-integrable functions, the Fourier series always converges to the function from which it was derived if we interpret convergence in the L^2 sense described earlier. As a final warning

about how fragile the situation is, there is an example due to A. Kolmogorov (1903–1987) of a Lebesgue-integrable function where the Fourier series fails to converge at any point.

Although all of these results require significantly more background to pursue in any rigorous way, we are in a position to prove some important theorems that require a few extra assumptions about the function in question. We will content ourselves with two interesting results in this area.

Theorem 8.5.3. *Let $f(x)$ be continuous on $(-\pi, \pi]$, and let $S_N(x)$ be the N th partial sum of the Fourier series described in equation (7), where the coefficients (a_n) and (b_n) are given by equations (9) and (10). It follows that*

$$\lim_{N \rightarrow \infty} S_N(x) = f(x)$$

pointwise at any $x \in (-\pi, \pi]$ where $f'(x)$ exists.

Proof. Cataloging a few preliminary facts makes for a smoother argument.

Fact 1: (a) $\cos(\alpha - \theta) = \cos(\alpha)\cos(\theta) + \sin(\alpha)\sin(\theta)$.

(b) $\sin(\alpha + \theta) = \sin(\alpha)\cos(\theta) + \cos(\alpha)\sin(\theta)$.

Fact 2: $\frac{1}{2} + \cos(\theta) + \cos(2\theta) + \cos(3\theta) + \cdots + \cos(N\theta) = \frac{\sin((N + 1/2)\theta)}{2 \sin(\theta/2)}$ for any $\theta \neq 2n\pi$.

Facts 1(a) and 1(b) are familiar trigonometric identities. Fact 2 is not as familiar. Its proof (which we omit) is most easily derived by taking the real part of a geometric sum of complex exponentials. The function in Fact 2 is called the *Dirichlet kernel* in honor of the mathematician responsible for the first rigorous convergence proof of this kind. Integrating both sides of this identity leads to our next important fact.

Fact 3: Setting

$$D_N(\theta) = \begin{cases} \frac{\sin((N+1/2)\theta)}{2 \sin(\theta/2)}, & \text{if } \theta \neq 2n\pi \\ 1/2 + N, & \text{if } \theta = 2n\pi \end{cases}$$

from Fact 2, we see that

$$\int_{-\pi}^{\pi} D_N(\theta) d\theta = \pi.$$

Although we will not restate it, the last fact we will use is the Riemann–Lebesgue Lemma.

Fix a point $x \in (-\pi, \pi]$. The first step is to simplify the expression for $S_N(x)$. Now x is a fixed constant at the moment, so we will write the integrals in equations (9) and (10) using t as the variable of integration. Keeping an eye

on Facts 1(a) and (2), we get that

$$\begin{aligned}
 S_N(x) &= a_0 + \sum_{n=1}^N a_n \cos(nx) + b_n \sin(nx) \\
 &= \left[\frac{1}{2\pi} \int_{-\pi}^{\pi} f(t) dt \right] + \sum_{n=1}^N \left[\frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \cos(nt) dt \right] \cos(nx) \\
 &\quad + \sum_{n=1}^N \left[\frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \sin(nt) dt \right] \sin(nx) \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\frac{1}{2} + \sum_{n=1}^N \cos(nt) \cos(nx) + \sin(nt) \sin(nx) \right] dt \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\frac{1}{2} + \sum_{n=1}^N \cos(nt - nx) \right] dt \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) D_N(t - x) dt.
 \end{aligned}$$

As one final simplification, let $u = t - x$. Then,

$$S_N(x) = \frac{1}{\pi} \int_{-\pi-x}^{\pi-x} f(u+x) D_N(u) du = \frac{1}{\pi} \int_{-\pi}^{\pi} f(u+x) D_N(u) du.$$

The last equality is a result of our agreement to extend f to be 2π -periodic. Because D_N is also periodic (it is the sum of cosine functions), it does not matter over what interval we compute the integral as long as we cover exactly one full period.

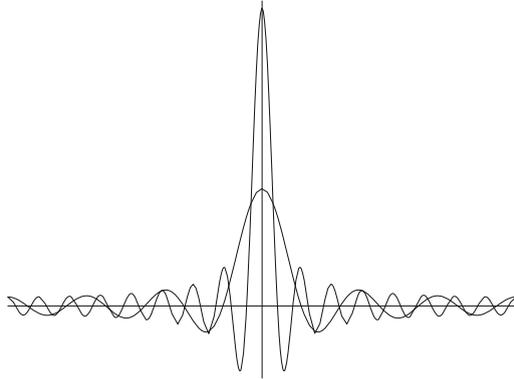
To prove $S_N(x) \rightarrow f(x)$, we must show that $|S_N(x) - f(x)|$ gets arbitrarily small when N gets large. Having expressed $S_N(x)$ as an integral involving $D_N(u)$, we are motivated to do a similar thing for $f(x)$. By Fact 3,

$$f(x) = f(x) \frac{1}{\pi} \int_{-\pi}^{\pi} D_N(u) du = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) D_N(u) du,$$

and it follows that

$$(11) \quad S_N(x) - f(x) = \frac{1}{\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) D_N(u) du.$$

Our goal is to show this quantity tends to zero as $N \rightarrow \infty$. A sketch of $D_N(u)$ (Fig. 8.5) for a few values of N reveals why this might happen. For large N , the Dirichlet kernel $D_N(u)$ has a tall, thin spike around $u = 0$, but this is precisely where $f(u+x) - f(x)$ is small (because f is continuous). Away from zero, $D_N(u)$ exhibits the fast oscillations that hearken back to the Riemann–Lebesgue Lemma (Theorem 8.5.2). Let's see how to use this theorem to finish the argument.

Figure 8.5: $D_6(u)$ AND $D_{16}(u)$.

Using Fact 1(b), we can rewrite the Dirichlet kernel as

$$D_N(u) = \frac{\sin((N + 1/2)u)}{2 \sin(u/2)} = \frac{1}{2} \left[\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} + \cos(Nu) \right].$$

Then, equation (11) becomes

$$\begin{aligned} S_N(x) - f(x) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) \left[\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} + \cos(Nu) \right] du \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) \left(\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} \right) \\ &\quad + (f(u+x) - f(x)) \cos(Nu) du \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} p_x(u) \sin(Nu) du + \frac{1}{2\pi} \int_{-\pi}^{\pi} q_x(u) \cos(Nu) du, \end{aligned}$$

where in the last step we have set

$$p_x(u) = \frac{(f(u+x) - f(x)) \cos(u/2)}{\sin(u/2)} \quad \text{and} \quad q_x(u) = f(u+x) - f(x).$$

Exercise 8.5.7. (a) First, argue why the integral involving $q_x(u)$ tends to zero as $N \rightarrow \infty$.

(b) The first integral is a little more subtle because the function $p_x(u)$ has the $\sin(u/2)$ term in the denominator. Use the fact that f is differentiable at x (and a familiar limit from calculus) to prove that the first integral goes to zero as well. \square

This completes the argument that $S_N(x) \rightarrow f(x)$ at any point x where f is differentiable. If the derivative exists everywhere, then we get $S_N \rightarrow f$

pointwise. If we add the assumption that f' is continuous, then it is not too difficult to show that the convergence is uniform. In fact, there is a very strong relationship between the speed of convergence of the Fourier series and the smoothness of f . The more derivatives f possesses, the faster the partial sums S_N converge to f .

Cesaro Mean Convergence

Rather than pursue the proofs in this interesting direction, we will finish this very brief introduction to Fourier series with a look at a different type of convergence called Cesaro mean convergence.

Exercise 8.5.8. Prove that if a sequence of real numbers (x_n) converges, then the arithmetic means

$$y_n = \frac{x_1 + x_2 + x_3 + \cdots + x_n}{n}$$

also converge to the same limit. Give an example to show that it is possible for the sequence of means (y_n) to converge even if the original sequence (x_n) does not.

The discussion preceding Theorem 8.5.3 is intended to create a kind of reverence for the difficulties inherent in deciphering the behavior of Fourier series, especially in the case where the function in question is not differentiable. It is from this humble frame of mind that the following elegant result due to L. Fejér in 1904 can best be appreciated.

Theorem 8.5.4 (Fejér's Theorem). Let $S_n(x)$ be the n th partial sum of the Fourier series for a function f on $(-\pi, \pi]$. Define

$$\sigma_N(x) = \frac{1}{N+1} \sum_{n=0}^N S_n(x).$$

If f is continuous on $(-\pi, \pi]$, then $\sigma_N(x) \rightarrow f(x)$ uniformly.

Proof. This argument is patterned after the proof of Theorem 8.5.3 but is actually much simpler. In addition to the trigonometric formulas listed in Facts 1 and 2, we are going to need a version of Fact 2 for the sine function, which looks like

$$\sin(\theta) + \sin(2\theta) + \sin(3\theta) + \cdots + \sin(N\theta) = \frac{\sin\left(\frac{N\theta}{2}\right) \sin\left((N+1)\frac{\theta}{2}\right)}{\sin\left(\frac{\theta}{2}\right)}.$$

Exercise 8.5.9. Use the previous identity to show that

$$\frac{1/2 + D_1(\theta) + D_2(\theta) + \cdots + D_N(\theta)}{N+1} = \frac{1}{2(N+1)} \left[\frac{\sin\left((N+1)\frac{\theta}{2}\right)}{\sin\left(\frac{\theta}{2}\right)} \right]^2.$$

The expression in Exercise 8.5.9 is called the *Fejér kernel* and will be denoted by $F_N(\theta)$. Analogous to the Dirichlet kernel $D_N(\theta)$ from the proof of Theorem 8.5.3, F_N is used to greatly simplify the formula for $\sigma_N(x)$.

Exercise 8.5.10. (a) Show that

$$\sigma_N(x) = \frac{1}{\pi} \int_{-\pi}^{\pi} f(u+x)F_N(u) du.$$

(b) Graph the function $F_N(u)$ for several values of N . Where is F_N large, and where is it close to zero? Compare this function to the Dirichlet kernel $D_N(u)$. Now, prove that $F_N \rightarrow 0$ uniformly on any set of the form $\{u : |u| \geq \delta\}$, where $\delta > 0$ is fixed (and u is restricted to the interval $(-\pi, \pi]$).

(c) Prove that $\int_{-\pi}^{\pi} F_N(u) du = \pi$.

(d) To finish the proof of Fejér's Theorem, first choose a $\delta > 0$ so that

$$|u| < \delta \quad \text{implies} \quad |f(x+u) - f(x)| < \epsilon.$$

Set up a single integral that represents the difference $\sigma_N(x) - f(x)$ and divide this integral into sets where $|u| \leq \delta$ and $|u| \geq \delta$. Explain why it is possible to make each of these integrals sufficiently small, independently of the choice of x . \square

Weierstrass Approximation Theorem

The hard work of proving Fejér's Theorem has many rewards, one of which is access to a relatively short argument for a profoundly important theorem discovered by Weierstrass in 1885. The Weierstrass Approximation Theorem (WAT) is studied in depth in Section 6.7 and is restated here for ease of reference.

Theorem 6.7.1 (Weierstrass Approximation Theorem). *Let $f : [a, b] \rightarrow \mathbf{R}$ be continuous. Given $\epsilon > 0$, there exists a polynomial $p(x)$ satisfying*

$$|f(x) - p(x)| < \epsilon$$

for all $x \in [a, b]$.

Proof. We have actually seen a few special cases of this result before in Section 6.6 on Taylor series. For instance, we showed that

$$\sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \frac{x^9}{9!} - \cdots,$$

where the series converges uniformly on any bounded subset of \mathbf{R} . Uniform convergence of a series means the partial sums converge uniformly, and the

partial sums in this case are polynomials. Notice that this is precisely what WAT asks us to prove, only we must do it for an arbitrary, continuous function in place of $\sin(x)$.

Using Taylor series does not work in general. To construct a Taylor series we need the function to be infinitely differentiable—not just continuous—and even in this case we might get a series that either does not converge or converges to the wrong thing. Taylor series are a valuable tool, however. In Section 6.7 we used the Taylor series for $\sqrt{1-x}$ as the starting point for a proper proof of WAT. Fejér's Theorem, in conjunction with the Taylor series for $\sin(x)$ and $\cos(x)$, provides a significant shortcut to the same result.

Exercise 8.5.11. (a) Use the fact that the Taylor series for $\sin(x)$ and $\cos(x)$ converge uniformly on any compact set to prove WAT under the added assumption that $[a, b]$ is $[0, \pi]$.

(b) Show how the case for an arbitrary interval $[a, b]$ follows from this one. □

A comment from Section 6.7 that bears repeating relates to the striking contrast between this result and Weierstrass's demonstration of a continuous nowhere-differentiable function. Although there exist continuous functions that oscillate so wildly that they fail to have a derivative at any point, these unruly functions are always uniformly within ϵ of an infinitely differentiable polynomial.

Approximation as a Unifying Theme

Viewing the last section of this chapter as a kind of appendix (included to clear up some loose ends from Chapter 1 regarding the definition of the real numbers), the Weierstrass' Approximation Theorem makes for a fitting close to our introductory survey of some of the gems of analysis.

The idea of approximation permeates the entire subject. Every real number can be approximated with rational ones. The value of an infinite sum is approximated with partial sums, and the value of a continuous function can be approximated with its values nearby. A function is differentiable when a straight line is a good approximation to the curve, and it is integrable when finite sums of rectangles are a good approximation to the area under the curve. Now, we learn that every continuous function can be approximated arbitrarily well with a polynomial. In every case, the approximating objects are tangible and well-understood, and the issue is how well these properties survive the limiting process. By viewing the different infinities of mathematics through pathways crafted out of finite objects, Weierstrass and the other founders of analysis created a paradigm for how to extend the scope of mathematical exploration deep into territory previously unattainable. Although our journey ends here, the road is long and continues to be written.

8.6 A Construction of \mathbf{R} From \mathbf{Q}

This entire section is devoted to constructing a proof for the following theorem:

Theorem 8.6.1 (Existence of the Real Numbers). *There exists an ordered field in which every nonempty set that is bounded above has a least upper bound. In addition, this field contains \mathbf{Q} as a subfield.*

There are a few terms to define before this statement can be properly understood and proved, but it can essentially be paraphrased as “the real numbers exist.” In Section 1.1, we encountered a major failing of the rational number system as a place to do analysis. Without the square root of 2 (and uncountably many other irrational numbers) we cannot confidently move from a Cauchy sequence to its limit because in \mathbf{Q} there is no guarantee that such a number exists. (A review of Sections 1.1 and 1.3 is highly recommended at this point.) The resolution we proposed in Chapter 1 came in the form of the Axiom of Completeness, which we restate.

Axiom of Completeness. *Every nonempty set of real numbers that is bounded above has a least upper bound.*

Now let’s be clear about how we actually proceeded in Chapter 1. This is the property that distinguishes \mathbf{Q} from \mathbf{R} , but by referring to this property as an *axiom* we were making the point that it was not something to be proved. The real numbers were defined simply as an extension of the rational numbers in which bounded sets have least upper bounds, but no attempt was made to demonstrate that such an extension is actually possible. Now, the time has finally come. By explicitly building the real numbers from the rational ones, we will be able to demonstrate that the Axiom of Completeness does not need to be an axiom at all; it is a theorem!

There is something ironic about having the final section of this book be a construction of the number system that has been the underlying subject of every preceding page, but there is something perfectly apt about it as well. Through eight chapters stretching from Cantor’s Theorem to the Baire Category Theorem, we have come to see how profoundly the addition of completeness changes the landscape. We all grow up believing in the existence of real numbers, but it is only through a study of classical analysis that we become aware of their elusive and enigmatic nature. It is because completeness matters so much, and because it is responsible for such perplexing phenomena, that we should now feel obliged—compelled really—to go back to the beginning and verify that such a thing really exists.

As we mentioned in Chapter 1, proceeding in this order puts us in good historical company. The pioneering work of Cauchy, Bolzano, Abel, Dirichlet, Weierstrass, and Riemann preceded—and in a very real sense led to—the host of rigorous definitions for \mathbf{R} that were proposed in the last half of the 19th century. Georg Cantor is a familiar name responsible for one of these definitions, but alternate constructions of the real number system also came from Charles

Meray (1835–1911), Eduard Heine (1821–1881), and Richard Dedekind (1831–1916). The formulation that follows is the one due to Dedekind. In a sense it is the most abstract of the approaches, but it is the most appropriate for us because the verification of completeness is done in terms of least upper bounds.

Dedekind Cuts

We begin this discussion by assuming that the rational numbers and all of the familiar properties of addition, multiplication, and order are available to us. At the moment, there is no such thing as a real number.

Definition 8.6.2. A subset A of the rational numbers is called a *cut* if it possesses the following three properties:

- (c1) $A \neq \emptyset$ and $A \neq \mathbf{Q}$.
- (c2) If $r \in A$, then A also contains every rational $q < r$.
- (c3) A does not have a maximum; that is, if $r \in A$, then there exists $s \in A$ with $r < s$.

Exercise 8.6.1. (a) Fix $r \in \mathbf{Q}$. Show that the set $C_r = \{t \in \mathbf{Q} : t < r\}$ is a cut.

The temptation to think of all cuts as being of this form should be avoided. Which of the following subsets of \mathbf{Q} are cuts?

- (b) $S = \{t \in \mathbf{Q} : t \leq 2\}$
- (c) $T = \{t \in \mathbf{Q} : t^2 < 2 \text{ or } t < 0\}$
- (d) $U = \{t \in \mathbf{Q} : t^2 \leq 2 \text{ or } t < 0\}$

Exercise 8.6.2. Let A be a cut. Show that if $r \in A$ and $s \notin A$, then $r < s$.

To dispel any suspense, let's get right to the point.

Definition 8.6.3. Define the *real numbers* \mathbf{R} to be the set of all cuts in \mathbf{Q} .

This may feel awkward at first—real numbers should be numbers, not sets of rational numbers. The counterargument here is that when working on the foundations of mathematics, sets are about the most basic building blocks we have. We have defined a set \mathbf{R} whose elements are subsets of \mathbf{Q} . We now must set about the task of imposing some algebraic structure on \mathbf{R} that behaves in a way familiar to us. What exactly does this entail? If we are serious about constructing a proof for Theorem 8.6.1, we need to be more specific about what we mean by an “ordered field.”

Field and Order Properties

Given a set F and two elements $x, y \in F$, an *operation* on F is a function that takes the ordered pair (x, y) to a third element $z \in F$. Writing $x + y$ or xy to represent different operations reminds us of the two operations that we are trying to emulate.

Definition 8.6.4. A set F is a *field* if there exist two operations—addition $(x + y)$ and multiplication (xy) —that satisfy the following list of conditions:

- (f1) (commutativity) $x + y = y + x$ and $xy = yx$ for all $x, y \in F$.
- (f2) (associativity) $(x + y) + z = x + (y + z)$ and $(xy)z = x(yz)$ for all $x, y, z \in F$.
- (f3) (identities exist) There exist two special elements 0 and 1 with $0 \neq 1$ such that $x + 0 = x$ and $x1 = x$ for all $x \in F$.
- (f4) (inverses exist) Given $x \in F$, there exists an element $-x \in F$ such that $x + (-x) = 0$. If $x \neq 0$, there exists an element x^{-1} such that $xx^{-1} = 1$.
- (f5) (distributive property) $x(y + z) = xy + xz$ for all $x, y, z \in F$.

Exercise 8.6.3. Using the usual definitions of addition and multiplication, determine which of these properties are possessed by \mathbf{N} , \mathbf{Z} , and \mathbf{Q} , respectively.

Although we will not pursue this here in any depth, all of the familiar algebraic manipulations in \mathbf{Q} (e.g., $x + y = x + z$ implies $y = z$) can be derived from this short list of properties.

Definition 8.6.5. An *ordering* on a set F is a relation, represented by \leq , with the following three properties:

- (o1) For arbitrary $x, y \in F$, at least one of the statements $x \leq y$ or $y \leq x$ is true.
- (o2) If $x \leq y$ and $y \leq x$, then $x = y$.
- (o3) If $x \leq y$ and $y \leq z$, then $x \leq z$.

We will sometimes write $y \geq x$ in place of $x \leq y$. The strict inequality $x < y$ is used to mean $x \leq y$ but $x \neq y$.

A field F is called an *ordered field* if F is endowed with an ordering \leq that satisfies

- (o4) If $y \leq z$, then $x + y \leq x + z$.
- (o5) If $x \geq 0$ and $y \geq 0$, then $xy \geq 0$.

Let's take stock of where we are. To prove Theorem 8.6.1, we are accepting as given that the rational numbers are an ordered field. We have defined the real numbers \mathbf{R} to be the collection of cuts in \mathbf{Q} , and the challenge now is to invent addition, multiplication, and an ordering so that each possesses the properties

outlined in the preceding two definitions. The easiest of these is the ordering. Let A and B be two arbitrary elements of \mathbf{R} .

Define $A \leq B$ to mean $A \subseteq B$.

Exercise 8.6.4. Show that this defines an ordering on \mathbf{R} by verifying properties (o1), (o2), and (o3) from Definition 8.6.5.

Algebra in \mathbf{R}

Given A and B in \mathbf{R} , define

$$A + B = \{a + b : a \in A \text{ and } b \in B\}.$$

Before checking properties (f1)–(f4) for addition, we must first verify that our definition really defines an operation. Is $A + B$ actually a cut? To get the flavor of how these arguments look, let's verify property (c2) of Definition 8.6.2 for the set $A + B$.

Let $a + b \in A + B$ be arbitrary and let $s \in \mathbf{Q}$ satisfy $s < a + b$. Then, $s - b < a$, which implies that $s - b \in A$ because A is a cut. But then

$$s = (s - b) + b \in A + B,$$

and (c2) is proved.

Exercise 8.6.5. (a) Show that (c1) and (c3) also hold for $A + B$. Conclude that $A + B$ is a cut.

(b) Check that addition in \mathbf{R} is commutative (f1) and associative (f2).

(c) Show that property (o4) holds.

(d) Show that the cut

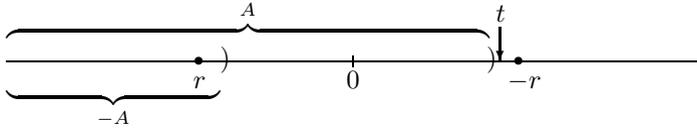
$$O = \{p \in \mathbf{Q} : p < 0\}$$

successfully plays the role of the additive identity (f3). (Showing $A + O = A$ amounts to proving that these two sets are the same. The standard way to prove such a thing is to show two inclusions: $A + O \subseteq A$ and $A \subseteq A + O$.)

What about additive inverses? Given $A \in \mathbf{R}$, we must produce a cut $-A$ with the property that $A + (-A) = O$. This is a bit more difficult than it sounds. Conceptually, the cut $-A$ consists of all rational numbers less than $-\sup A$. The problem is how to define this set without using suprema, which are strictly off limits at the moment. (We are building the field in which they exist!)

Given $A \in \mathbf{R}$, define

$$-A = \{r \in \mathbf{Q} : \text{there exists } t \notin A \text{ with } t < -r\}.$$



Exercise 8.6.6. (a) Prove that $-A$ defines a cut.

(b) What goes wrong if we set $-A = \{r \in \mathbf{Q} : -r \notin A\}$?

(c) If $a \in A$ and $r \in -A$, show $a + r \in O$. This shows $A + (-A) \subseteq O$. Now, finish the proof of property (f4) for addition in Definition 8.6.4.

Although the ideas are similar, the technical difficulties increase when we try to create a definition for multiplication in \mathbf{R} . This is largely due to the fact that the product of two negative numbers is positive. The standard method of attack is first to define multiplication on the non-negative cuts.

Given $A \geq O$ and $B \geq O$ in \mathbf{R} , define the product

$$AB = \{ab : a \in A, b \in B \text{ with } a, b \geq 0\} \cup \{q \in \mathbf{Q} : q < 0\}.$$

Exercise 8.6.7. (a) Show that AB is a cut and that property (o5) holds.

(b) Propose a good candidate for the multiplicative identity (1) on \mathbf{R} and show that this works for all cuts $A \geq O$.

(c) Show the distributive property (f5) holds for non-negative cuts.

Products involving at least one negative factor can be defined in terms of the product of two positive cuts by observing that $-A \geq 0$ whenever $A \leq O$. (Given $A \leq O$, property (o4) implies $A + (-A) \leq O + (-A)$, which yields $O \leq -A$.)

For any A and B in \mathbf{R} , define

$$AB = \begin{cases} \text{as given} & \text{if } A \geq O \text{ and } B \geq O \\ -[A(-B)] & \text{if } A \geq O \text{ and } B < O \\ -[(-A)B] & \text{if } A < O \text{ and } B \geq O \\ (-A)(-B) & \text{if } A < O \text{ and } B < O. \end{cases}$$

Verifying that multiplication defined in this way satisfies all the required field properties is important but uneventful. The proofs generally fall into cases for when terms are positive or negative and follow a pattern similar to those for addition. We will leave them as an unofficial exercise and move on to the punch line.

Least Upper Bounds

Having proved that \mathbf{R} is an ordered field, we now set our sights on showing that this field is complete. We defined completeness in Chapter 1 in terms of least upper bounds. Here is a summary of the relevant definitions from that discussion.

Definition 8.6.6. A set $\mathcal{A} \subseteq \mathbf{R}$ is *bounded above* if there exists a $B \in \mathbf{R}$ such that $A \leq B$ for all $A \in \mathcal{A}$. The number B is called an *upper bound* for \mathcal{A} .

A real number $S \in \mathbf{R}$ is the *least upper bound* for a set $\mathcal{A} \subseteq \mathbf{R}$ if it meets the following two criteria:

- (i) S is an upper bound for \mathcal{A} and
- (ii) if B is any upper bound for \mathcal{A} , then $S \leq B$.

Exercise 8.6.8. Let $\mathcal{A} \subseteq \mathbf{R}$ be nonempty and bounded above, and let S be the *union* of all $A \in \mathcal{A}$.

- (a) First, prove that $S \in \mathbf{R}$ by showing that it is a cut.
- (b) Now, show that S is the least upper bound for \mathcal{A} .

This finishes the proof that \mathbf{R} is complete. Notice that we could have proved that least upper bounds exist immediately after defining the ordering on \mathbf{R} , but saving it for last gives it the privileged place in the argument it deserves. There is, however, still one loose end to sew up. The statement of Theorem 8.6.1 mentions that our complete ordered field contains \mathbf{Q} as a subfield. This is a slight abuse of language. What it should say is that \mathbf{R} contains a subfield that looks and acts exactly like \mathbf{Q} .

Exercise 8.6.9. Consider the collection of so-called “rational” cuts of the form

$$C_r = \{t \in \mathbf{Q} : t < r\}$$

where $r \in \mathbf{Q}$. (See Exercise 8.6.1.)

- (a) Show that $C_r + C_s = C_{r+s}$ for all $r, s \in \mathbf{Q}$. Verify $C_r C_s = C_{rs}$ for the case when $r, s \geq 0$.
- (b) Show that $C_r \leq C_s$ if and only if $r \leq s$ in \mathbf{Q} .

Cantor’s Approach

As a way of giving Georg Cantor the last word, let’s briefly look at his very different approach to constructing \mathbf{R} out of \mathbf{Q} . One of the many equivalent ways to characterize completeness is with the assertion that “Cauchy sequences converge.” Given a Cauchy sequence of rational numbers, we are now well aware that this sequence may converge to a value not in \mathbf{Q} . Just as before, the goal is to create something, which we will call a *real number*, that can serve as the limit of this sequence. Cantor’s idea was essentially to define a real number to be the entire Cauchy sequence. The first problem one encounters with this approach is the realization that two different Cauchy sequences can converge to the same real number. For this reason, the elements in \mathbf{R} are more appropriately defined as *equivalence classes* of Cauchy sequences where two sequences (x_n) and (y_n) are in the same equivalence class if and only if $(x_n - y_n) \rightarrow 0$.

As with Dedekind's approach, it can be momentarily disorienting to supplant our relatively simple notion of a real number as a decimal expansion with something as unruly as an equivalence class of Cauchy sequences. But what exactly do we mean by a decimal expansion? And how are we to understand the number $1/2$ as both $.5000\dots$ and $.4999\dots$? We leave it as an exercise.