

Chapter 2


Sequences and Series

2.1 Discussion: Rearrangements of Infinite Series

Consider the infinite series

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \dots$$

If we naively begin adding from the left-hand side, we get a sequence of what are called *partial sums*. In other words, let s_n equal the sum of the first n terms of the series, so that $s_1 = 1$, $s_2 = 1/2$, $s_3 = 5/6$, $s_4 = 7/12$, and so on. One immediate observation is that the successive sums oscillate in a progressively narrower space. The odd sums decrease ($s_1 > s_3 > s_5 > \dots$) while the even sums increase ($s_2 < s_4 < s_6 < \dots$).



$$s_2 < s_4 < s_6 < \dots < S < \dots < s_5 < s_3 < s_1$$

It seems reasonable—and we will soon prove—that the sequence (s_n) eventually hones in on a value, call it S , where the odd and even partial sums “meet.” At this moment, we cannot compute S precisely, but we know it falls somewhere between $7/12$ and $5/6$. Summing a few hundred terms reveals that $S \approx .69$. Whatever its value, there is now an overwhelming temptation to write

$$(1) \quad S = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \dots$$

meaning, perhaps, that if we could indeed add up *all* infinitely many of these numbers, then the sum would *equal* S . A more familiar example of an equation of this type might be

$$2 = 1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} + \frac{1}{64} + \cdots,$$

the only difference being that in the second equation we have a more recognizable value for the sum.

But now for the crux of the matter. The symbols $+$, $-$, and $=$ in the preceding equations are deceptively familiar notions being used in a very unfamiliar way. The crucial question is whether or not properties of addition and equality that are well understood for finite sums remain valid when applied to infinite objects such as equation (1). The answer, as we are about to witness, is somewhat ambiguous.

Treating equation (1) in a standard algebraic way, let's multiply through by $1/2$ and add it back to equation (1):

$$\begin{aligned} \frac{1}{2}S &= \quad \frac{1}{2} & -\frac{1}{4} & +\frac{1}{6} & -\frac{1}{8} & +\frac{1}{10} & -\frac{1}{12} & +\cdots \\ + S &= 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \frac{1}{9} - \frac{1}{10} + \frac{1}{11} - \frac{1}{12} + \frac{1}{13} - \cdots \\ \hline (2) \quad \frac{3}{2}S &= 1 & +\frac{1}{3} - \frac{1}{2} + \frac{1}{5} & +\frac{1}{7} - \frac{1}{4} + \frac{1}{9} & +\frac{1}{11} - \frac{1}{6} + \frac{1}{13} & \cdots \end{aligned}$$

Now, look carefully at the result. The sum in equation (2) consists *precisely* of the same terms as those in the original equation (1), only in a different order. Specifically, the series in (2) is a rearrangement of (1) where we list the first two positive terms ($1 + \frac{1}{3}$) followed by the first negative term ($-\frac{1}{2}$), followed by the next two positive terms ($\frac{1}{5} + \frac{1}{7}$) and then the next negative term ($-\frac{1}{4}$). Continuing this, it is apparent that every term in (2) appears in (1) and vice versa. The rub comes when we realize that equation (2) asserts that the sum of these rearranged, but otherwise unaltered, numbers is equal to $3/2$ its original value. Indeed, adding a few hundred terms of equation (2) produces partial sums in the neighborhood of 1.03. Addition, in this infinite setting, is not commutative!

Let's look at a similar rearrangement of the series

$$\sum_{n=0}^{\infty} (-1/2)^n.$$

This series is geometric with first term 1 and common ratio $r = -1/2$. Using the formula $1/(1-r)$ for the sum of a geometric series (Example 2.7.5), we get

$$1 - \frac{1}{2} + \frac{1}{4} - \frac{1}{8} + \frac{1}{16} - \frac{1}{32} + \frac{1}{64} - \frac{1}{128} + \frac{1}{256} \cdots = \frac{1}{1 - (-\frac{1}{2})} = \frac{2}{3}.$$

This time, some computational experimentation with the "two positives, one negative" rearrangement

$$1 + \frac{1}{4} - \frac{1}{2} + \frac{1}{16} + \frac{1}{64} - \frac{1}{8} + \frac{1}{256} + \frac{1}{1024} - \frac{1}{32} \cdots$$

yields partial sums quite close to $2/3$. The sum of the first 30 terms, for instance, equals .666667. Infinite addition is commutative in some instances but not in others.

Far from being a charming theoretical oddity of infinite series, this phenomenon can be the source of great consternation in many applied situations. How, for instance, should a double summation over two index variables be defined? Let's say we are given a *grid* of real numbers $\{a_{ij} : i, j \in \mathbf{N}\}$, where $a_{ij} = 1/2^{j-i}$ if $j > i$, $a_{ij} = -1$ if $j = i$, and $a_{ij} = 0$ if $j < i$.

$$\begin{bmatrix} -1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{8} & \frac{1}{16} & \cdots \\ 0 & -1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{8} & \cdots \\ 0 & 0 & -1 & \frac{1}{2} & \frac{1}{4} & \cdots \\ 0 & 0 & 0 & -1 & \frac{1}{2} & \cdots \\ 0 & 0 & 0 & 0 & -1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

We would like to attach a mathematical meaning to the summation

$$\sum_{i,j=1}^{\infty} a_{ij}$$

whereby we intend to include every term in the preceding array in the total. One natural idea is to temporarily fix i and sum across each row. A moment's reflection (and a fact about geometric series) shows that each row sums to 0. Summing the sums of the rows, we get

$$\sum_{i,j=1}^{\infty} a_{ij} = \sum_{i=1}^{\infty} \left(\sum_{j=1}^{\infty} a_{ij} \right) = \sum_{i=1}^{\infty} (0) = 0.$$

We could just as easily have decided to fix j and sum down each column first. In this case, we have

$$\sum_{i,j=1}^{\infty} a_{ij} = \sum_{j=1}^{\infty} \left(\sum_{i=1}^{\infty} a_{ij} \right) = \sum_{j=1}^{\infty} \left(\frac{-1}{2^{j-1}} \right) = -2.$$

Changing the order of the summation changes the value of the sum! One common way that double sums arise (although not this particular one) is from the multiplication of two series. There is a natural desire to write

$$\left(\sum a_i \right) \left(\sum b_j \right) = \sum_{i,j} a_i b_j,$$

except that the expression on the right-hand side makes no sense at the moment.

It is the pathologies that give rise to the need for rigor. A satisfying resolution to the questions raised will require that we be absolutely precise about what we mean as we manipulate these infinite objects. It may seem that progress is slow at first, but that is because we do not want to fall into the trap of letting the biases of our intuition corrupt our arguments. Rigorous proofs are meant to be a check on intuition, and in the end we will see that they vastly improve our mental picture of the mathematical infinite.

As a final example, consider something as intuitively fundamental as the associative property of addition applied to the series $\sum_{n=1}^{\infty} (-1)^n$. Grouping the terms one way gives

$$(-1 + 1) + (-1 + 1) + (-1 + 1) + (-1 + 1) + \cdots = 0 + 0 + 0 + 0 + \cdots = 0,$$

whereas grouping in another yields

$$-1 + (1 - 1) + (1 - 1) + (1 - 1) + \cdots = -1 + 0 + 0 + 0 + \cdots = -1.$$

Manipulations that are legitimate in finite settings do not always extend to infinite settings. Deciding when they do and why they do not is one of the central themes of analysis.

2.2 The Limit of a Sequence

An understanding of infinite series depends heavily on a clear understanding of the theory of sequences. In fact, most of the concepts in analysis can be reduced to statements about the behavior of sequences. Thus, we will spend a significant amount of time investigating sequences before taking on infinite series.

Definition 2.2.1. A *sequence* is a function whose domain is \mathbf{N} .

This formal definition leads immediately to the familiar depiction of a sequence as an ordered list of real numbers. Given a function $f : \mathbf{N} \rightarrow \mathbf{R}$, $f(n)$ is just the n th term on the list. The notation for sequences reinforces this familiar understanding.

Example 2.2.2. Each of the following are common ways to describe a sequence.

- (i) $(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \cdots)$,
- (ii) $(\frac{1+n}{n})_{n=1}^{\infty} = (\frac{2}{1}, \frac{3}{2}, \frac{4}{3}, \cdots)$,
- (iii) (a_n) , where $a_n = 2^n$ for each $n \in \mathbf{N}$,
- (iv) (x_n) , where $x_1 = 2$ and $x_{n+1} = \frac{x_n+1}{2}$.

On occasion, it will be more convenient to index a sequence beginning with $n = 0$ or $n = n_0$ for some natural number n_0 different from 1. These minor variations should cause no confusion. What is essential is that a sequence be an *infinite* list of real numbers. What happens at the beginning of such a list is of

little importance in most cases. The business of analysis is concerned with the behavior of the infinite “tail” of a given sequence.

We now present what is arguably the most important definition in the book.

Definition 2.2.3 (Convergence of a Sequence). A sequence (a_n) converges to a real number a if, for every positive number ϵ , there exists an $N \in \mathbf{N}$ such that whenever $n \geq N$ it follows that $|a_n - a| < \epsilon$.

To indicate that (a_n) converges to a , we usually write either $\lim a_n = a$ or $(a_n) \rightarrow a$. The notation $\lim_{n \rightarrow \infty} a_n = a$ is also standard.


In an effort to decipher this complicated definition, it helps first to consider the ending phrase “ $|a_n - a| < \epsilon$,” and think about the points that satisfy an inequality of this type.

Definition 2.2.4. Given a real number $a \in \mathbf{R}$ and a positive number $\epsilon > 0$, the set

$$V_\epsilon(a) = \{x \in \mathbf{R} : |x - a| < \epsilon\}$$


is called the ϵ -neighborhood of a .

Notice that $V_\epsilon(a)$ consists of all of those points whose distance from a is less than ϵ . Said another way, $V_\epsilon(a)$ is an interval, centered at a , with radius ϵ .



Recasting the definition of convergence in terms of ϵ -neighborhoods gives a more geometric impression of what is being described.

Definition 2.2.3B (Convergence of a Sequence: Topological Version). A sequence (a_n) converges to a if, given any ϵ -neighborhood $V_\epsilon(a)$ of a , there exists a point in the sequence after which all of the terms are in $V_\epsilon(a)$. In other words, every ϵ -neighborhood contains all but a finite number of the terms of (a_n) .



Definition 2.2.3 and Definition 2.2.3B say precisely the same thing; the natural number N in the original version of the definition is the point where the sequence (a_n) enters $V_\epsilon(a)$, never to leave. It should be apparent that *the value of N depends on the choice of ϵ* . The smaller the ϵ -neighborhood, the larger N may have to be.

Example 2.2.5. Consider the sequence (a_n) , where $a_n = 1/\sqrt{n}$.

Our intuitive understanding of limits points confidently to the conclusion that

$$\lim \left(\frac{1}{\sqrt{n}} \right) = 0.$$

Before trying to prove this not too impressive fact, let's first explore the relationship between ϵ and N in the definition of convergence. For the moment, take ϵ to be $1/10$. This defines a sort of "target zone" for the terms in the sequence. By claiming that the limit of (a_n) is 0, we are saying that the terms in this sequence eventually get arbitrarily close to 0. How close? What do we mean by "eventually"? We have set $\epsilon = 1/10$ as our standard for closeness, which leads to the ϵ -neighborhood $(-1/10, 1/10)$ centered around the limit 0. How far out into the sequence must we look before the terms fall into this interval? The 100th term $a_{100} = 1/10$ puts us right on the boundary, and a little thought reveals that

$$\text{if } n > 100, \quad \text{then } a_n \in \left(-\frac{1}{10}, \frac{1}{10} \right).$$

Thus, for $\epsilon = 1/10$ we choose $N = 101$ (or anything larger) as our response.

Now, our choice of $\epsilon = 1/10$ was rather whimsical, and we can do this again, letting $\epsilon = 1/50$. In this case, our target neighborhood shrinks to $(-1/50, 1/50)$, and it is apparent that we must travel farther out into the sequence before a_n falls into this interval. How far? Essentially, we require that

$$\frac{1}{\sqrt{n}} < \frac{1}{50} \quad \text{which occurs as long as } n > 50^2 = 2500.$$

Thus, $N = 2501$ is a suitable response to the challenge of $\epsilon = 1/50$.

It may seem as though this duel could continue forever, with different ϵ challenges being handed to us one after another, each one requiring a suitable value of N in response. In a sense, this is correct, except that the game is effectively over the instant we recognize a *rule* for how to choose N given an *arbitrary* $\epsilon > 0$. For this problem, the desired algorithm is implicit in the algebra carried out to compute the previous response of $N = 2501$. Whatever ϵ happens to be, we want

$$\frac{1}{\sqrt{n}} < \epsilon \quad \text{which is equivalent to insisting that } n > \frac{1}{\epsilon^2}.$$

With this observation, we are ready to write the formal argument.

We claim that

$$\lim \left(\frac{1}{\sqrt{n}} \right) = 0.$$

Proof. Let $\epsilon > 0$ be an arbitrary positive number. Choose a natural number N satisfying

$$N > \frac{1}{\epsilon^2}.$$

We now verify that this choice of N has the desired property. Let $n \geq N$. Then,

$$n > \frac{1}{\epsilon^2} \quad \text{implies} \quad \frac{1}{\sqrt{n}} < \epsilon, \quad \text{and hence} \quad |a_n - 0| < \epsilon. \quad \square$$

Quantifiers

The definition of convergence given earlier is the result of hundreds of years of refining the intuitive notion of limit into a mathematically rigorous statement. The logic involved is complicated and is intimately tied to the use of the quantifiers “for all” and “there exists.” Learning to write a grammatically correct convergence proof goes hand in hand with a deep understanding of why the quantifiers appear in the order that they do.

The definition begins with the phrase,

“For all $\epsilon > 0$, there exists $N \in \mathbf{N}$ such that ...”

Looking back at our first example, we see that our formal proof begins with, “Let $\epsilon > 0$ be an arbitrary positive number.” This is followed by a construction of N and then a demonstration that this choice of N has the desired property. This, in fact, is a basic outline for how every convergence proof should be presented.

TEMPLATE FOR A PROOF THAT $(x_n) \rightarrow x$:

- “Let $\epsilon > 0$ be arbitrary.”
- Demonstrate a choice for $N \in \mathbf{N}$. This step usually requires the most work, almost all of which is done prior to actually writing the formal proof.
- Now, show that N actually works.
- “Assume $n \geq N$.”
- With N well chosen, it should be possible to derive the inequality $|x_n - x| < \epsilon$.

Example 2.2.6. Show

$$\lim \left(\frac{n+1}{n} \right) = 1.$$

As mentioned, before attempting a formal proof, we first need to do some preliminary scratch work. In the first example, we experimented by assigning specific values to ϵ (and it is not a bad idea to do this again), but let us skip straight to the algebraic punch line. The last line of our proof should be that for suitably large values of n ,

$$\left| \frac{n+1}{n} - 1 \right| < \epsilon.$$

Because

$$\left| \frac{n+1}{n} - 1 \right| = \frac{1}{n},$$

this is equivalent to the inequality $1/n < \epsilon$ or $n > 1/\epsilon$. Thus, choosing N to be an integer greater than $1/\epsilon$ will suffice.

With the work of the proof done, all that remains is the formal writeup.

Proof. Let $\epsilon > 0$ be arbitrary. Choose $N \in \mathbf{N}$ with $N > 1/\epsilon$. To verify that this choice of N is appropriate, let $n \in \mathbf{N}$ satisfy $n \geq N$. Then, $n \geq N$ implies $n > 1/\epsilon$, which is the same as saying $1/n < \epsilon$. Finally, this means

$$\left| \frac{n+1}{n} - 1 \right| < \epsilon,$$

as desired. □

It is instructive to see what goes wrong in the previous example if we try to prove that our sequence converges to some limit other than 1.

Theorem 2.2.7 (Uniqueness of Limits). *The limit of a sequence, when it exists, must be unique.*

Proof. Exercise 2.2.6. □

Divergence

Significant insight into the role of the quantifiers in the definition of convergence can be gained by studying an example of a sequence that does not have a limit.

Example 2.2.8. Consider the sequence

$$\left(1, -\frac{1}{2}, \frac{1}{3}, -\frac{1}{4}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \dots \right).$$

How can we argue that this sequence does not converge to zero? Looking at the first few terms, it seems the initial evidence actually supports such a conclusion. Given a challenge of $\epsilon = 1/2$, a little reflection reveals that after $N = 3$ all the terms fall into the neighborhood $(-1/2, 1/2)$. We could also handle $\epsilon = 1/4$. (What is the smallest possible N in this case?)

But the definition of convergence says “*For all* $\epsilon > 0 \dots$,” and it should be apparent that there is no response to a choice of $\epsilon = 1/10$, for instance. This leads us to an important observation about the logical negation of the definition of convergence of a sequence. To prove that a particular number x is *not* the limit of a sequence (x_n) , we must produce a single value of ϵ for which no $N \in \mathbf{N}$ works. More generally speaking, the negation of a statement that begins “For all P, there exists Q...” is the statement, “For at least one P, no Q is possible...” For instance, how could we disprove the spurious claim that “At every college in the United States, there is a student who is at least seven feet tall”?

We have argued that the preceding sequence does not converge to 0. Let's argue against the claim that it converges to $1/5$. Choosing $\epsilon = 1/10$ produces the neighborhood $(1/10, 3/10)$. Although the sequence continually revisits this neighborhood, there is no point at which it enters and never leaves as the definition requires. Thus, no N exists for $\epsilon = 1/10$, so the sequence does not converge to $1/5$.

Of course, this sequence does not converge to any other real number, and it would be more satisfying to simply say that this sequence does not converge.

Definition 2.2.9. A sequence that does not converge is said to *diverge*.

Although it is not too difficult, we will postpone arguing for divergence in general until we develop a more economical divergence criterion later in Section 2.5.

Exercises

Exercise 2.2.1. What happens if we reverse the order of the quantifiers in Definition 2.2.3?

Definition: A sequence (x_n) *verconges* to x if there exists an $\epsilon > 0$ such that for all $N \in \mathbf{N}$ it is true that $n \geq N$ implies $|x_n - x| < \epsilon$.

Give an example of a vercongent sequence. Is there an example of a vercongent sequence that is divergent? Can a sequence verconge to two different values? What exactly is being described in this strange definition?

Exercise 2.2.2. Verify, using the definition of convergence of a sequence, that the following sequences converge to the proposed limit.

(a) $\lim_{n \rightarrow \infty} \frac{2n+1}{5n+4} = \frac{2}{5}$.

(b) $\lim_{n \rightarrow \infty} \frac{2n^2}{n^3+3} = 0$.

(c) $\lim_{n \rightarrow \infty} \frac{\sin(n^2)}{\sqrt[3]{n}} = 0$.

Exercise 2.2.3. Describe what we would have to demonstrate in order to disprove each of the following statements.

- At every college in the United States, there is a student who is at least seven feet tall.
- For all colleges in the United States, there exists a professor who gives every student a grade of either A or B.
- There exists a college in the United States where every student is at least six feet tall.

Exercise 2.2.4. Give an example of each or state that the request is impossible. For any that are impossible, give a compelling argument for why that is the case.

- A sequence with an infinite number of ones that does not converge to one.

- (b) A sequence with an infinite number of ones that converges to a limit not equal to one.
- (c) A divergent sequence such that for every $n \in \mathbf{N}$ it is possible to find n consecutive ones somewhere in the sequence.

Exercise 2.2.5. Let $\llbracket x \rrbracket$ be the greatest integer less than or equal to x . For example, $\llbracket \pi \rrbracket = 3$ and $\llbracket 3 \rrbracket = 3$. For each sequence, find $\lim a_n$ and verify it with the definition of convergence.

- (a) $a_n = \llbracket 5/n \rrbracket$,
- (b) $a_n = \llbracket (12 + 4n)/3n \rrbracket$.

Reflecting on these examples, comment on the statement following Definition 2.2.3 that “the smaller the ϵ -neighborhood, the larger N may have to be.”

Exercise 2.2.6. Prove Theorem 2.2.7. To get started, assume $(a_n) \rightarrow a$ and also that $(a_n) \rightarrow b$. Now argue $a = b$.

Exercise 2.2.7. Here are two useful definitions:

- (i) A sequence (a_n) is *eventually* in a set $A \subseteq \mathbf{R}$ if there exists an $N \in \mathbf{N}$ such that $a_n \in A$ for all $n \geq N$.
- (ii) A sequence (a_n) is *frequently* in a set $A \subseteq \mathbf{R}$ if, for every $N \in \mathbf{N}$, there exists an $n \geq N$ such that $a_n \in A$.
- (a) Is the sequence $(-1)^n$ eventually or frequently in the set $\{1\}$?
- (b) Which definition is stronger? Does frequently imply eventually or does eventually imply frequently?
- (c) Give an alternate rephrasing of Definition 2.2.3B using either frequently or eventually. Which is the term we want?
- (d) Suppose an infinite number of terms of a sequence (x_n) are equal to 2. Is (x_n) necessarily eventually in the interval $(1.9, 2.1)$? Is it frequently in $(1.9, 2.1)$?

Exercise 2.2.8. For some additional practice with nested quantifiers, consider the following invented definition:

Let’s call a sequence (x_n) *zero-heavy* if there exists $M \in \mathbf{N}$ such that for all $N \in \mathbf{N}$ there exists n satisfying $N \leq n \leq N + M$ where $x_n = 0$.

- (a) Is the sequence $(0, 1, 0, 1, 0, 1, \dots)$ zero heavy?
- (b) If a sequence is zero-heavy does it necessarily contain an infinite number of zeros? If not, provide a counterexample.
- (c) If a sequence contains an infinite number of zeros, is it necessarily zero-heavy? If not, provide a counterexample.
- (d) Form the logical negation of the above definition. That is, complete the sentence: A sequence is *not* zero-heavy if . . .