Appendix A

Proof Techniques

In this chapter we'll discuss how to read a proof, some of the more common ways to prove statements, and highlight a few ways that do not work and should be avoided at all costs! Students are often frustrated when they transition from the more standard courses such as calculus, where there aren't too many theorems and most of the exercises are mechanical and straightforward (if the homework problem is from the integration by parts section, it's pretty clear what you're going to need to do to evaluate the integral!), to upper level classes, where frequently the proof of the main theorem of a section is left as an exercise! Even if you happen to be lucky enough to have a book which gives a proof, it's easy to lose the forest in the trees. What this means is that as you're reading the proof you can understand each line in isolation. You can understand how they go from one line to the next; however, it's a complete mystery how the author decided that it would be good to go from *this* line to *that* line, and you're rightly a bit terrified about your turn at proving something, as then you'll be responsible for directing the flow. Learning how to see these paths, learning what's a good next step, is hard, but doing so is essential for your growth in mathematics. The aim below is to describe in detail many of the common methods, in the hope that learning these will help you in following and creating proofs.

We cover the following proof techniques below.

1. Proof by Induction.
2. Proof by Grouping.
3. Proof by Exploiting Symmetry.
4. Proof by Brute Force.
5. Proof by Comparison or Story.
6. Proof by Contradiction.
7. Proof by Exhaustion (also known as Divide and Conquer).
8. Proof by Counterexample.

9. Proof by Generalizing Example.
10. Proof by Pigeon-Hole Principle.
11. Proof by Adding Zero or Multiplying by One.

A.1 How to read a proof

Frequently in books you'll find a square, such as □, at the end of the proof. This is meant to alert you that the argument is done, and the claim has been shown. This is done because all too often we're so caught up in following the arguments from line to line that we don't realize we've reached the end! Other people write qed or Q.E.D., which is an abbreviation of the Latin phrase quod erat demonstrandum, which means *that which was to be demonstrated*. Sometimes authors also write '*Proof*' at the start of the argument. These are done to help clue you in to what's going on. This helps prevent the proof from blending in with the rest of the text.

Before diving into proof techniques, here's some general advice on reading proofs. On a first pass through a proof don't be too concerned with mastering all the details. Rather, just look for a broad overview of what's happening or being discussed, and don't worry if you're unable to follow the argument from line to line.

Step one is to make sure you understand the conditions and the claim. If you can, take a few examples of objects that satisfy the conditions, and see that the claim is true for them. Sometimes it's particularly good, or at least easy, to try extreme examples. If you have a continuous function, try a constant function. Try a wildly oscillating one like $x^2 \sin(1/x)$, or perhaps one that isn't differentiable at a point, such as $|x|$. Also try a few examples that don't satisfy the assumptions of the theorem. In this case, the claim may or may not hold. Doing a few checks like this can give you a feel for what's going on, and as you do your checks you might start to see what will be needed in the proof. This is especially true when you find examples that don't satisfy the claim, as somehow the assumptions of the theorem must prohibit bad cases like this from happening.

After trying to get a feel of which examples work and which don't, return and think deeply about the assumptions. How are the assumptions used in the argument? When you read the assumptions, your first thoughts should be: okay, so what theorems do I know that require these conditions? For example, if one of your assumptions is that your function is differentiable, maybe you start to think about using the Mean Value Theorem. Or perhaps you've assumed f is a polynomial of degree n; in that case, the Fundamental Theorem of Algebra tells you that f has n complex roots. The more you know, the easier this becomes. I do a lot of work in number theory; if I'm told I have two relatively prime numbers x and y, my first thought is the Euclidean algorithm, which says there are integers a and b such that $ax + by = 1$ (for example, 17 and 11 are relatively prime, and $2 \cdot 17 - 3 \cdot 11 = 1$). Why is this my first thought? Experience – I've done so many problems that I know this is often a great way to start. The more you do, the easier it becomes. Assumptions are signposts, they're markers to help direct the flow of the proof. Time spent thinking about them and what they entail is time well spent.

What if the assumptions aren't used? Well, something strange is happening, as why would they be given if they're not used? What's more likely to happen is that sometimes the assumptions aren't truly needed, but are given to allow you access to powerful other theorems to simplify the proof. For example, the proof of the Fundamental Theorem of Calculus only needs our function f to be continuous on a finite interval $[a, b]$; however, whenever I teach calculus I always assume f' exists, is continuous, and is bounded. This is *not* needed, but assuming it simplifies the proof. If you continue deeper in mathematics, you'll revisit old theorems in a quest to have the weakest conditions possible and still get the same result.

Finally, when reading the proof don't worry about understanding every justification. First skim the argument, trying to get a sense of the main ideas. What results were used in the proof? Roughly, why were we able to use these? Sometimes there are lots of technical conditions that need to be met to invoke a theorem. In these cases, a lot of the proof is devoted to showing these conditions are met. When reading the proof for the first time, it's fine to gloss over these parts. Think something like: okay, we need to show the quotient is a finitely generated Abelian group, and the next few lines do that, I'll take their word on it for now. Later, of course, you should go back and try to understand these justifications, but don't obsess too much, as that can lead you to losing the flow of the proof. Often books and papers remove these mini-arguments and isolate them, either before or after the proof, calling them lemmas. A lemma is a smaller result, a building block to the proof of the main claim. Sometimes authors put these first so that by the time you get to the theorem you've seen everything you'll need. Other times these are placed afterwards, to avoid interrupting the flow with technicalities. Each approach is fine.

One last remark about reading proofs. Eventually you'll come across the phrase **without loss of generality**. Typically, this is followed by the author doing one case and saying the other cases follow similarly. If you're new to doing proofs, you should do all the cases in full glory (or is it gory?). These four words can be very dangerous, as sometimes there *are* differences between the various cases, and the only way to be sure is to do each case. If everything checks out and the arguments really are the same, mathematicians will often just give the details for one to save space and time. If you read a proof invoking this claim, it's good practice to fill in the details for the other cases.

Okay, we're now ready to explore different proof method!

A.2 Proofs by Induction

This section is an expansion of an appendix from [MT-B] on Proofs by Induction. This method is designed to handle the following situation: for each positive integer n we have some statement $P(n)$, and we desire to show that $P(n)$ is true for all n. For example, maybe $P(n)$ is the statement that the sum of the first n odd numbers is always a perfect square. One possibility is to start evaluating it for different choices of n. In this case, we get 1, 4, 9 and 16 for the first four values, and we're feeling confident that the result is true. However, confidence is *not* the same as a proof, and just because it worked for the first few values doesn't mean it'll continue to work. If we eventually find an n such that $P(n)$ fails, then we know the statement is false. What if, however, it always holds for every value we check. Does this mean it must be true? Sadly, it doesn't. It may be we just haven't checked far enough.

For an example of what can go wrong, let's consider a famous polynomial,

n	f(n)	Primality of $f(n)$
1	41	prime
2	43	prime
3	47	prime
$\vdots$	$\vdots$	$\vdots$
37	1447	prime
38	1523	prime

Table A.1: Values of the polynomial $f(x) = x^2 + x + 1$.

n	Sum of first n odd numbers	Value of the sum
1	1	1
2	1+3	4
3	1+3+5	9
4	$1 + 3 + 5 + 7$	16
$\vdots$	$\vdots$	$\vdots$
100	$1 + 3 + \cdots + 197 + 199$	10000

Table A.2: Sums of odd integers.

$f(x) = x^2 + x + 41$. Euler was interested in this polynomial, and you'll see why in a moment. Let's look at some of its values, which we record in Table A.1.

In the interest of space we only recorded a subset of the values of $f(n)$; a little work shows that $f(n)$ is also prime for *all* n up to 38. Based on the data above, it's natural to conjecture that $f(n)$ is *always* prime for any positive integer n. While the data suggests this, testing some values isn't a proof.

How should we proceed? We can take larger and larger values of n and see what happens. In this case, we would find $f(39) = 1601$ is prime, but $f(40) = 1681 = 41^2$ is composite, as is $f(49) = 2491 = 47 \cdot 53$. Here, we were able to go far enough to see the pattern break down, and once we have one value that fails we know the claim cannot always hold.

For another example, let's revisit the sum of the first n odd integers. We can make a similar table as before, which we do in Table A.2.

Do you see the pattern? It looks like the sum of the first n odd integers is just n^2. Unlike the previous example, this time our conjecture is true. No matter how far we check, we'll see the pattern hold; however, just observing this equality *is not* a proof.

We need a way to prove statements like this and others. We quickly describe a powerful method, called Proofs by Induction, that works for a variety of problems. The general framework is that we have some statement $P(n)$ which we want to determine whether or not it holds for all positive integers n.

Proof Technique: Proofs by Induction: A statement $P(n)$ is true for all positive integers n if the following two conditions hold.

- **Basis Step:** $P(1)$ is true;

- **Inductive Step**: whenever $P(n)$ is true, $P(n+1)$ is true.

Proof by Induction is a very useful method for proving results; we'll see many instances of this in this appendix. The reason the method works follows from basic logic. We assume the following two sentences are true:

$$P(1) \text{ is true.}$$
$$\text{For all } n \geq 1, P(n) \text{ is true implies } P(n+1) \text{ is true.}$$

Set $n = 1$ in the second statement. As $P(1)$ is true, and $P(1)$ implies $P(2)$, $P(2)$ must be true. Now set $n = 2$ in the second statement. As $P(2)$ is true, and $P(2)$ implies $P(3)$, $P(3)$ must be true. And so on, completing the proof. Verifying the first statement is called **basis step**, and the second the **inductive step**. In verifying the inductive step, note we assume $P(n)$ is true; this is called the **inductive assumption**. Sometimes instead of starting at $n = 1$ we start at $n = 0$, although in general we could start at any n_0 and then prove for all $n \geq n_0$, $P(n)$ is true.

We give four of the more standard examples of proofs by induction in the next subsections, and one false example; the first example is the most typical. When you have mastered proofs by induction, you might want to return to the problem below. It's a fun problem involving the Fibonacci numbers.

Problem A.2.1 (Zeckendorf's Theorem) *Consider the set of distinct Fibonacci numbers: $\{1, 2, 3, 5, 8, 13, \ldots\}$, where $F_{n+2} = F_{n+1} + F_n$. Show every positive integer can be written uniquely as a sum of distinct Fibonacci numbers where we do not allow two consecutive Fibonacci numbers to occur in the decomposition. Equivalently, for any n there are choices of $\epsilon_i(n) \in \{0, 1\}$ such that*

$$n = \sum_{i=1}^{\ell(n)} \epsilon_i(n) F_i, \quad \epsilon_i(n)\epsilon_{i+1}(n) = 0 \text{ for } i \in \{1, \ldots, \ell(n) - 1\}.$$

Does a similar result hold for all recurrence relations? If not, can you find another recurrence relation where such a result holds?

A.2.1 Sums of Integers

Let $P(n)$ be the statement

$$\sum_{k=1}^{n} k = \frac{n(n+1)}{2}.$$

Here we're using summation notation, which is a very compact way of writing expressions. Unwinding, the left hand side is just

$$\sum_{k=1}^{n} k = 1 + 2 + \cdots + k.$$

More generally,

$$\sum_{k=1}^{n} a_k \ = \ a_1 + a_2 + \cdots + a_n.$$

This is probably the most famous of all examples for proofs by induction. The great Gauss is said to have successfully evaluated this sum when he was five years old. According to the story, his teacher was having a bad day (we all do), and wanted some busywork to occupy the children; he did not count on having a budding master mathematician in the room!

Anyway, let's show that the statement is true by induction. We have two things to check, the basis step (or the base case), and the inductive step (or induction case). Let's go!

Proof: We proceed by induction.
Basis Step: $P(1)$ is clearly true, as both sides equal 1.

Inductive Step: Assuming $P(n)$ is true, we must show $P(n+1)$ is true. By the inductive assumption, $\sum_{k=1}^{n} k = \frac{n(n+1)}{2}$. Thus

$$\begin{aligned}
\sum_{k=1}^{n+1} k &= 1 + 2 + \cdots + n + (n+1) \\
&= (1 + 2 + \cdots + n) + (n+1) \\
&= \left(\sum_{k=1}^{n} k\right) + (n+1) \\
&= \frac{n(n+1)}{2} + (n+1) \\
&= \frac{(n+1)(n+1+1)}{2}.
\end{aligned}$$

Thus, given $P(n)$ is true, then $P(n+1)$ is true. □

You might have seen the above example in a calculus class when studying area under curves. This (and the sum in the exercise below) arise in computing the upper and lower sums.

Note how the argument proceeded above. The hard part was showing that if $P(n)$ held then $P(n+1)$ holds too. The way we did this was to look at our expression for $P(n+1)$ and note that there was a $P(n)$ hiding inside it. We then used the fact that $P(n)$ was assumed true to rewrite $P(n+1)$, and then did some simple algebra. Many, many inductions proceed like this. The trick is finding out how to easily work in the induction assumption; however, if you're attempting a proof by induction then you should be on the watch for such an opportunity. The whole point of induction is to build on results for smaller n, so you should try to find the $P(n)$ case lurking in the $P(n+1)$ expression.

Problem A.2.2 *Prove*

$$\sum_{k=1}^{n} k^2 \ = \ \frac{n(n+1)(2n+1)}{6}.$$

Find a similar formula for the sum of k^3. For the brave, find a similar formula for the sum of k^4. Hint: the sum of the d^{th} powers of integers up to n is a polynomial in n of degree $d+1$.

Problem A.2.3 *Show the sum of the first n odd numbers is n^2, i.e.,*

$$\sum_{k=1}^{n}(2k-1) \;=\; n^2.$$

In the last exercise above, there are two ways to write an odd number. We chose to write the odd numbers as $2k-1$, as this allowed our index k to range from 1 to n (if we want the first n odd numbers). If instead we write odd numbers as $2m+1$, then m would range from 0 to $n-1$ to give the first n odd integers. Either method is fine; the only difference is whether or not you want the index of summation to be nice (from 1 to n) or if you want to avoid the minus sign in the summands.

A.2.2 Divisibility

We now consider a divisibility problem. This is another example of a proof by induction, but the algebra and analysis is a little different, which is why we want to give these arguments too.

Let $P(n)$ be the statement 133 *divides* $11^{n+1}+12^{2n-1}$. *We prove this claim by induction.*

Proof: We proceed by induction.
Basis Step: A straightforward calculation shows $P(1)$ is true: $11^{1+1}+12^{2-1} = 121+12 = 133$.

Inductive Step: Assume $P(n)$ is true, i.e., 133 divides $11^{n+1}+12^{2n-1}$. We must show $P(n+1)$ is true, or that 133 divides $11^{(n+1)+1}+12^{2(n+1)-1}$. But

$$\begin{aligned}
11^{(n+1)+1}+12^{2(n+1)-1} \quad &= \quad 11^{n+1+1}+12^{2n-1+2}\\
&= \quad 11\cdot 11^{n+1}+12^2\cdot 12^{2n-1}\\
&= \quad 11\cdot 11^{n+1}+(133+11)12^{2n-1}\\
&= \quad 11\left(11^{n+1}+12^{2n-1}\right)+133\cdot 12^{2n-1}.
\end{aligned}$$

By the inductive assumption 133 divides $11^{n+1}+12^{2n-1}$; therefore, 133 divides $11^{(n+1)+1}+12^{2(n+1)-1}$, completing the proof. □

The difficulty in this proof was noting that 133 and 11 were lurking together in 144. Specifically, we could write 144 as 133 plus 11. The reason this helps is that the other term is multiplied by 11, and by cleverly re-grouping we saw $11^{n+1}+12^{2n-1}$. It was a very good idea to rewrite 11^{n+2} as $11\cdot 11^{n+1}$ (and similarly for the expression involving 12), and it was reasonable to try this as we wanted to 'see' $P(n)$. In fact, staring at this and thinking back to the sum of integers, we see that *both* proofs had us finding $P(n)$ somewhere in $P(n+1)$. Many induction problems require you to find $P(n)$ lurking in $P(n+1)$; it's not surprising that this happens, as the whole point of induction arguments is to assume $P(n)$ is true and then show this implies $P(n+1)$ holds.

Problem A.2.4 *Prove* 4 *divides* $1 + 3^{2n+1}$.

Problem A.2.5 *Find a positive integer* a *such that* 5 *divides* $1 + 4^{an}$ *for all* n*, and prove your claim.*

A.2.3 The Binomial Theorem

We end with one more example of a proof by induction, the proof of the Binomial Theorem. This time the result is *clearly* of importance for a probability class. The Binomial Theorem is used all the time; in fact, we even have binomial random variables!

Before stating and proving the result, we first recall the definition and some properties of binomial coefficients.

Definition A.2.6 (Binomial Coefficients) *Let* n *and* k *be integers with* $0 \le k \le n$. *We set*

$$\binom{n}{k} \ = \ \frac{n!}{k!(n-k)!}.$$

Note that $0! = 1$. *We set* $\binom{n}{k} = 0$ *if* $k > n$.

The combinatorial interpretation of $\binom{n}{k}$ is that this is the number of ways of choosing k people from n when order doesn't matter, and $m!$ is the number of ways of ordering m. It may seem strange to say $0! = 1$, but if we use these interpretations we could read this as saying there are no ways to order an empty set of people. If you don't remember the proofs of these statements, they're given below in §A.5 and §A.6.

We're now ready to state the Binomial Theorem.

Theorem A.2.7 (The Binomial Theorem) *For all positive integers* n *we have*

$$(x+y)^n \ = \ \sum_{k=0}^{n} \binom{n}{k} x^{n-k} y^k.$$

Proof of the Binomial Theorem: We proceed by induction.

Basis Step: For $n = 1$ we have

$$\sum_{k=0}^{1} \binom{1}{k} x^{1-k} y^k \ = \ \binom{1}{0} x + \binom{1}{1} y \ = \ (x+y)^1.$$

Inductive Step: Suppose

$$(x+y)^n \;=\; \sum_{k=0}^{n} \binom{n}{k} x^{n-k} y^k. \tag{A.1}$$

Then using Lemma A.5.1 we find that

$$\begin{aligned}
(x+y)^{n+1} \quad &= \quad (x+y)(x+y)^n \\
&= \quad (x+y)\sum_{k=0}^{n}\binom{n}{k}x^{n-k}y^k \\
&= \quad \sum_{k=0}^{n}\binom{n}{k}\left[x^{n+1-k}y^k + \binom{n}{k}x^{n-k}y^{k+1}\right] \\
&= \quad x^{n+1} + \sum_{k=1}^{n}\left[\binom{n}{k} + \binom{n}{k-1}\right]x^{n+1-k}y^k + y^{n+1} \\
&= \quad \sum_{k=0}^{n+1}\binom{n+1}{k}x^{n+1-k}y^k,
\end{aligned}$$

as $x^{n+1} = \binom{n+1}{0}x^{n+1}$ and $y^{n+1} = \binom{n+1}{n+1}y^{n+1}$. This establishes the induction step, and hence the theorem. □

As always, the hardest part of the proof is figuring out how to use the inductive assumption. The main idea here was to write $(x+y)^{n+1}$ as $(x+y)(x+y)^n$; this is a 'natural' thing to do, as we now have a factor of $(x+y)^n$, which by the inductive assumption we know how to handle. Of course, we could also have written it as $(x+y)^n(x+y)$, and the proof would have similar. This is almost always the goal: find a way to rewrite the expression so you can exploit the inductive assumption. The most troublesome part of this problem is having to adjust the index of summation (if you continue to differential equations, you'll get a lot of practice with this when you do series expansions). A good guideline is to try to make all terms look the same. We thus want the powers of x and y to look the same in each expression, and this helps us figure out how to shift. Typically it's preferable to have the powers of x and y the same and the index of the coefficients different than the other way around.

There are other ways to prove the Binomial Theorem; we'll see one in §A.6 where we do proofs by comparison.

A.2.4 Fibonacci numbers modulo 2

If it's 10 o'clock now, most people would have no difficulty saying that in 5 hours it'll be 3 o'clock. If we look at what we've just said, are we saying 10 plus 5 is 3? On a clock with twelve hours: yes! The idea of **clock** or **modulo arithmetic** plays a central role in much of number theory, and generalizes nicely. We say x **is congruent to** y **modulo** n if $x-y$ is divisible by n. Thus, 15 is congruent to 3 modulo 12, as 12 divides 15-3. Similarly we find 67 is equivalent to 7 modulo 12, as $60 = 5 \cdot 12 + 7$. We write 67 modulo $12 = 7$ or $67 = 7 \bmod 12$.

Let's look at a fun problem involving the Fibonacci numbers. Recall these are defined by $F_0 = 0$, $F_1 = 1$ and $F_{n+2} = F_{n+1} + F_n$. The first few are

$$0,\ 1,\ 1,\ 2,\ 3,\ 5,\ 8,\ 13,\ 21,\ 34,\ 55,\ 89,\ \dots.$$

Let's look at these numbers modulo 2. A little inspection shows us that x modulo 2 is 0 if x is even (and thus a multiple of 2) and 1 if x is odd. The Fibonacci numbers modulo 2 are

$$0,\ 1,\ 1,\ 0,\ 1,\ 1,\ 0,\ 1,\ 1,\ 0,\ 1,\ 1,\ \dots.$$

Looking at this, we see the beginning of a pattern. It seems to be repeating blocks of 0, 1, 1. Does this always continue? It does, and one nice way to prove this is by induction.

Proof that the Fibonacci numbers modulo 2 are the repeating sequence 0, 1, 1, 0, 1, 1, ... We proceed by induction.

Base Step: Just calculating the first few terms verifies that it does start 0, 1, 1, 0, 1, 1.

Inductive Step: The defining property of the Fibonacci numbers is that the two previous terms are added to get the next. It's thus natural to investigate whether or not this holds modulo 2. In other words, is F_{n+2} modulo 2 the same as F_{n+1} modulo 2 plus F_n modulo 2, all of this modulo 2? Unwinding, there are four cases:

- If F_n is even and F_{n-1} is even, is F_{n+2} even?
- If F_n is even and F_{n-1} is odd, is F_{n+2} odd?
- If F_n is odd and F_{n-1} is even, is F_{n+2} odd?
- If F_n is odd and F_{n-1} is odd, is F_{n+2} even?

The four statements are true, and can be verified with a little bit of algebra (at the level of odd plus odd is even, odd plus even is even, even plus even is even). Armed with this, we can now complete the proof. Assume the first k blocks of three are 0, 1, 1; we'll denote this by

$$0,\ 1,\ 1,\ \dots,\ 0,\ 1,\ 1.$$

Let's look at the next three terms of the Fibonacci numbers modulo 2. The next number is the sum of the two previous modulo 2, so the next number is $1 + 1$ modulo 2, which is zero. Thus our sequence is now

$$0,\ 1,\ 1,\ \dots,\ 0,\ 1,\ 1,\ 0.$$

The next term is just 1+0 modulo 2, which is one, implying our sequence is

$$0,\ 1,\ 1,\ \dots,\ 0,\ 1,\ 1,\ 0,\ 1.$$

The next term is just 1 modulo 2, which is 1 again, giving us

$$0,\ 1,\ 1,\ \dots,\ 0,\ 1,\ 1,\ 0,\ 1,\ 1.$$

This is exactly what we wanted to prove – we just showed that if the first k blocks of three are 0, 1, 1 then the next block is also 0, 1, 1. This completes the proof. □

There are lots of wonderful patterns that emerge when looking at interesting sequences modulo primes (2 is the smallest prime). We urge you to google the pattern for Pascal's triangle modulo 2 – the resulting pattern is quite surprising!

A.2.5 False Proofs by Induction

After seeing how powerful proofs by induction can be, it's a good idea to be aware of the pitfalls. If you're not careful, you can convince yourself that you've proven many statements that are, in fact, false! Below is a favorite of mine.

Consider the following: let $P(n)$ be the statement that in any group of n people, everyone has the same name. We give a (false!) proof by induction that $P(n)$ is true for all n!

Proof: We proceed by induction.
Basis Step: Clearly, in any group with just 1 person, every person in the group has the same name.

Inductive Step: Assume $P(n)$ is true, namely, in any group of n people, everyone has the same name. We now prove $P(n+1)$. Consider a group of $n+1$ people:

$$\{1, 2, 3, \dots, n-1, n, n+1\}.$$

The first n people form a group of n people; by the inductive assumption, they all have the same name. So, the name of 1 is the same as the name of 2 is the same as the name of 3 ... is the same as the name of n.

Similarly, the last n people form a group of n people; by the inductive assumption they all have the same name. So, the name of 2 is the same as the name of 3 ... is the same as the name of n is the same as the name of $n+1$. Combining yields everyone has the same name! □

Where is the error? Even Borg drones have different designations; it's unlikely that everyone reading this book shares my name! Clearly we've done something terribly wrong, but where? Let's go through the above argument slowly and carefully. Rather than trying to follow the proof for an arbitrary n, let's run through it with specific values of n and see what happens.

If $n = 4$, we would have the set $\{1, 2, 3, 4, 5\}$, and the two sets of 4 people would be $\{1, 2, 3, 4\}$ and $\{2, 3, 4, 5\}$. We see that persons 2, 3 and 4 are in both sets, providing the necessary link. If $n = 3$ our set would be $\{1, 2, 3, 4\}$, and the two sets of 3 people would be $\{1, 2, 3\}$ and $\{2, 3, 4\}$. Again we find people in common, providing the necessary link.

What about smaller n? Eventually we reach $n = 1$. Then our set would be $\{1, 2\}$, and the two sets of 1 person would be $\{1\}$ and $\{2\}$; there is no overlap! The error was that we assumed n was "large" in our proof of $P(n) \Rightarrow P(n+1)$. Yes, in this problem, 2 is large. Terms like large and small are relative. The problem was we accidentally used some facts that only hold for $n \geq 2$. It's very easy to fall into this trap.

Problem A.2.8 *Similar to the above, give a false proof that any sum of integer squares is an integer square, i.e., $x_1^2 + \cdots + x_n^2 = x^2$. In particular, this would prove all positive integers are squares as $m = 1^2 + \cdots + 1^2$.*

A.3 Proof by Grouping

Our next technique is close to induction. I call it **proof by grouping**. A great example is the rule from calculus that the derivative of a sum is the sum of the derivatives. Most books prove this carefully for a sum of two functions, but then ignore the proof in general. Some care is needed; sadly, the derivative of an infinite sum need not equal the sum of the derivatives; however, if we have a finite sum of differentiable functions, then the derivative of the sum is the sum of the derivative.

We'll give the proof, assuming we know that whenever we have two differentiable functions then the derivative of their sum is the sum of their derivatives. What follows is essentially an induction argument, but I think it's nice to see how we win by cleverly adding parentheses and grouping terms.

Proof: Let

$$g(x) \quad = \quad f_1(x) + f_2(x) + f_3(x) \ = \ (f_1(x) + f_2(x)) + f_3(x)$$

be a sum of three differentiable functions; note we've grouped the first two functions together, and written g as a sum of *two* functions (the first is $f_1 + f_2$ and the second is f_3). Taking the derivative, we find

$$\begin{aligned} \frac{dg}{dx}(x) &= \frac{d}{dx}\left[f_1(x) + f_2(x) + f_3(x)\right] \\ &= \frac{d}{dx}\left[(f_1(x) + f_2(x)) + f_3(x)\right]. \end{aligned}$$

We now have the derivative of the sum of two functions, which we know is the sum of the two derivatives. We thus obtain

$$\frac{dg}{dx}(x) \ = \ \frac{d}{dx}\left(f_1(x) + f_2(x)\right) + \frac{df_3}{dx}(x).$$

We now use the derivative of the sum of two functions is the sum of the derivatives again. We thus obtain

$$\frac{dg}{dx}(x) \ = \ \frac{df_1}{dx}(x) + \frac{df_2}{dx}(x) + \frac{df_3}{dx}(x),$$

completing the proof. □

More generally, this type of argument shows the derivative of any finite sum is the sum of the derivatives, extending the common sum rule from calculus. Sadly, most calculus classes gloss over this point, and never remark that you need to be a bit careful as technically we only proved the derivative of a sum of two functions is the sum of the derivatives.

We'll see this method again in §6.2.2, where we meet the multinomial coefficients (a generalization of binomial coefficients), and in §14.3 (when we show sums of normal random variables are normal).

A.4 Proof by Exploiting Symmetries

There are infinitely many similar integrals that calculus professors love to give students. Here's on version: find

$$\int_{-2}^{2}(x^8 - 1701x^6 + 24601)\cos^3 x \sin(x^3 + 2x)\log(x^2 + 4)dx.$$

Good luck finding an anti-derivative for that! Class problems have a huge advantage over the real world: you know there has to be a solution using just the methods you know. Thus, this has to be doable using just Calculus I and II knowledge. The 'trick' is to notice that we are *not* being asked to find the anti-derivative. Yes, if we let $f(x)$ be the integrand and $F(x)$ an anti-derivative, then the answer is just $F(2) - F(-2)$. *If* we know an anti-derivative *then* we can evaluate the integral; however, maybe it's possible to evaluate the integral without finding F. It's *helpful* to know F, but it's not always essential. Sometimes all it does is help with the algebra (If you've done multivariable calculus, this is similar to Lagrange multipliers; often we can find the maximum / minimum values without finding the multipliers.).

The key observation is to note that it's not an arbitrary integral, but an integral from -2 to 2. Note that this is a *symmetric* region about 0. Further, the integrand is an odd function about 0. Recall that $f(x)$ is an **even function (about the point** a**)** if $f(a + x) = f(a - x)$, while it is an **odd function (about the point** a**)** if $f(a + x) = -f(a - x)$. The integral of an odd function about a over a symmetric interval centered at a is zero: this is because the contribution on one side is negated by the contribution on the other side. See Figure A.1 (left) for an example.

Another nice application of exploiting symmetries is to simplify integrations. The arguments above show that the integral of an odd function over a symmetric region is zero. What if we have an even function about a and we integrate over the interval $[a - b, a + b]$? In that case, the integral is double that of the integral over $[a, a + b]$, as the first half has the same contribution as the second half. See Figure A.1 (right) for an example.

We record our results.

Exploiting Symmetries: Integration of odd and even functions. Let $f(x)$ be an odd function about a, and $g(x)$ an even function about a. Then

$$\int_{a-b}^{a+b} f(x)dx \ = \ 0, \quad \int_{a-b}^{a+b} g(x)dx \ = \ 2\int_{a}^{a+b} g(x)dx.$$

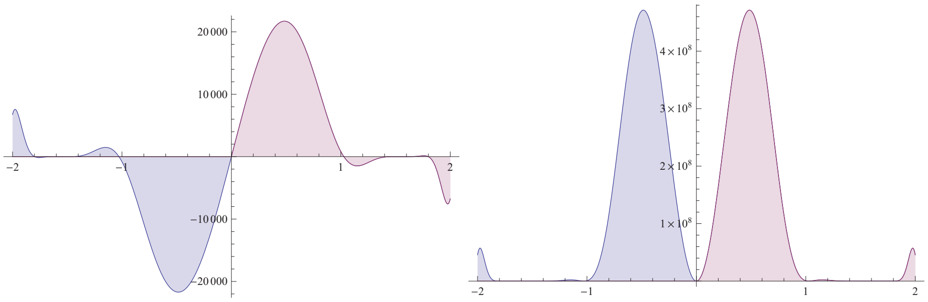

Figure A.1: Let $f(x) = (x^8 - 1701x^6 + 24601)\cos^3 x \sin(x^3 + 2x)\log(x^2 + 4)$. (Left) Area under $f(x)$ (an odd function about 0) from -2 to 2. (Right) Area under $f(x)^2$ (an even function about 0) from -2 to 2. Note the region is symmetric about 0.

These are two of the most common symmetries worth exploiting, but there are many others, and you should keep your eyes open for them. We'll do one more, which is useful when we prove the cosecant identity of the Gamma function. The following is a gem of mathematics:

$$\sum_{n=1}^{\infty} \frac{1}{n^2} \ = \ \frac{\pi^2}{6}.$$

There are lots of proofs of this, and it has probabilistic interpretations (for example, it's the reciprocal of the probability two random numbers are relatively prime). It's often proved in a Fourier Analysis or Complex Analysis course (see [SS1, SS2]).

Let's take this result as a given, and deduce the sum of the reciprocals of the odd squares. We find

$$\begin{aligned}
\sum_{n=1}^{\infty} \frac{1}{n^2} &= \sum_{\substack{n=1 \\ n \text{ even}}} \frac{1}{n^2} + \sum_{\substack{n=1 \\ n \text{ odd}}} \frac{1}{n^2} \\
&= \sum_{n=1}^{\infty} \frac{1}{(2n)^2} + \sum_{n=1}^{\infty} \frac{1}{(2n-1)^2} \\
&= \frac{1}{4}\sum_{n=1}^{\infty} \frac{1}{n^2} + \sum_{n=1}^{\infty} \frac{1}{(2n-1)^2} \\
\frac{3}{4}\sum_{n=1}^{\infty} \frac{1}{n^2} &= \sum_{n=1}^{\infty} \frac{1}{(2n-1)^2},
\end{aligned}$$

which means

$$\sum_{n=1}^{\infty} \frac{1}{(2n-1)^2} \ = \ \frac{3}{4}\sum_{n=1}^{\infty} \frac{1}{n^2} \ = \ \frac{3}{4}\frac{\pi^2}{6} \ = \ \frac{\pi^2}{8}.$$

This is another great example of the powerful consequences if you can **exploit symmetry** properly. The key observation is that the sum over the even terms is just one-fourth of the total sum. We then **brought it over** (we'll see more of this technique in the calculus review problems, especially Question F.2.40).

A.5 Proof by Brute Force

There are several ways to attack problems by brute force. They all share a common feature: rolling up your sleeves and diving into the algebra. Sometimes we're lucky and there are only a few items to check, but often there are so many cases that it just isn't feasible. Below we'll give an example to give a flavor of this method.

Recall the lemma on binomial coefficients.

Lemma A.5.1 *We have*

$$\binom{n}{k} = \binom{n}{n-k}, \quad \binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k}.$$

Proof of Lemma A.5.1: First Part: The first claim is just the fact that multiplication is commutative:

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} = \frac{n!}{(n-k)!k!} = \binom{n}{k}.$$

□

The second claim is more interesting. Here's a 'brute force' proof:

Proof of Lemma A.5.1: Second Part: We have

$$\begin{aligned}
\binom{n}{k} + \binom{n}{k-1} &= \frac{n!}{k!(n-k)!} + \frac{n!}{(k-1)!(n-k+1)!} \\
&= \frac{n!}{(k-1)!(n-k)!}\left[\frac{1}{k} + \frac{1}{n-k+1}\right] \\
&= \frac{n!}{(k-1)!(n-k)!}\left[\frac{n-k+1+k}{k(n-k+1)}\right] \\
&= \frac{n!}{(k-1)!(n-k)!}\frac{n+1}{k(n-k+1)} \\
&= \frac{(n+1)!}{k!(n-k+1)!} = \binom{n+1}{k}.
\end{aligned}$$

□

While the above argument *is* a proof, in some sense it's a terrible one. Yes, it's logically sound, yes, all the steps are correct, yes, it does give us the result; however, after reading it do you have any sense of *why* the result is true? It's just a long list of algebraic manipulations. It's great to be able to do this, but for many problems the algebra will be significantly worse, and it won't be clear at all how to proceed. For this problem, we were lucky. The algebra wasn't too bad, and it was pretty clear what to do: collect common factors and simplify. An alternative algebraic approach to this problem would have been to clear the denominators and then simplify. Is there another way to approach this problem, one which is more enlightening? Fortunately, the answer is a resounding yes, and we give it in the next section.

A.6 Proof by Comparison or Story

We return to the problem from the previous section, where we want to prove $\binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k}$. We've seen an unenlightening proof; now we'll see a better one that highlights what's really happening. The idea of this method, **Proof by Comparison**, is to compute the desired quantity two different ways. As we're calculating the same thing, these two expressions must be equal. Though the idea is easy to state, in practice it's often very hard to find a viewpoint that leads to an easy calculation. Combinatorial problems are some of the hardest you'll find (both in probability and in mathematics), and you often need a flash of insight (or a lot of experience) to suggest a good way to look at a problem.

Another way to say what we're going to do is that we'll **count the same quantity two different ways**. If we're counting the same quantity two different ways, then the two answers must agree; many identities are derived this way. We're essentially telling a story, with the exciting conclusion that the two main characters are actually one and the same. **Proof by Story** doesn't sound as academic as Proof by Comparison, but that's really what we're doing.

Let's do a simple warm-up example, inspired by the Dr. Seuss story *The Sneetches*. *Proof that* $\binom{n}{k} = \binom{n}{n-k}$*:* Imagine we have a group of n Sneetches, and we want to give some of them stars on their bellies. If exactly k are going to get stars, there are $\binom{n}{k}$ ways to choose k of the n Sneetches to be starred. Alternatively, we could look at this as *excluding* $n - k$ of the Sneetches from getting stars, and there are $\binom{n}{n-k}$ ways to choose $n - k$ Sneetches *not* to be starred. We've counted the same thing two different ways, so $\binom{n}{k} = \binom{n}{n-k}$.

I prefer this proof to the algebra one, as it illustrates what's really going on and *why* there's an equivalence.

Proof that $\binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k}$*:* We find a combinatorial interpretation for all these quantities. Imagine we have $n + 1$ marbles; n of these marbles are red and 1 marble is blue. This is the hardest part of the proof, figuring out what story to tell. While this gets easier with practice, there's at least a reason for doing this. We have $n + 1$ objects, so perhaps n of them are of one type, and 1 is of another.

One half of our story isn't too bad. There are $\binom{n+1}{k}$ ways to choose k marbles from the $n + 1$ marbles when we do not care about order of choice.

How else could we count the number of ways of choosing k marbles from $n+1$? Well, we could look at how many ways there are to choose k marbles from our $n+1$ marbles when order doesn't matter, keeping track of whether or not we choose the blue marble. If we don't have the blue marble, then we must choose k marbles from the n red ones; there are $\binom{n}{k}$ ways to do this. If we do have the blue marble (and there is $\binom{1}{1}$ way to do this as we only have one blue marble) then we must choose $k - 1$

red marbles from n red marbles, and there are $\binom{n}{k-1}$ ways of doing this. Collecting, we find our two counts must be equal, so

$$\binom{n}{k} + \binom{n}{k-1} \;=\; \binom{n+1}{k}.$$

□

This is a much better proof; it highlights what is going on, and gives a reason for the algebraic miracle.

There's a better way to view the calculation. We should *really* write it as

$$\binom{1}{0}\binom{n}{k} + \binom{1}{1}\binom{n}{k-1} \;=\; \binom{n+1}{k}.$$

Looking at it this way, the first factor on the left is the story: we choose 0 of the 1 blue marbles and then k of the n red marbles; the second factor represents choosing 1 of the 1 blue marbles and $k-1$ of the n red marbles. These two expressions are equal as $\binom{1}{0} = \binom{1}{1} = 1$, but I prefer the second. What's nice now is that a certain symmetry has been restored to both sides of the equation. On both the left and the right hand side, the sum of the 'top' parts of the terms add up to $n+1$, and the sum of the 'bottom' parts of the terms add up to k. To me, it's a little clearer how we're partitioning, and anything that can highlight what's going on is good! It also decreases the chance that we'll forget a factor, as in other problems these terms won't always be 1.

At the risk of beating the problem to death, it's worth chatting about why we're adding the two terms and not multiplying them. Often in probability we multiply the probabilities of events. Here, what we're doing is partitioning our event, which is choosing k of $n+1$, into disjoint possibilities (having 0 blue, having 1 blue). For finite sets, the probability of a disjoint union is the sum of the probabilities. This forces us to add the two probabilities together.

Let's do another example. We'll prove $\sum_{k=0}^{n}\binom{n}{k}\binom{n}{n-k} = \binom{2n}{n}$. We'll do this by calculating the same quantity two different ways. This essentially means we need to make up a story, where the expressions above are the quantities involved. Imagine we have n men and n women who want to take a probability class; unfortunately, the classroom is small and only n people can enroll in the class. There are $\binom{2n}{n}$ ways to choose a class of n people from our $2n$ people (n men and n women). That's the right hand side – what about the left hand side? Note in any class of n people there must be some number of men and some number of women. If there are k men, there must be $n-k$ women. The number of ways of choosing k men from n men is just $\binom{n}{k}$; similarly there are $\binom{n}{n-k}$ ways to choose $n-k$ women from n women. Thus, the number of ways to have a class of n people with exactly k men is $\binom{n}{k}\binom{n}{n-k}$. There must be some number of men in the class; that number ranges from 0 to n. Thus the total number of possible classes is just $\sum_{k=0}^{n}\binom{n}{k}\binom{n}{n-k}$, which must equal $\binom{2n}{n}$. □

As an aside, since $\binom{n}{n-k} = \binom{n}{k}$, the above implies $\sum_{k=0}^{n} \binom{n}{k}^2 = \binom{2n}{n}$.

For a nice challenge, try to find a simple formula involving a triple product of binomial coefficients. You'll have to think a bit and find a good story. *Hint: it isn't* $\binom{n}{k}^3$.

We end with one last example. Let's see how the Binomial Theorem can be proved in this manner. We want to show

$$(x+y)^n \;=\; \sum_{k=0}^{n} \binom{n}{k} x^{n-k} y^k.$$

We have n factors of $x + y$. For each factor, we choose either x or y. We see that $(x + y)^n$ will be a polynomial in x and y, involving terms like $x^j y^k$. What are the possible pairs of (i, j) that work, and what are the coefficients of these terms?

Well, we have n factors and for each factor we *choose* either an x or a y. Thus $j + k$ must equal n, so j must be $n - k$. What about the coefficient of $x^{n-k}y^k$? Every time we choose y from exactly k of the factors (which then forces us to have exactly $n - k$ factors of x), we get a $x^{n-k}y^k$. How many ways are there to choose k of the n factors to be y? Why, this is just the definition of the binomial coefficient, $\binom{n}{k}$, which completes the proof. □

It's worthwhile to see different proofs of the same result, especially if it's an important result. Each of these proofs highlights a different feature. These different approaches will help you not only in understanding the theorem, but in attacking future problems. What can make many math problems seem exceptionally difficult is that it isn't always clear how to start. The more methods you see, the more ideas you have for tackling future problems.

A.7 Proof by Contradiction

Proof by Contradiction is one of my favorite ways of proving statements. Sometimes, instead of trying to directly show that something is true, it's easier to assume it fails, and go for a contradiction. Let's look at an example. Remember that a number is **rational** if we can write it as a ratio of two integers (with the denominator nonzero); if we cannot do this, the number is **irrational**.

The square-root of 2 is irrational.

We proceed by assuming it is not irrational, and look for a contradiction; see [MilMo] for a more geometric proof by contradiction. If it isn't irrational then it's rational, and we have $\sqrt{2} = p/q$, and we may assume p and q are relatively prime (this means that no integer 2 or more divides both). If there were a common divisor, we could remove it and get a new fraction p'/q', with $p' < p$ and $q' < q$.

Since we're assuming $\sqrt{2} = p/q$, then $2q^2 = p^2$. We claim that 2 divides p^2. While this appears obvious, this must be proved. It's clearly true if p is even, as an even times any integer is still even. If p is odd, we may write $p = 2m + 1$. Then

$p^2 = 4m^2 + 4m + 1 = 2(2m^2 + 2m) + 1$, which is clearly not divisible by 2. Thus p is even, say $p = 2p_1$. Then $2q^2 = p^2$ becomes $2q^2 = 4p_1^2$. We now have $2p_1^2 = q^2$, and a similar argument yields q is even. Hence p and q have a common factor, which contradicts p and q are relatively prime. We were led to this falsehood by assuming $\sqrt{2}$ is rational. Thus that assumption must be false, and $\sqrt{2}$ must be irrational. □

As proofs by counterexample occur so frequently, it's worth doing another example. This one is more involved, and uses some results from analysis and calculus.

Let $f(x)$ be a continuous function on the real line. If the integral of $f(x)$ vanishes for every interval $[a, b]$ with $a < b$ then $f(x)$ is identically zero.

If we try to prove this directly we might run into some trouble, for we're given information on $f(x)$ over intervals, but must prove something over a point. What if, perhaps, we try to prove by contradiction? We assume for the sake of argument that the result is false: all the hypotheses hold, but there's a counterexample, say f, that is *not* zero everywhere. Now we have something to work with, and we try to show that if such a function existed, then it couldn't possibly satisfy all of our hypotheses. This contradiction means that our initial assumption that there was a counterexample *is false*, and thus the theorem does hold.

Let's try this here. So let's assume we have a continuous function which integrates to zero over any interval, but isn't identically zero. So there's some point, say x_0, where the function isn't zero. Without loss of generality, let's assume our function is positive at the point x_0 (a similar proof works for $f(x_0) < 0$).

Well, let's glean all the information we can out of our hypotheses on f. We assumed f is continuous. So, if we choose any $\epsilon > 0$ then we know there is a δ such that, if $|x - x_0| < \delta$ then $|f(x) - f(x_0)| < \epsilon$.

But, we have freedom in choosing ϵ! We know that our f must integrate to zero over any interval, so we have $\int_{x_0-\delta}^{x_0+\delta} f(x)dx = 0$. But we have $f(x_0)$ is positive! If ϵ is sufficiently small, by continuity $f(x)$ will be positive around x_0. For example: taking $\epsilon < f(x_0)/20$, we get there is a δ such that $f(x) > 19f(x_0)/20 > 0$.

Now we can get a contradiction. As $f(x) > 19f(x_0)/20$ on this interval but we've assumed the integral on this interval vanishes, standard results from calculus give us

$$0 \;=\; \int_{x_0-\delta}^{x_0+\delta} f(x)dx < \int_{x_0-\delta}^{x_0+\delta} \frac{19f(x)}{20}dx \;=\; \frac{19f(x_0)\cdot 2\delta}{20},$$

where the first equality follows from our assumption that f integrates to zero over any interval. But $f(x_0)\delta > 0$, and we've reached a contradiction! Basically the above is just a rigorous way of saying that if a continuous function is positive at some point, it's positive in a neighborhood of the point and thus cannot integrate to zero there. □

When you were reading this proof, you probably raised your eyebrows or wondered a bit when numbers like 19/20 entered the proof. Quantities like this are common in analysis. The idea is we want enough control to show that our function is positive in a small interval. We didn't need to take 19/20; many other numbers would've worked too.

A.8 Proof by Exhaustion (or Divide and Conquer)

The more assumptions and hypotheses we have on objects, the more (detailed) theorems and results we should know about them. Often it helps in proving theorems to break the proof up into several cases, covering all possibilities. We call this method **Divide and Conquer**. It's *essential* in using this method that you cover all cases: *make sure you consider all possibilities*. For example, you might do: *Case 1: the function is continuous*. Now you have all the theorems about continuity at your disposal. And then: *Case 2: the function is not continuous*. Now you have a special point where the function is discontinuous, and theorems and results about such points. The advantage is that, before, you couldn't use either set of results. The disadvantage is that you now have to give two proofs. Often, it's worthwhile having to prove more claims because for each claim you have more at your disposal. Let's do an example.

For f, g real valued functions, $|f(x) + g(x)| \leq |f(x)| + |g(x)|$.

If we can show this holds for an arbitrary point x, then we're done. Let's fix an x and investigate.

Case 1: Assume $f(x), g(x) \geq 0$.

Under this assumption, we have

$$|f(x) + g(x)| \;=\; f(x) + g(x) \;=\; |f(x)| + |g(x)|,$$

which is what we needed to show.

Case 2: Assume $f(x) \geq 0$, $g(x) < 0$.

We want to somehow get $f(x) + g(x)$. We can add them together, and get

$$f(x) + g(x) \;<\; 0 + f(x) \;=\; |f(x)|,$$

but when we take absolute values of both sides, the inequality could change ($-5 < 4$ but $|-5| > |4|$). So, a standard trick is to break this case into *subcases!*

- *Subcase A: Assume* $0 \leq f(x) + g(x)$. Then as $g(x) < 0$, $f(x) + g(x) < f(x)$. So $0 \leq |f(x) + g(x)| < f(x) \leq |f(x)| + |g(x)|$, which is what we needed to show.

- *Subcase B: Assume* $f(x) + g(x) < 0$. Then $0 < -1\,(f(x) + g(x)) \leq -g(x)$ as $f(x) \geq 0$. So $0 < |f(x) + g(x)| \leq |g(x)| \leq |f(x)| + |g(x)|$, which is what we needed to show.

This completes the analysis of Case 2. Unfortunately, we are not done as Cases 1 and 2 do not exhaust all possibilities.

Case 3: Assume $f(x) < 0$, $g(x) \geq 0$.

This is proved similarly as in Case 2; essentially we just switch the roles of f and g.

Case 4: Assume $f(x) < 0, g(x) < 0$. This is proved almost identically as in Case 1. □

Frequently in proofs by exhaustion many of the cases are essentially the same. For example, in the problem above it doesn't really matter if $f(x) \geq g(x)$ or $g(x) \geq f(x)$, as we can always change the label names of the functions. Because of this, you'll often see proofs using the phrase **without loss of generality**, which means that as it makes no difference in the proof, for definiteness we'll assume a certain ordering or certain values. Be careful, though, as sometimes the different names are important. For example, if we're studying the function $f(x, y) = x^2y^4 + x^4y^2$, then once we compute $\partial f/\partial x$ we know $\partial f/\partial y$ by interchanging the roles of x and y. This is not the case for the function $g(x, y) = x^2y^4 + x^3y^3$; here there is a real difference between the x-behavior and the y-behavior.

A.9 Proof by Counterexample

One of the most common mistakes students make is to assume that **Proof by Example** is a valid way to prove a relation. This isn't true; just because something sometimes works doesn't mean it will always work. We saw a great example in §A.2 when we looked at Euler's polynomial $x^2 + x + 41$; it was always prime for $n \in \{0, 1, \dots, 39\}$ but failed to be prime for many n afterwards.

While it's often useful to check a special case and build intuition on how to tackle the general case, checking a few examples isn't a proof. For another example, because $16/64 = 1/4$ and $19/95 = 1/5$, one might think that in dividing two digit numbers if two numbers on a diagonal are the same one just cancels them. Skeptical? Let's test it again. If we look at $49/98$, canceling the 9's gives $4/8$, which simplifies to 1/2. Convinced? Probably not. A little experimentation brings us to $12/24$. If we really could just cancel the 2's we'd get this equals 1/4, but it's 1/2. Of course this is *not* how one divides two digit numbers, but it is interesting to see how many times it works!

However, if we are trying to disprove some statement, this means that if we are able to find just one example where the statement fails under the necessary assumptions of the statement, then we have in fact disproved it, as we have shown that it does not hold for all cases. This is the essence of **Proof by Counterexample**.

Problem A.9.1 *How many pairs of three digit numbers with the same middle are there such that the ratio of these two numbers is the same as the ratio with the middle digit removed? For example, one pair is* $(561, 462)$, *as* $561/462 = 51/42 = 17/14$.

A.10 Proof by Generalizing Example

Another great way to prove a result is to look at a special case, detect a pattern, and try to generalize what you see. Let's look at an example you may have seen years ago when learning how to multiply and divide. You may remember the rule for divisibility by 3: if the sum of the digits of your number is divisible by three, then so is your number. We check this with 231 (yes, as 2+3+1 = 6 which is divisible by 3, as is $231 = 3 \cdot 77$), 9444 (yes, as 9+4+4+4 = 21 which is divisible by 3, as is $9444 = 3 \cdot 3148$), and 1717 (no, as $1 + 7 + 1 + 7 = 16$ which is not divisible by 3, nor is 1717). Now, while the rule is true, checking a few examples doesn't constitute a proof. We haven't checked *every* number, only three specific numbers. We would have to show that, given an arbitrary number with digits $a_n \dots a_3a_2a_1a_0$, then if $a_0 + a_1 + \cdots + a_n$ is divisible by 3, so is $a_n \dots a_3a_2a_1a_0$.

This leads us to proving claims by generalizing an example or known case. Often the way the theorem is stated, it tries to guide you as to what to do. For instance, in the theorem we're trying to prove on divisibility by three, it tells us that divisibility by three is related to the sum of the digits of our number. So, we ask ourselves: how can we get the sum of the digits, given the number $a_n \dots a_3a_2a_1a_0$?

For example, 314 would be $a_2a_1a_0$, with $a_2 = 3$, $a_1 = 1$, $a_0 = 4$, and the sum of digits would be 3+1+4. Well, we might try looking at other ways of writing our number. Often there are different forms that are equivalent, but bring out different properties. For digits, we recall this comes from powers of 10: our number 314 can be written as $314 = 3 \cdot 100 + 1 \cdot 10 + 4 \cdot 1$.

So, notice what happens if we subtract from 314 the sum of its digits:

$$\begin{aligned}
314 - (3 + 1 + 4) &= 3 \cdot 100 + 1 \cdot 10 + 4 \cdot 1 - (3 + 1 + 4) \\
314 - (3 + 1 + 4) &= (3 \cdot 100 - 3) + (1 \cdot 10 - 1) + (4 \cdot 1 - 4) \\
314 - (3 + 1 + 4) &= (3) \cdot 99 + (1) \cdot 9 + (4) \cdot 0.
\end{aligned}$$

Ah. Notice that the right hand side is clearly divisibly by 3, as each term is multiplied by 0 or 9 or 99. If $3 + 1 + 4$ is divisible by 3, when we bring it over to the right hand side we find 314 equals number divisible by three! If $3 + 1 + 4$ is not divisible by three, when we bring it over we get 314 equals a number *not* divisible by three!

Now we've done this proof in the special case when our number is 314. There's nothing wrong with first proving something for a specific case or number of function, as long as we then generalize. We see that the exact same proof would carry through if instead we considered the number: $a_n \dots a_3a_2a_1a_0 = a_n \cdot 10^n + \cdots + a_1 \cdot 10^1 + a_0 \cdot 10^0$.

A.11 Dirichlet's Pigeon-Hole Principle

The following seemingly trivial observation appears in a variety of problems, and is a very powerful way to prove many claims.

Dirichlet's Pigeon-Hole Principle Let $A_1, A_2, \dots, A_n$ be a collection of sets with the property that $A_1 \cup \cdots \cup A_n$ has at least $n + 1$ elements. Then at least one of the sets A_i has at least two elements.

This is called the Pigeon-Hole Principle for the following reason. Imagine we have $n + 1$ pigeons and n boxes, and we put each pigeon in exactly one box. Then at least one box must have two pigeons. If not, then each box has at most 1 pigeon, and as there are n boxes, this can account for at most n pigeons – at least one pigeon is missing! In a more mathematical prose, if we distribute k objects in n boxes and $k > n$, one of the boxes contains at least two objects. The Pigeon-Hole Principle is also known as the **Box Principle**. While there are many applications in number theory, there are a few in probability as well. For example, it's used in the Birthday Problem in Chapter 1 to see that once we have 366 people then we must have at least two sharing a birthday (we assumed no one was born on February 29th). Let's do one more example. We'll first give the slick proof, then talk a bit about how to find such arguments.

Let S be any subset of $\{1, 2, \dots, 2n\}$ with $n+1$ elements. Then S contains at least two elements a, b with a dividing b.

To see this, we write each element $s \in S$ as $s = 2^\sigma s_0$ with s_0 odd. There are n odd numbers in the set $\{1, 2, \dots, 2n\}$, and as the set S has $n+1$ elements, the Pigeon-Hole Principle implies that there are at least two elements a, b with the same odd part. Without loss of generality, we might as well assume $a < b$, and write the numbers as $a = 2^i(2m+1)$ and $b = 2^j(2m+1)$. As $a < b$, $i < j$, we see $b = 2^{j-i}a$, proving a does indeed divide b.

The hard part of this problem is figuring out *how* to use the Pigeon-Hole Principle. The phrasing gives us some clues that we *should* use it. We have a collection of objects and we want to show that if we take a large enough subset, then at least two of those have a special relation. The Pigeon-Hole Principle is all about forced relations when we have enough item, so this is a natural approach.

The trick or difficulty is realizing that we should write our numbers as a power of two times an odd number, and then there must be two odd components that are equal. How can we figure out that *this* is what we should try? One way is to take special values of n and look at some sets, and see which elements have things in common. Related to this, try to take sets with just n numbers and see whether or not you can make the claim fail (since we're not taking $n+1$ objects, it's fine to have the conclusion fail). After some experimentation, you might hit upon looking at the n odd numbers less than $2n$. If $2n = 8$ then this is a good choice, as $\{1, 3, 5, 7\}$ is such that no number divides another in this list; however, if $2n = 10$ we'd have $\{1, 3, 5, 7, 9\}$, and 3 divides 9. This illustrates the dangers of looking at small cases; we might see something that doesn't persist.

Returning to the drawing board, what other good sets are there of $\{1, 2, \dots, 2n\}$ with n items? Perhaps a good choice is $\{n+1, n+2, \dots, 2n\}$. This set always has n elements, and for all n we never have one element in the list dividing another. This is a great example, and we now know that we can't replace the $n+1$ in the theorem with n. If the theorem is true, any element x added to $\{n+1, n+2, \dots, 2n\}$ gives two numbers where one divides another. Further, as $x \le n$, it must be the case that x divides something already in our list. It's not immediately clear, though, what it should divide. If x is large, say $n/2 < x \le n$, then $2x$ is in our list $\{n+1, n+2, \dots, 2n\}$. If $n/4 < x \le n/2$, then $4x$ is in our list. This is probably the hardest jump to make, seeing the powers of two come into play. What we're trying to do is gather data and use that to guide us. From here, we somehow have to make the leap to noticing that our special pair differ by a power of 2.

A.12 Proof by Adding Zero or Multiplying by One

I've saved my personal favorite for last: **adding zero** and **multiplying by one**. At first glance, neither of these seem capable of being that useful. After all, if we multiply by one, we're back where we started. The same goes for adding zero. Neither of these operations changes our expression.

Exactly! These are powerful methods *because* they don't change anything. We can't modify one side of an equality and not the other. We can't discriminate mathematically: whatever we do to one side, we must do to the other. The reason these are useful methods is that we can write 1 or 0 in many different ways, and we don't have to use the same representation on both sides. The point of this is to arrange the algebra in a more illuminating manner, to extract out sub-expressions that we know. Let's do a few examples. I chose these examples from calculus, but all we really need is the definition of the derivative, which states

$$f'(x) \;=\; \lim_{h\to 0}\frac{f(x+h)-f(x)}{h} \;=\; \lim_{x'\to x}\frac{f(x')-f(x)}{x'-x}$$

(both variants are used below), and that the derivative of x^n is nx^{n-1}. For an example in probability, go to §2.5.2.

Our first example is the proof of the product rule in calculus. Imagine f and g are differentiable functions, and set $A(x) = f(x)g(x)$. It's not unreasonable to hope that there's a nice formula for the derivative of A in terms of f, f', g and g'. A great way to guess this relationship is to take some special examples. If we try

$$f(x) \;=\; x^3 \quad \text{and} \quad g(x) \;=\; x^4,$$

then

$$A(x) \;=\; x^7 \quad \text{so} \quad A'(x) \;=\; 7x^6.$$

At the same time,

$$f'(x) \;=\; 3x^2 \quad \text{and} \quad g'(x) \;=\; 4x^3.$$

There's only two ways to combine $f(x)$, $f'(x)$, $g(x)$ and $g'(x)$ and get x^6: $f'(x)g(x)$ and $f(x)g'(x)$. (Okay, there are more ways if we allow divisions; there's only two ways if we restrict ourselves to addition and multiplication.) Interestingly, if we add these together we get $3x^2 \cdot x^4 + x^3 \cdot 4x^3 = 7x^6$, which is just $A'(x)$. This *suggests* that $A'(x) = f'(x)g(x) + f(x)g'(x)$. If we try more and more examples, we'll see this formula keeps working. While this is strong evidence, it's not a proof; however, it *will* suggest the key step in our proof.

From the definition of the derivative and substitution,

$$A'(x) \;=\; \lim_{h\to 0}\frac{A(x+h)-A(x)}{h} \;=\; \lim_{h\to 0}\frac{f(x+h)g(x+h)-f(x)g(x)}{h}.$$

From our investigations above, we think the answer should be $f'(x)g(x)+f(x)g'(x)$. We can begin to see an $f'(x)$ and a $g'(x)$ lurking above. Imagine the last term were $f(x)g(x+h)$ instead of $f(x)g(x)$. If this were the case, the limit would equal $f'(x)g(x)$ (we pull out the $g(x+h)$, which tends to $g(x)$, and what's left is the definition of $f'(x)$). Similarly, if the first piece were instead $f(x)g(x+h)$, then we'd get $f(x)g'(x)$. What we see is that our expression is *trying* to look like the right things, but we're missing pieces. This can be remedied by adding zero, in the form $f(x)g(x+h) - f(x)g(x+h)$. Let's see what this does. In the algebra below we use the limit of a sum is the sum of the limits and the limit of a product is the product of the limits; we can use these results as all these limits exist. We find

$$\begin{aligned}
A'(x) &= \lim_{h\to 0} \frac{f(x+h)g(x+h) - f(x)g(x+h) \ + \ f(x)g(x+h) - f(x)g(x)}{h} \\
&= \lim_{h\to 0} \frac{f(x+h)-f(x)}{h} g(x+h) + \lim_{h\to 0} f(x) \frac{g(x+h)-g(x)}{h} \\
&= \lim_{h\to 0} \frac{f(x+h)-f(x)}{h} \lim_{h\to 0} g(x+h) + \lim_{h\to 0} f(x) \lim_{h\to 0} \frac{g(x+h)-g(x)}{h} \\
&= f'(x)g(x) + f(x)g'(x).
\end{aligned}$$

□

The above proof has a lot of nice features. First off, it's the proof of a result you should know (at least if you've taken a calculus class). Second, we were able to guess the form of the answer by exploring some special cases. Finally, the proof was a natural outgrowth of these cases. We saw terms like $f'(x)g(x)$ and $f(x)g'(x)$ appearing, and thus asked ourselves: *So, what can we do to bring out these terms from what we have?* This led to adding zero in a clever way. It's fine to add zero, as it doesn't change the value. The advantage is we ended up with a new expression where we could now do some great simplifications.

For our second example, we'll look at the chain rule, one of the most dreaded rules from calculus. Now we take $B(x) = f(g(x))$. We assume f and g are differentiable, that $f(g(x))$ is defined, and for convenience we assume $g'(x)$ is continuous and never zero. This assumption isn't needed, but it'll simplify the argument so we make it as our point here is not to prove your old calculus results but rather to highlight the power of multiplying by 1. Since it worked so well last time, let's try to build some intuition from looking at

$$f(x) \ = \ x^3 \quad \text{and} \quad g(x) \ = \ x^4.$$

Again we have

$$f'(x) \ = \ 3x^2 \quad \text{and} \quad g'(x) \ = \ 4x^3;$$

however, we need to remember that we're supposed to evaluate f and f' not at x but at $g(x)$, so the relevant quantities are

$$\begin{aligned}
B(x) &= f(g(x)) \ = \ (x^4)^3 \ = \ x^{12} \\
f'(g(x)) &= 3(x^4)^2 = 3x^8 \\
g'(x) &= 4x^3.
\end{aligned}$$

Since $B'(x) = 12x^{11}$, looking at out building blocks we see that $12x^{11} = 3x^8 \cdot 4x^3$, or in this case we have $B'(x) = f'(g(x)) \cdot g'(x)$. So, just like the product rule, we have a candidate for the derivative. Knowing our goal is a great aid in suggesting the right way to manipulate expressions.

From the definition of the derivative, we have

$$B'(x) \;=\; \lim_{h\to 0} \frac{B(x+h)-B(x)}{h} \;=\; \lim_{h\to 0} \frac{f(g(x+h))-f(g(x))}{h}.$$

We're searching for $f'(g(x))$ and $g'(x)$. Note the numerator almost looks like the derivative of f at the point $g(x)$; the reason it isn't is that we evaluate f at $g(x+h)$ rather than $g(x)+h$. What if we use the second variant for the definition of the derivative? In that case, $x' = g(x+h)$ tends to x, but the denominator isn't right. It should be $x' - g(x) = g(x+h) - g(x)$, but it's only h. To remedy this, we multiply by 1 in the form of $\frac{g(x+h)-g(x)}{g(x+h)-g(x)}$, and find

$$\begin{aligned} B'(x) &= \lim_{h\to 0} \frac{f(g(x+h))-f(g(x))}{h} \frac{g(x+h)-g(x)}{g(x+h)-g(x)} \\ &= \lim_{h\to 0} \frac{f(g(x+h))-f(g(x))}{g(x+h)-g(x)} \frac{g(x+h)-g(x)}{h} \\ &= \lim_{h\to 0} \frac{f(g(x+h))-f(g(x))}{g(x+h)-g(x)} \lim_{h\to 0} \frac{g(x+h)-g(x)}{h} \\ &= f'(g(x)) \cdot g'(x). \end{aligned}$$

The last few lines deserve some justification. We're using the second variant of the definition of the derivative. Since the derivative of f exists, $\lim_{x'\to x} \frac{f(x')-f(g(x))}{x'-g(x)}$ equals $f'(g(x))$ for *any* sequence of x' tending to $g(x)$, and thus for the particular sequence where $x' = g(x+h)$.

Where did we use our assumption that $g'(x)$ is continuous and never zero? That assumption implies $g(x) = g(y)$ if and only if $x = y$. It's essential that $g(y) \neq g(x)$ for x and y distinct as otherwise $\frac{g(x+h)-g(x)}{g(x+h)-g(x)}$ could be 0/0.

We end with one last remark on these techniques. It's kind of like drawing auxiliary lines in geometry or trigonometry to highlight relationships. Drawing these lines doesn't change anything, but it often draws our attention to certain aspects of the problem.

Appendix B

Analysis Results

Not surprisingly, the tools from calculus (and more generally, real analysis) play a big role in probability. The reason, of course, is that to each random variable we attach a probability distribution. Often that distribution is continuous and even differentiable, and the quantities we want to study can be expressed in terms of our density and its integrals and derivatives.

We quickly review some of the key results from analysis below, and give some idea of how these are used.

B.1 The Intermediate and Mean Value Theorems

This section involves two of the biggest theorems from calculus, the Intermediate and the Mean Value Theorem. We'll use the Intermediate Value Theorem to prove the Mean Value Theorem, which can then be used to approximate numerous probabilities. First, we quickly review some notation. We write (a, b) for the interval $\{x : a < x < b\}$, and call this an **open interval**; by $[a, b]$ we mean $\{x : a \leq x \leq b\}$, and we call this a **closed interval**. We could of course have a half-open interval $[a, b)$ (which is also a half-closed interval!).

Theorem B.1.1 (Intermediate Value Theorem (IVT)) *Let f be a continuous function on $[a, b]$. For all C between $f(a)$ and $f(b)$ there exists $c \in [a, b]$ such that $f(c) = C$. In other words, all intermediate values of a continuous function are obtained.*

If we convert from mathspeak to English, the theorem is a lot clearer, and quite reasonable. One way to do this is with the following example. Imagine we're driving our car. We start off traveling at 20 mph (about 32 kph), and later in the trip we're cruising at 100 mph (about 161 kph). As this example is for math, the police will kindly look the other way this one time. The Intermediate Value Theorem asserts that, at some time in our trip, we must've been traveling 50 mph (about 80 kph). This should be reasonable; we're assuming our speed is given by a nice, continuous function, and thus we can't get from the slow starting speed to the fast final speed without passing through all *intermediate* speeds.

Sketch of the proof: We proceed by **Divide and Conquer**. Without loss of generality, we can assume $f(a) < C < f(b)$, as the proof is trivial if $f(a) = C$ or $f(b) = C$. Many proofs start like this – first get rid of the straightforward cases, and then move on to the heart of the argument.

Let x_1 be the midpoint of $[a, b]$. If $f(x_1) = C$ we're done. If not, there are two cases: either $f(x_1) < C$ or $f(x_1) > C$. If $f(x_1) < C$, we look at the interval $[x_1, b]$. If $f(x_1) > C$ we look at the interval $[a, x_1]$.

In either case, we have a new interval, call it $[a_1, b_1]$, such that $f(a_1) < C < f(b_1)$ and the interval has half the size of $[a, b]$. We continue in this manner, repeatedly taking the midpoint and looking at the appropriate half-interval.

For example, imagine that our function is $f(x) = x^2 + x + 1$, $a = 0$, $b = 1$ and $C = 2$. We have $f(0) = 1$ and $f(1) = 3$. We look at the midpoint and find $f(1/2) = 1.75$, thus our next interval is $[a_1, b_1] = [1/2, 1]$. We continue; we have $f(1/2) = 1.75$, $f(1) = 3$ and at the midpoint $3/4$ we find $f(3/4) = 37/16 = 2.3125$. This means that our next interval is $[a_2, b_2] = [1/2, 3/4]$.

To recap, we have a sequence of intervals

$$[a, b] \ \supset \ [a_1, b_1] \ \supset \ [a_2, b_2] \ \supset \ \cdots$$

such that $f(a_n) \le C \le f(b_n)$, and each a_n and b_n is either an endpoint from the previous interval, or the midpoint of the previous interval. If any of these satisfy $f(x_n) = C$, we're done. If no midpoint works, we divide infinitely often and obtain a sequence of points x_n in intervals $[a_n, b_n]$. This is where rigorous mathematical analysis is required (see for example [Rud] for details). In a real analysis class you'll show that

$$\bigcap_{n=1}^{\infty} [a_n, b_n] \ = \ [a_1, b_1] \cap [a_2, b_2] \cap [a_3, b_3] \cap \cdots$$

is just a point, say $\{x_0\}$. This is intuitively plausible; at each stage we have an open interval, and we cut its length in half when we go to the next level. Thus the final result cannot have any positive length. It should be nonempty as we have the chain

$$a_1 \ \le \ a_2 \ \le \ a_3 \ \le \ \cdots \ \le b_3 \ \le \ b_2 \ \le \ b_1.$$

Let's assume there's a unique point in the intersection. Since f is continuous and $a_n \to x_0$ and $b_n \to x_0$,

$$\lim_{n\to\infty} f(a_n) \ = \ f(x_0) \ = \ \lim_{n\to\infty} f(b_n);$$

this is just a restatement of what it means for f to be continuous at x_0. But

$$f(a_n) \ \le \ C \ \le f(b_n) \quad \text{and} \quad f(a_n) \ \le \ f(x_0) \ \le f(b_n).$$

This implies that $f(x_0) = C$. Why? They are both 'squeezed' to the same thing. Specifically, as $\lim_{n\to\infty} f(a_n) = \lim_{n\to\infty} f(b_n)$, we see that both of these limits equal C as well as $f(x_0)$. Thus, we have found our point! (For the example $f(x) = x^2 + x + 1$ on $[0, 1]$ with $C = 2$, we would find $x_0 = \frac{\sqrt{5}-1}{2} \approx 0.618034$.) □

Theorem B.1.2 (The Mean Value Theorem (MVT)) *Let $f(x)$ be differentiable on $[a, b]$. Then there exists $c \in (a, b)$ such that*

$$f(b) - f(a) \;=\; f'(c) \cdot (b - a).$$

Let's give an interpretation of the Mean Value Theorem. Let $f(x)$ represent the distance our car has traveled from the starting point at time x. The average speed from a to b is the distance traveled, $f(b) - f(a)$, divided by the elapsed time, $b - a$. As $f'(x)$ represents the speed at time x, the Mean Value Theorem says that there's some intermediate time at which we're traveling at the average speed.

For example, imagine that our average speed is 50 mph (about 80 kph). If our speed is always below 50 mph, there's no way that our average speed could be 50 mph; similarly if our speed is always above 50 mph there's no way our average speed could be 50 mph. Thus either our speed is always 50 mph (in which case the conclusion is trivial), or we can deduce that at some point in time we were traveling slower than 50 mph and at another point in time we were traveling faster. We can now use the Intermediate Value Theorem to prove that at some point we must be traveling at 50 mph, as that is an *intermediate* speed. This is essentially the proof; the only difference is that usually in a math book one sees impressive looking math symbols rather than text about cars!

To prove the Mean Value Theorem in familiar math language, it suffices to consider the special case when $f(a) = f(b) = 0$; this case is known as Rolle's Theorem.

Theorem B.1.3 (Rolle's Theorem) *Let f be differentiable on $[a, b]$, and assume $f(a) = f(b) = 0$. Then there exists $c \in (a, b)$ such that $f'(c) = 0$.*

Show the Mean Value Theorem follows from Rolle's Theorem. *Hint:* Consider

$$h(x) \;=\; f(x) - \frac{f(b) - f(a)}{b - a}(x - a) - f(a).$$

Note $h(a) = f(a) - f(a) = 0$ and $h(b) = f(b) - (f(b) - f(a)) - f(a) = 0$. The conditions of Rolle's Theorem are satisfied for $h(x)$, and

$$h'(c) \quad = \quad f'(c) - \frac{f(b) - f(a)}{b - a}.$$

Proof of Rolle's Theorem: Step one is to handle some special cases. We'll assume that $f'(a)$ and $f'(b)$ are non-zero. If one of these is zero we sadly aren't quite done, as the theorem asserts there is a *c strictly between a* and b; however, as a similar proof to what we give below handles this case, we leave that case as an exercise to the reader.

Multiplying $f(x)$ by -1 if needed, we may assume $f'(a) > 0$. *For convenience, we assume $f'(x)$ is continuous.* This assumption simplifies the proof, but isn't necessary. As you read the proof below, try to see where we use f' is continuous.

Case 1: $f'(b) < 0$**:** As $f'(a) > 0$ and $f'(b) < 0$, the Intermediate Value Theorem applied to $f'(x)$ asserts that all intermediate values are attained. As $f'(b) < 0 < f'(a)$, this implies the existence of a $c \in (a, b)$ such that $f'(c) = 0$.

Case 2: $f'(b) > 0$**:** $f(a) = f(b) = 0$, and the function f is increasing at a and b. If x is real close to a then $f(x) > 0$ if $x > a$. This follows from the fact that

$$f'(a) \quad = \quad \lim_{x \to a} \frac{f(x) - f(a)}{x - a}.$$

As $f'(a) > 0$, the limit is positive. As the denominator is positive for $x > a$, the numerator must be positive. Thus $f(x)$ must be greater than $f(a)$ for such x. Similarly $f'(b) > 0$ implies $f(x) < f(b) = 0$ for x slightly less than b.

Therefore the function $f(x)$ is positive for x slightly greater than a and negative for x slightly less than b. If the first derivative were always positive then $f(x)$ could never be negative as it starts at 0 at a. This can be seen by again using the limit definition of the first derivative to show that if $f'(x) > 0$ then the function is increasing near x. Thus the first derivative cannot always be positive. Either there must be some point $y \in (a, b)$ such that $f'(y) = 0$ (and we're then done) or $f'(y) < 0$. By the Intermediate Value Theorem, as 0 is between $f'(a)$ (which is positive) and $f'(y)$ (which is negative), there's some $c \in (a, y) \subset [a, b]$ such that $f'(c) = 0$. □

Did you see where we used f' was continuous? It happened when we invoked the Intermediate Value Theorem. Whenever you use a theorem, you need to make sure all the conditions are satisfied. To use the IVT, we need our function to be continuous.

B.2 Interchanging Limits, Derivatives and Integrals

B.2.1 Interchanging Orders: Theorems

For the convenience of the reader we record exact statements of several standard results from advanced calculus that are used at various points of the text. As the Change of Variable Theorem is so important, it gets its own chapter (Appendix C).

Theorem B.2.1 (Fubini's Theorem) *Assume f is continuous and*

$$\int_a^b \int_c^d |f(x,y)| dx dy \ < \ \infty.$$

Then

$$\int_a^b \left[\int_c^d f(x,y) dy \right] dx \ = \ \int_c^d \left[\int_a^b f(x,y) dx \right] dy.$$

Similar statements hold if we instead have

$$\sum_{n=N_0}^{N_1} \int_c^d f(x_n, y) dy, \qquad \sum_{n=N_0}^{N_1} \sum_{m=M_0}^{M_1} f(x_n, y_m).$$

For a proof in special cases, see [BL, VG]; an advanced, complete proof is given in [Fol]. See Exercise B.7.2 for an example where the orders of integration cannot be changed.

Theorem B.2.2 (Interchanging Differentiation and Integration) *Let $f(x,t)$ be a continuous function whose partial derivatives with respect to x and with respect to t are continuous in the region $\{(x,t) : x \in [a,b], t \in [c,d]\}$ with a, b, c, d finite. Then*

$$\frac{d}{dx}\int_a^b f(x,t)dt \ = \ \int_a^b \frac{\partial f}{\partial x}(x,t)dt.$$

The above theorem holds in greater generality. We can allow the regions to be infinite, at the cost of requiring additional decay in the functions. For a proof and generalizations, see [La2].

Our last result is on interchanging limits and integrals. We state one of the most useful below, though *not* in its most general form (see [Fol] for the more general phrasing and a proof).

Theorem B.2.3 (Dominated Convergence Theorem) *Let $\{f_n\}$ be a sequence of piecewise continuous real-valued functions on $\mathbb{R}$, and assume there is a non-negative, piecewise continuous function g with $|f_n(x)| \le g(x)$ for all n. Assume $\lim_{n\to\infty} f_n(x)$ converges pointwise to a piecewise continuous function f. Then*

$$\lim_{n\to\infty}\int_{-\infty}^{\infty} f_n(x)dx \ = \ \int_{-\infty}^{\infty}\lim_{n\to\infty} f_n(x)dx;$$

in other words, we may interchange the limit and the integral.

B.2.2 Interchanging Orders: Examples

The purpose of this section is to give a quick crash course in using analysis to justify certain statements. What follows is essentially independent of the rest of the book. As it's important to know *how* to justify statements (this lessens the chance of accidentally using results that can't be justified!), it's fine to skim or skip what follows.

In general, we need to appeal to some advanced theorems in analysis to interchange the order of operations, such as switching the order of integration or interchanging a sum and a derivative. In the case of the geometric series, however, we can justify interchanging the sum and the derivative without appealing to advanced machinery. The reason is that if we truncate the geometric series

$$\sum_{n=0}^{\infty} \ = \ 1 \ + \ x \ + \ x^2 \ + \ x^3 \ + \ x^4 \ + \ \cdots \ = \ \frac{1}{(1-x)^2} \tag{B.1}$$

at any N, the geometric series formula gives us an explicit formula for the sum of the tail:

$$\begin{aligned}\sum_{n=0}^{\infty} \ &= \ \left(1 \ + \ x \ + \ x^2 \ + \ \cdots \ + \ x^N\right) \ + \ \left(x^{N+1} \ + \ x^{N+2} \ + \ \cdots\right) \\ &= \ 1 \ + \ x \ + \ x^2 \ + \ \cdots \ + \ x^N \ + \ \frac{x^{N+1}}{1-x} \ = \ \frac{1}{(1-x)^2}.\end{aligned}$$

We show that we may interchange differentiation and summation for the geometric series (assuming, of course, that $|x| < 1$). The derivative of the right hand side (with respect to x) of (B.1) is just $(1-x)^{-2}$. We want to say the derivative of the left hand side of (B.1) is

$$\sum_{n=0}^{\infty} n x^{n-1},$$

but do to so requires us to justify

$$\frac{d}{dx}\sum_{n=0}^{\infty} x^n \ = \ \sum_{n=0}^{\infty} \frac{d}{dx} x^n.$$

A standard way to justify statements like this is as follows. We note that $\sum_{n=0}^{\infty} n x^{n-1}$ converges for $|x| < 1$; if we can show that for any $\epsilon > 0$ that this is within ϵ of $(1-x)^{-2}$, then we will have justified the interchange.

To see this, fix an $\epsilon > 0$. For each N, as discussed above we may write

$$\begin{aligned}\sum_{n=0}^{\infty} x^n &= \sum_{n=0}^{N} x^n + \sum_{n=N+1}^{\infty} x^n \\ &= \sum_{n=0}^{N} x^n + \frac{x^{N+1}}{1-x} \ = \ \frac{1}{1-x}.\end{aligned}$$

We can differentiate each side, and we can justify interchanging the differentiation and the summation because we have *finitely many* sums. Specifically, there are only $N+2$ terms ($N+1$ from the sum and then one more, $\frac{x^{N+1}}{1-x}$). Therefore we have

$$\begin{aligned}\frac{d}{dx}\sum_{n=0}^{N} x^n + \frac{d}{dx}\frac{x^{N+1}}{1-x} &= \frac{d}{dx}\frac{1}{1-x} \\ \sum_{n=0}^{N} n x^{n-1} + \frac{(N+1)x^N(1-x) - x^{N+1}(-1)}{(1-x)^2} &= \frac{1}{(1-x)^2} \\ \sum_{n=0}^{N} n x^{n-1} + \frac{(N+1)(1-x)+x}{(1-x)^2}\, x^N &= \frac{1}{(1-x)^2}.\end{aligned}$$

As $|x| < 1$, given any $\epsilon > 0$ we can find an N_0 such that for all $N \geq N_0$,

$$\left|\frac{(N+1)(1-x)+x}{(1-x)^2}\, x^N\right| \ \leq \ \frac{\epsilon}{2}.$$

Similarly we can find an N_1 such that for all $N \geq N_1$ we have

$$\left|\sum_{n=N+1}^{\infty} n x^{n-1}\right| \ \leq \ \frac{\epsilon}{2}.$$

Therefore we have shown that for every $\epsilon > 0$ we have

$$\left|\frac{1}{(1-x)^2} - \sum_{n=0}^{\infty} n x^{n-1}\right| \ \leq \ \epsilon,$$

proving the claim. Instead of studying these sums for a specific x, we can consider $x \in [a, b]$ with $-1 < a \le b < 1$, and N_0, N_1 will just depend on a, b and ϵ.

One situation where we cannot interchange differentiation and summation is when we have series that are **conditionally convergent** but not absolutely convergent. This means $\sum a_n$ converges but $\sum |a_n|$ does not. For example, consider

$$\sum_{n=0}^{\infty} \frac{x^n}{n}. \tag{B.2}$$

If $x = -1$ this series conditionally converges but not absolutely; in fact, as

$$-\log(1-x) \;=\; x + \frac{x^2}{2} + \frac{x^3}{3} + \cdots \;=\; \sum_{n=1}^{\infty} \frac{x^n}{n},$$

then (B.2) with $x = -1$ is just $-\log 2$. What happens if we try to differentiate? We have

$$\frac{d}{dx}\left[-\log(1-x)\right] \;=\; \frac{d}{dx}\left[\sum_{n=1}^{\infty} \frac{x^n}{n}\right].$$

The left hand side is easy to differentiate for $x \in [-1, 0]$, giving $\frac{1}{1-x}$. But if we interchange the differentiation and summation we would have

$$\frac{d}{dx}\left[\sum_{n=1}^{\infty} \frac{x^n}{n}\right] \;=\; \sum_{n=1}^{\infty} x^{n-1},$$

and this does not converge when $x = -1$ (aside: the sum oscillates between 1 and 0; in some sense it can be interpreted as $\frac{1}{2}$, which is what $\frac{1}{1-x}$ equals when $x = -1$!).

Sometimes, however, conditionally convergent but absolutely divergent series can be managed. Consider

$$\sum_{n=1}^{\infty} \frac{x^n}{n \log n}.$$

This series converges conditionally when $x = -1$ but diverges upon inserting absolute values. If we interchange differentiation and summation we get

$$\sum_{n=1}^{\infty} \frac{x^{n-1}}{\log n},$$

and this sum does converge (conditionally, not absolutely) when $x = -1$.

B.3 Convergence Tests for Series

In calculus classes we learn various tests to determine whether or not a series converges or diverges. There are many reasons for all those hours you spent mastering these, as you're now perfectly prepared to actually *use* these for problems you might care about. In Chapter 19 we'll meet generating functions. These are series that encode a wealth of information about a probability distribution. If these sums converge and are differentiable, then simple differentiation gives us nice formulas for many properties; however, it's sadly not the case that these infinite series always converge. We thus (finally!) see applications of the various series convergence tests from calculus.

As with any result from an earlier course, if you haven't used it in awhile it's easy to be rusty. For completeness we quickly state some of the more popular and powerful tests, and give a few examples illustrating their use. Before doing so, we quickly recall some standard results and notation about series. First, the summation notation:

$$\sum_{n=0}^{N} a_n \;=\; a_0 + a_1 + \cdots + a_N;$$

if instead of N we had ∞ as the upper bound the sum would be $a_0 + a_1 + a_2 + \cdots$. For finite N, we have

$$\sum_{n=0}^{N} (a_n + b_n) \;=\; \sum_{n=0}^{N} a_n + \sum_{n=0}^{N} b_n;$$

if the two sums on the right are finite then this result also holds if $N = \infty$. If these two sums are infinite, however, things are trickier. The problem is one sum could be ∞ and the other $-\infty$, and $\infty - \infty$ is undefined. (Imagine the examples where $a_n = 2n$ and $b_n = -4n$, and $a_n = 2n$ and $b_n = -n$.) If c is any real number,

$$\sum_{n=0}^{N} c a_n \;=\; c \sum_{n=0}^{N} a_n.$$

Root Test: Assume $\lim_{n\to\infty} \sqrt[n]{|a_n|}$ exists, and denote this limit by ρ. Then the series $\sum_{n=0}^{\infty} a_n s^n$ converges for $|s| < 1/\rho$ and diverges for $|s| > 1/\rho$; if $\rho = 0$ we interpret $1/\rho$ as infinity, meaning the series converges for all s. If $\rho = 1$ then there's no information on whether or not it converges or diverges.

Ratio Test: Assume $\lim_{n\to\infty} |a_{n+1}/a_n|$ exists, and denote this limit by ρ. Then the series $\sum_{n=0}^{\infty} a_n s^n$ converges for $|s| < 1/\rho$ and diverges for $|s| > 1/\rho$; if $\rho = 0$ we interpret $1/\rho$ as infinity, meaning the series converges for all s. If $\rho = 1$ then there's no information on whether or not it converges or diverges.

For example, let $a_n = n^2/4^n$. By the ratio test, we have

$$\lim_{n\to\infty} \frac{a_{n+1}}{a_n} \;=\; \lim_{n\to\infty} \frac{(n+1)^2/4^{n+1}}{n^2/4^n} \;=\; \lim_{n\to\infty} \left(\frac{n+1}{n}\right)^2 \frac{1}{4} \;=\; \frac{1}{4}.$$

Thus

$$G(s) \;=\; \sum_{n=0}^{\infty} \frac{n^2}{4^n} s^n$$

converges for $|s| < 4$.

Two other tests that are frequently used are the comparison test and the integral test. These can be a little harder to use, as you need to choose a comparison sequence or function, while the ratio and root tests are automatic (simply compute the limit). That said, with experience these become easier to apply.

Comparison Test: Let $\{b_n\}_{n=1}^{\infty}$ be a sequence of non-negative terms (so $b_n \geq 0$). Assume the series $\sum_{n=0}^{\infty} b_n$ converges, and $\{a_n\}_{n=1}^{\infty}$ is another sequence such that $|a_n| \leq b_n$ for all n. Then the series $\sum_{n=0}^{\infty} a_n$ also converges. If instead $\sum_{n=0}^{\infty} b_n$ diverges and $a_n \geq b_n$, then the series $\sum_{n=0}^{\infty} a_n$ also diverges.

Integral Test: Consider a sequence $\{a_n\}_{n=1}^{\infty}$ of non-negative terms. Assume there's some function f such that $f(n) = a_n$ and f is non-increasing. Then the series

$$\sum_{n=1}^{\infty} a_n$$

converges if and only if the integral

$$\int_1^{\infty} f(x)dx$$

converges; thus if the integral diverges the series diverges.

Note: in both these tests, if instead of starting the sums at $n = 0$ we start at $n = N$, the conclusions still hold; this is because the convergence of series depend only on the tails, and we can add or remove finitely many terms without harm.

Let's determine if the series $\sum_{n=1}^{\infty} \frac{1}{2^n+\sqrt{n}}$ converges or diverges. We use the comparison test. The hardest part about using this test is figuring out what to compare our sequence to. If we think it converges we should find a series that converges that is always greater, while if it diverges we should look for a series that is always small and diverges. When n is large, 2^n is larger than $\sqrt{n}$, and thus the denominator essentially looks like 2^n. We thus expect our series to converge by a comparison with the geometric series $b_n = 1/2^n$. Writing down the algebra formally, we would argue that since $2^n + \sqrt{n} \geq 2^n$, we have

$$0 \leq \frac{1}{2^n + \sqrt{n}} \leq \frac{1}{2^n}.$$

Thus the series converges by the comparison test. We can easily modify this to a problem in generating functions. Consider

$$G(s) = \sum_{n=1}^{\infty} \frac{1}{2^n + \sqrt{n}} s^n.$$

A similar argument shows that we can compare this to $\sum_{n=1}^{\infty}(s/2)^n$, which is a geometric series converging for $|s| < 2$. Thus $G(s)$ converges for $|s| < 2$.

Now let's consider $a_n = \frac{1}{n \ln^p n}$ for some $p > 0$. For which p does it converge? Diverge? We know $\sum_{n=1}^{\infty} \frac{1}{n}$ diverges; unfortunately, this is useless for the comparison test as $\frac{1}{n \ln^p n} \leq \frac{1}{n}$ for n large. If we want to show a series diverges by the comparison test, we must compare it to something smaller that diverges, not something larger. It's hard to find a good series to compare this to, and unfortunately the ratio and root tests don't provide any useful information (as the limit in both cases is 1). We are left with trying the integral test.

The first step is to find a strictly decreasing function $f(x)$ that equals $\frac{1}{n \ln^p n}$ when $x = n$ for n large. Looking at what we've written, you should be able to hear the integral test screaming which function to use: $f(x) = \frac{1}{x \ln^p x}$; it's very common in these problems to just replace n with x. Thus the series converges or diverges depending on whether or not

$$\int_{x=\mathrm{BIG}}^{\infty} \frac{1}{x \ln^p x}\, dx$$

converges or diverges; we write 'BIG' to indicate that the lower bound doesn't really matter – what matters is the behavior at infinity. We use a u-substitution. This is a *very* natural thing to do. The reason is the derivative of $\ln x$ is $1/x$; looking at our integrand, we see it's begging us to change variables as we have $1/x$. We try $u = \ln x$. This gives $du = dx/x$, and thus our integral becomes

$$\int_{u=\ln(\mathrm{BIG})}^{\infty} u^{-p} du.$$

The integral of u^{-p} is $\frac{u^{1-p}}{p}$ if $p \neq 1$ and $\ln u$ if $p = 1$. Thus the integral converges if $p > 1$ and diverges $p \leq 1$.

We can turn this into a statement about the generating function

$$G(s) \ = \ \sum_{n=1}^{\infty} \frac{1}{n \ln^p n} s^n.$$

For any choice of p, with some work you can show the sum converges for $|s| \leq 1$ and diverges for $|s| > 1$.

B.4 Big-Oh Notation

The purpose of this section is to introduce some notation to make it easy for us to compare two quantities as some parameter tends to infinity. If the definition seems technical, there's a natural reason: it is! The entire point of this definition is to allow us to carefully discuss and compare two expressions in some limit situation. The point is to bypass handwaving arguments, to avoid using phrases such as "clearly" and "of course". This notation is used throughout analysis whenever one needs to make rigorous comparisons.

As a motivating example, think of the standard normal and the standard exponential. As the first has density function $\frac{1}{\sqrt{2\pi}}\exp(-x^2/2)$ while the second has the density function $\exp(-x)$, "clearly" the standard normal is decaying faster as $x \to \infty$ than the standard exponential. What we want to do now is clarify how much faster the standard normal decays, as well as avoid using the word "clearly".

Definition B.4.1 (Big-Oh Notation) *$A(x) = O(B(x))$, read "$A(x)$ is of order (or big-Oh) $B(x)$", means there's a $C > 0$ and an x_0 such that for all $x \geq x_0$, $|A(x)| \leq C\,B(x)$. This is also written $A(x) \ll B(x)$ or $B(x) \gg A(x)$.*

Let's unwind this. The part about C is no problem; it's just saying there's some positive constant which will surface later. The purpose of the x_0 constant is to define our universe of discourse. We're saying what happens from some point onward; we're making *no* claims about the behavior for 'small' x; all we're saying is that we know what happens as $x \to \infty$. Specifically, for all large x we have $|A(x)|$ is at most $CB(x)$. Frequently the actual value of C doesn't matter; what's important is the growth (or decay) in x. Additionally, in many problems the inequality holds for each and every x, and thus we don't need to worry about x_0.

Sometimes we use big-Oh notation for $x \to 0$ instead of $x \to \infty$; in that case we modify the definition to there's an x_0 such that for all x with $|x| \leq x_0$ we have $|A(x)| \leq CB(x)$.

In Figure B.1 we plot $A(x) = 2010\log x$ versus $B(x) = x^2$. For small values of x, we see that $A(x)$ is larger; however, as x increases we see eventually $B(x)$ is greater. The reason is that x^2 is growing faster than $\log x$, so in the limit x^2 dominates $\log x$. We can't, however, say that $2010\log x \leq x^2$, though, as this inequality fails for small x. It's only true for x large ($x \geq 100$ suffices). In many problems, we're only interested in making comparisons as our input parameter tends to infinity, and thus such restrictions are fine.

Big-Oh notation is a convenient way to handle lower order terms. For example, if we write $F(x) = x^5 + O(x^2)$, this means that as x tends to infinity, the main term of $F(x)$ grows like x^5, and the correction (or error) terms are at most some constant times x^2.

Not surprisingly, this is used all the time in Taylor series expansions. Consider the Taylor series expansion for $\cos x$:

$$\cos x \;=\; \sum_{n=0}^{\infty} \frac{(-1)^n x^{2n}}{(2n)!} \;=\; 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \cdots .$$

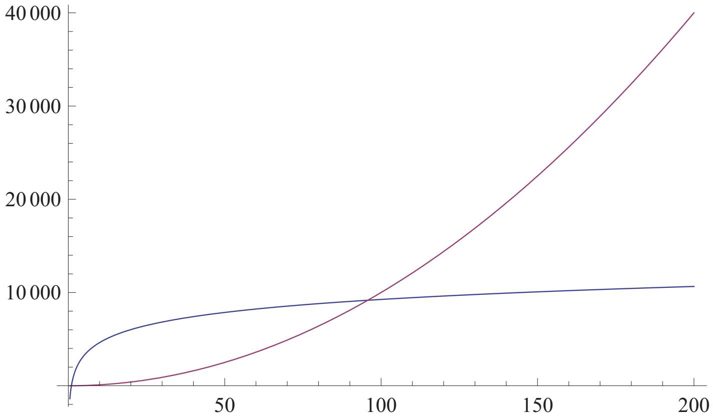

Figure B.1: Plot of $2010 \log x$ versus x^2.

Let's take $x \in [-\pi, \pi]$ near 0, and see how good of a job the various partial Taylor series expansions do of approximating $\cos x$. We have, for instance,

$$\cos x \;=\; 1 - \frac{x^2}{2} + O(x^4),$$

and we claim this works for all x. The reason is the error in the approximation is

$$-\frac{x^4}{4!} + \frac{x^6}{6!} - \frac{x^8}{8!} + \cdots .$$

We can trivially bound this by dropping all the minus signs, and thus the error is at most

$$\frac{x^4}{4!} + \frac{x^6}{6!} + \frac{x^8}{8!} + \cdots .$$

How big is this sum? Remember we plan on taking x near 0, so the higher the power of x, the smaller the contribution. Thus the 'main' term in the error comes from the $x^4/4!$ piece. Pulling that out, we find the error is at most

$$\frac{x^4}{4!}\left(1 + x^2 + x^4 + \cdots\right).$$

For x close to zero, we clearly have $|x| \leq 1/2$ and thus we may use the geometric series formula to evaluate the sum (the ratio is just x^2); note the sum is largest when $|x| = 1/2$ (that's the worst case). We finally see that the error is at most

$$\frac{x^4}{4!}\frac{1}{1-x^2};$$

if we assume $|x| \leq 1/2$ then we finally obtain

$$\left|\cos x - \left(1 - \frac{x^2}{2}\right)\right| \;\leq\; \frac{4}{3}\frac{x^4}{4!} \;=\; \frac{x^4}{18}.$$

In other words, if $|x| \leq 1/2$ the error in using the second order Taylor series to approximate $\cos x$ is quite small, as it's at most $x^4/18$. For example, if we take $x = .1$

then we would say $\cos(.1)$ is approximately $1 - \frac{.1^2}{2} = .995$, with an error that is at most $.1^4/18 \approx 5.5556 \cdot 10^{-6}$. The actual value of $\cos(.1)$ (to ten decimal places) is 0.995004165, which means the true error is about $4.16528 \cdot 10^{-6}$. Note the true error is less than our theoretical bound, so it's likely we have done the algebra correctly!

Two very important relations are that x^r grows slower than e^x for any fixed r as $x \to \infty$, and $\log x$ grows slower than x^c for any $c > 0$ as $x \to \infty$. There are many ways to prove these relations. We prove the first one now to highlight the method, and leave the second one for you. Let's consider x^r versus e^x as $x \to \infty$. We want to show $x^r = O(e^x)$. Clearly this is true if r is negative, so we need only look at $r \geq 0$. If r happened to be an integer, we can use L'Hopital's rule:

$$\lim_{x\to\infty} \frac{x^r}{e^x} = \lim_{x\to\infty} \frac{rx^{r-1}}{e^x} = \lim_{x\to\infty} \frac{r(r-1)x^{r-2}}{e^x} = \cdots = \lim_{x\to\infty} \frac{r!}{e^x} = 0.$$

Why did we assume r was an integer? This is just to make applying L'Hopital a little cleaner; if r is an integer then after applying L'Hopital r times the numerator is just $r!$. As this limit is zero, by definition there's some x_0 such that for $x \geq x_0$ we have $x^r/e^x \leq 1/2$, which gives $x^r = O(e^x)$ if r is a positive integer. For general r, we can either use L'Hopital (ending up with a power of x in the denominator of the fraction), or note that $x^r \leq x^{\lceil r \rceil}$, where $\lceil r \rceil$ represents the smallest integer at least r.

As $x^r = O(e^x)$ is used in numerous problems, we give one more proof. For convenience, let's assume r is an integer. From the Taylor series expansion of e^x, we know $e^x > x^{r+1}/(r+1)!$ (this is because we're keeping just one term). If $x > (r+1)!$, then

$$x^r < \frac{x^{r+1}}{(r+1)!} < e^x.$$

B.5 The exponential function

In this section we study some of the basic properties of the number e. There are many ways to define the number e, the base of the natural logarithm. From the point of view of calculus, the most convenient is through an infinite series:

$$e = \sum_{n=0}^{\infty} \frac{1}{n!}.$$

If we denote the partial sums of the above series by

$$s_m = \sum_{n=0}^{m} \frac{1}{n!},$$

we see e is the limit of the convergent sequence s_m. This representation is one of the main tool in analyzing the nature of e.

We generalize the above and write

The exponential function. Let x be any real (or complex) number. The exponential function e^x (which for typographical purposes is sometimes written $\exp(x)$ when the argument is complicated) is defined at

$$e^x \;=\; \sum_{n=0}^{\infty} \frac{x^n}{n!}.$$

Further, $e^{x+y} = e^x e^y$.

We call the above the exponential function. As remarked, we frequently use the $\exp$ notation for typographic purposes; for example, $\exp(-x^2/2)$ is a little easier to read than $e^{-x^2/2}$ or, even worse, $e^{-\frac{x^2}{2}}$!

The series defining the exponential function converges so rapidly that almost any test works. Let's use the Ratio Test, as it's easy to apply. We have

$$\begin{aligned}
\rho &= \lim_{n\to\infty} \frac{|a_{n+1}|}{|a_n|} \\
&= \lim_{n\to\infty} \frac{|x|^{n+1}/(n+1)!}{|x|^n/n!} \\
&= \lim_{n\to\infty} \frac{|x|}{n+1} \;=\; 0.
\end{aligned}$$

Thus, the series converges for all x.

This notation is meant to be highly suggestive, and is designed to make you think about raising numbers to powers. We read e^x as e raised to the x power. If asked what is $e^x e^y$, you should immediately answer e^{x+y}; however, it's very important to note that this is *not* obvious and this needs to be proved! Technically e^x, e^y and e^{x+y} are three different infinite sums, and we must show the product of the first two equals the third. Of course, if this were not true then our notation would suck (no other word feels right for how horrible our notation would be); unfortunately, math does occasionally have bad notation. I've always hated that cosecant is one over sine and not one over cosine.

The proof that $e^x e^y = e^{x+y}$ is a nice application of the binomial theorem. We have

$$e^x e^y \;=\; \sum_{m=0}^{\infty} \frac{x^m}{m!} \sum_{n=0}^{\infty} \frac{y^n}{n!}.$$

Note that we used two different letters for our summations. It's a very common mistake to use the same letter twice; we can't and shouldn't do this. The reason it's wrong to use the same letter is that we have two sums, and each sum has a dummy variable for summation (similar to the dummy variables of integration). Consider for example

$$(1+2+3)\cdot(1^2+2^2+3^2) \;=\; 84.$$

If we use the same dummy variable, we might be led to the following flawed calculation:

$$\sum_{n=1}^{3} n \sum_{n=1}^{3} n^2 \ = \ \sum_{n=1}^{3} n^3 \ = \ 1^3 + 2^3 + 3^3 \ = \ 36.$$

Using a different letter for each sum minimizes our chance of making such a mistake.

Returning to our analysis of $e^x e^y$, we see we have a sum over terms of the form $\frac{x^m y^n}{m!n!}$, with $m, n \geq 0$. What we will do now is collect all terms where the sum of the power of x plus the power of y is constant. In other words, for a given $k \geq 0$ let's look at all pairs (m, n) with $m + n = k$. We need to introduce one more dummy variable. Let's let ℓ equal the power of x. If the power of x plus the power of y is k, this means that the power of y is $k - \ell$ whenever the power of x is ℓ; furthermore, ℓ ranges from 0 to k (as the powers of x and y are non-negative integers). Collecting, we find

$$e^x e^y \quad = \quad \sum_{k=0}^{\infty} \sum_{\ell=0}^{k} \frac{x^\ell y^{k-\ell}}{\ell!(k-\ell)!}.$$

We now need to do some pattern recognition. Note the denominator looks a lot like a binomial coefficient; it's the bottom of $\binom{k}{\ell}$. This suggests **multiplying by one** (see §A.12 for more examples), in this case $k!/k!$.

Note that the denominator invokes thoughts of binomial coefficients. Specifically, $\binom{k}{\ell} = \frac{k!}{\ell!(k-\ell)!}$. If we multiply by 1 in the form $k!/k!$, we'll see the binomial coefficient emerge:

$$\begin{aligned} e^x e^y \quad &= \quad \sum_{k=0}^{\infty} \sum_{\ell=0}^{k} \frac{1}{k!} \frac{k!}{\ell!(k-\ell)!} x^\ell y^{k-\ell} \\ &= \quad \sum_{k=0}^{\infty} \frac{1}{k!} \sum_{\ell=0}^{k} \binom{k}{\ell} x^\ell y^{k-\ell} \\ &= \quad \sum_{k=0}^{\infty} \frac{1}{k!} (x+y)^k \\ &= \quad \sum_{k=0}^{\infty} \frac{(x+y)^k}{k!} \ = \ e^{x+y}, \end{aligned}$$

where we used the binomial theorem to replace the ℓ sum with $(x+y)^\ell$ and we used the series expansion to replace the k sum with e^{x+y}. □

All that matters from the above discussion is that our intuition is correct, and our notation is good. It's also worth noting the power of multiplying by 1. This is one of the hardest math skills to learn, but one of the most important. We can always multiply by 1 (or do something similar, add zero); the trick is finding *good* ways to do this which lead to simpler expressions.

There is another definition of e^x, which also arises in probability. You might remember it from compound interest problems where the money is compounded instantaneously. This definition is very useful in proving the Central Limit Theorem for certain sums of independent random variables.

An alternative definition of e^x is

$$e^x \;=\; \lim_{n\to\infty}\left(1+\frac{x}{n}\right)^n.$$

A nice exercise is to show that this definition agrees with the series expansion.

No introduction to e^x would be complete without a few words about its derivative. Using the series expansion, the natural temptation is to differentiate term by term, which gives

$$\begin{aligned}\frac{d}{dx}e^x &= \frac{d}{dx}\left(1+x+\frac{x^2}{2!}+\frac{x^3}{3!}+\cdots\right)\\ &= 1+\frac{2x}{2!}+\frac{3x^2}{3!}+\frac{4x^3}{4!}+\cdots\\ &= 1+x+\frac{x^2}{2!}+\frac{x^3}{3!}+\cdots \;=\; e^x.\end{aligned}$$

Of course, we need to justify interchanging a sum and a derivative. This is typically done in an advanced analysis course.

Without using a calculator or computer, determine which is larger: e^π or π^e. *Hint:* One approach is to study the function $x^{1/x}$ (take the $e\pi$ root of both sides to reduce the problem to comparing $e^{1/e}$ and $\pi^{1/\pi}$. Use calculus to find the maximum value. One could also study $f(x) = e^x - x^e$ and try to show $f(x) > 0$ when $x > e$; however, it's hard to analyze all the critical points. It's easier to study $g(x) = e^{x/e} - x$, and show $g(x) > 0$ for $x > e$.

B.6 Proof of the Cauchy-Schwarz Inequality

Our last analysis result is the Cauchy-Schwarz inequality, which is very useful in bounding certain integrals.

Lemma B.6.1 (Cauchy-Schwarz Inequality) *For complex-valued functions f and g,*

$$\int_{-\infty}^{\infty}|f(x)g(x)|dx \;\le\; \left(\int_{-\infty}^{\infty}|f(x)|^2dx\right)^{1/2}\cdot\left(\int_{-\infty}^{\infty}|g(x)|^2dx\right)^{1/2}. \tag{B.3}$$

Proof of the Cauchy-Schwarz inequality: For notational simplicity, assume f and g are non-negative functions. Working with $|f|$ and $|g|$ we see there's no harm in the above assumption. As the proof is immediate if either of the integrals on the right hand side of (B.3) is zero or infinity, we assume both integrals are non-zero and finite. Let

$$h(x) = \; f(x) - \lambda g(x), \;\; \lambda \;=\; \frac{\int_{-\infty}^{\infty} f(x)g(x)dx}{\int_{-\infty}^{\infty} g(x)^2dx}.$$

As $\int_{-\infty}^{\infty} h(x)^2 dx \geq 0$ we have

$$\begin{aligned} 0 &\leq \int_{-\infty}^{\infty} \left(f(x) - \lambda g(x)\right)^2 dx \\ &= \int_{-\infty}^{\infty} f(x)^2 dx \;-\; 2\lambda \int_{-\infty}^{\infty} f(x)g(x)dx \;+\; \lambda^2 \int_{-\infty}^{\infty} g(x)^2 dx \\ &= \int_{-\infty}^{\infty} f(x)^2 dx \;-\; 2\frac{\left(\int_{-\infty}^{\infty} f(x)g(x)dx\right)^2}{\int_{-\infty}^{\infty} g(x)^2 dx} \;+\; \frac{\left(\int_{-\infty}^{\infty} f(x)g(x)dx\right)^2}{\int_{-\infty}^{\infty} g(x)^2 dx} \\ &= \int_{-\infty}^{\infty} f(x)^2 dx \;-\; \frac{\left(\int_{-\infty}^{\infty} f(x)g(x)dx\right)^2}{\int_{-\infty}^{\infty} g(x)^2 dx}. \end{aligned}$$

This implies

$$\frac{\left(\int_{-\infty}^{\infty} f(x)g(x)dx\right)^2}{\int_{-\infty}^{\infty} g(x)^2 dx} \quad \leq \quad \int_{-\infty}^{\infty} f(x)^2 dx,$$

or equivalently

$$\left(\int_{-\infty}^{\infty} f(x)g(x)dx\right)^2 \quad \leq \quad \int_{-\infty}^{\infty} f(x)^2 dx \cdot \int_{-\infty}^{\infty} g(x)^2 dx.$$

Taking square roots completes the proof. □

This proof uses one of the most important identities in all of mathematics: if u is a real number then $u^2 \geq 0$. The clever part is in choosing u. For those loving a challenge, think why this works. Why is this a good choice? A good starting point is to determine when the Cauchy-Schwarz inequality is an equality.

B.7 Exercises

Problem B.7.1 *In our proof of Rolle's Theorem we assumed $f'(a)$ and $f'(b)$ were non-zero; handle the case when one of these vanish.*

Problem B.7.2 *One cannot always interchange orders of integration. For simplicity, we give a sequence a_{mn} such that $\sum_m(\sum_n a_{m,n}) \neq \sum_n(\sum_m a_{m,n})$. For $m, n \geq 0$ let*

$$a_{m,n} \;=\; \begin{cases} 1 & \text{if } n = m \\ -1 & \text{if } n = m + 1 \\ 0 & \text{otherwise.} \end{cases}$$

Show that the two different orders of summation yield different answers (the reason for this is that the sum of the absolute value of the terms diverges).

Problem B.7.3 *In justifying interchanging a derivative and a sum we needed the existence of an N_0 such that for all $N \geq N_0$,*

$$\left| \frac{(N+1)(1-x)+x}{(1-x)^2}\, x^N \right| \;\leq\; \frac{\epsilon}{2}$$

(where ϵ is a fixed positive number). Find an N_0 that works (your answer should depend on ϵ).

Problem B.7.4 *Consider the dominated convergence theorem. Show that its conclusion need not hold if there is no non-negative, piecewise continuous function g such that $|f_n(x)| \leq g(x)$ for all n.*