

Proofwriting Checklist

Over the years, we've found many common proofwriting errors that can easily be spotted once you know how to look for them. In this handout, we've distilled seven major points about proofwriting that we will specifically be looking for when grading your assignments. They are as follows:

- Clearly articulate your start and end points.*
- Make each sentence "load-bearing."*
- Scope and properly introduce variables.*
- Make specific claims about specific variables.*
- Don't repeat definitions; use them instead.*
- Write in complete sentences and complete paragraphs.*
- Distinguish between proofs and disproofs.*

Some of the items on this list, like "write in complete sentences and complete paragraphs," are purely stylistic requirements on proofs. They're there because they ensure that you're writing proofs in the expected mathematical style. Other items on this list, like "scope and properly introduce variables," are there because they're often comorbid with more serious logic errors that can derail a proof. Our hope is that by providing these specific items to look for when checking your proofs, you'll be able to check your own work more effectively and build a better intuition for when there's something in a proof that just doesn't feel right.

We will be applying this checklist to the proofs that you submit. We strongly recommend that you work through this checklist on every proof that you write. Doing so will help you improve your proofwriting and possibly smoke out some underlying logic errors.

The remainder of this handout goes into more detail about what each of these rules mean.

Clearly Articulate Your Start and End Points

When you're writing a proof, you're laying out an argument that explains why a certain result is true. Most proofs have a number of intermediate steps that build up toward a larger result. When writing a proof, it's important to make sure that the reader has a clear sense of where it is that you're going and how you're going to arrive there. Otherwise, your proofs will be extremely hard to read, since while the reader might follow each *individual* step, they might have no idea where you're going with things. Think about how you might write an argumentative essay – if you just list a series of facts without giving some idea of where you're ultimately going, your readers are going to have a heck of a time trying to make sense of what you're doing!

Let's illustrate this with an example. Consider the following proof:

Incorrect! **Proof:** Consider an arbitrary $x \in A$. Since $x \in A$ and $A \subseteq B$, we see that $x \in B$. And, since $x \in B$ and $B \subseteq C$, we see that $x \in C$, as required. ■

Here's a question for you – what exactly is this proof trying to accomplish? It's hard to say, since we don't know that A , B , and C are, it seems like the statements $A \subseteq B$ and $B \subseteq C$ come out of nowhere, and the conclusion doesn't say exactly why any of this matters.

The above proof was written for the following theorem:

Theorem: If $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.

With knowledge of the theorem in mind, the proof makes more sense. We know that $A \subseteq B$ and that $B \subseteq C$ by assumption, and we're looking at elements of A and trying to get them as elements of C because we're trying to prove something about the subset relation. But that still shifts a lot of work to the person reading the proof. A better proof would provide more guidance about where everything comes from and where everything is going. Here's what that might look like:

Proof: Let A , B , and C be arbitrary sets where $A \subseteq B$ and $B \subseteq C$. We will prove that $A \subseteq C$. To do so, choose an arbitrary $x \in A$. We will prove that $x \in C$.

Since $x \in A$ and $A \subseteq B$, we see that $x \in B$. And, since $x \in B$ and $B \subseteq C$, we see that $x \in C$, which is what we needed to show. ■

Compare this proof to the one before it. Even if you had no idea what the theorem was when going into this proof, you could still see exactly what's being done – what's being assumed, what's being proved, how the logic flows, etc. There's no more mystery about why $A \subseteq B$ and $B \subseteq C$ are true: we can see that they're true by assumption.

There's a number of reasons why it's worthwhile to set up your proofs this way. First, when you're still working through the problem and trying to figure out why exactly the result is true, this step forces you to write out exactly what it is that you're assuming and what you need to prove. That makes it much easier to figure out what directions you should consider. It also forces you to articulate very precisely what it is that you need to establish. If you look at the overall theorem to prove here, it might seem, well, kinda *obvious*. Like, “well, *of course* if A is a subset of B and B is a subset of C , then A is a subset of C – that's just what subset *means!*” But if you start unpacking the definitions and articulating where specifically you're going to start and end, it becomes much easier to see what you need to do.

Why we enforce this rule: When you're first learning how to write proofs, one of the biggest challenges is simply figuring out what it is that you're supposed to assume and what it is that you're supposed to prove. By requiring that you articulate this information clearly, we hope to reduce the likelihood that you submit a proof that is a completely correct proof of the completely wrong theorem.

Make Each Sentence Load-Bearing

When you're writing a proof, you are trying to convey a mathematical argument, and each step in what you write should advance your argument. As a general rule, every statement in a proof should do one of the following things:

- **Set up a goal.** As mentioned in the preceding pages, your proof should start off with an introduction that clearly articulates a start and end point. In larger proofs, you might find yourself needing to prove an auxiliary result that you'll use to build up to the larger result, and when you do that, you'll similarly want to set up what it is that you're trying to prove.
- **Introduce a new variable.** Sometimes, in the course of a proof, you'll need to introduce new variables. If you're proving something universally-quantified, you might want to say something like "let x be an arbitrary bananafish," and if you're proving something existentially-quantified you might want to say something like "since n is even, we know there is an integer k such that $n = 2k$."
- **Combine preceding results into something new.** Any sentence that doesn't set up a new goal or introduce a new variable should make progress toward the result by taking some number of preceding statements and deriving some new, *mathematically rigorous* result from those preceding statements. For example, you might say something like "since $n = 2k$, we see that $n^2 = 2(2k)^2$ " or "since $A \subseteq B$ and $x \in A$, we learn that $x \in B$."

If you find yourself reading over a sentence that doesn't accomplish any of these goals, it is likely unnecessary and should be eliminated. This is a great way to reduce the size of your proofs and to make sure that you're being rigorous.

This is a particularly useful check to apply to a proof after you've first finished writing it, since often times in the course of solving a problem and writing up a first proof draft you'll go in a direction that ultimately ends up not being necessary, or write out some high-level lines of reasoning that you then make more rigorous later on.

Why we enforce this rule: We enforce this rule for a number of reasons.

First, this rule is designed to get you to review your proofs after having written a first draft. It's common, in writing up the first version of a proof, to include statements that aren't actually needed later on, and by requiring each statement to be load-bearing we hope to encourage you to closely review your work to make sure that everything you've included ends up getting used.

Second, this rule is there to make sure that you are being precise with your reasoning. If you find that your proofs include sentences that talk about how things tend to work in general, or which describe a mathematical situation without the precision required above, it might indicate that you haven't pinned things down as tightly as you may have expected.

Finally, this rule is here because this is just how proofs are expected to be written these days. It's common in mathematics to separate mathematical proofwriting from mathematical expository writing. In an exposition, an author might talk about various intuitions they've had, various insights that will make things easier to understand, etc., but in the proof itself it's common for sentences to be fairly direct and to the point.

Scope and Properly Introduce Variables

In programming languages like C, C++, and Java, you're required to declare variables before you use them. The type of the variable lets the reader (and the compiler!) know what sort of thing the variable can hold and what it represents. If you try to use a variable you haven't declared, or if you try to treat a variable of one type as though it had a different type, you get a compiler error because there's something amiss with what you've done.

Variables in mathematical proofs obey a similar sort of convention. When writing proofs, it's important that you clearly and precisely articulate what each variable stands for and, additionally, where it comes from. When you use a variable in a proof, you should explicitly articulate

- the name of the variable,
- what value it represents, and
- where it comes from.

Those last two points are critical in writing proofs. Every variable that you use should be of one of the following types:

- **An arbitrarily-chosen value.** A variable like this doesn't represent some specific number, set, or quantity, but rather an arbitrarily-chosen value. Variables like these often arise in the context of proving universally-quantified statements. For example, if you want to prove the claim "for any natural number n , if n is even, then n^2 is even," you might introduce a variable n like this:

Let n be an arbitrary even natural number.
Consider an even natural number n .
Let n be an even natural number.

Here, we're indicating that the variable is named n , its value is some even natural number, and that it's chosen arbitrarily.

- **An existentially instantiated value.** Sometimes, you know that some quantity must exist, but you don't know what it is. For example, if you know that n is an even natural number, you know that n must be twice some other natural number, and so you might give it a name, as shown here:

Since n is even, there is some integer k such that $n = 2k$.

It's important to note that this number k is *not* chosen arbitrarily. That would imply that any choice of k would work here, but that's not true: there's only one choice of k you can pick where $n = 2k$. Rather, k is called an **existentially instantiated** variable, because we know that *there exists* some value with some property and we're introducing the variable k as a way of saying what that value is.

- **An explicitly chosen value.** Sometimes, you'll introduce a variable simply as a simpler way of referring to some other quantity. For example, we might say something like this:

Let $m = 2k^2$.

Or, we could say something like this:

Consider the set $D = \{ x \in S \mid x \notin f(x) \}$.

Here, we're just giving a name to an existing quantity, which functions like a constant in a language like C, C++, or Java.

When you write up a proof (or, more generally, when you're reading something mathematical), it's important to make sure that you can look at each variable and clearly tell whether that variable is arbitrarily chosen, existentially instantiated, or explicitly chosen. Just like variables in C, C++, or Java, this helps you clearly indicate what your variables mean, what they store, and where they're coming from.

One particular caveat to watch out for: some variables in mathematics are true placeholders that don't actually stand for anything. For example, in set-builder notation, we use placeholder variables to denote the name of some unknown quantity:

$$\{ n \in \mathbb{N} \mid n \text{ is even and } n^2 > 48 \}$$

In this context, n does not represent a value. It's just a placeholder so that we can write the expression “ n is even and $n^2 > 48$ ” in a way that's clear and easy to follow. It's an error to try to reference the number n out of this context.

To see how these rules come into play, let's look at one possible proof of this result:

For any sets A , B , and C , if $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.

Here's a not-so-great proof of this result:

Incorrect! **Proof:** Let A , B , and C be arbitrary sets where $A \subseteq B$ and $B \subseteq C$. This means that for any choice of x , if $x \in A$, then $x \in B$. Similarly, for any choice of x , if $x \in B$, then $x \in C$. We need to prove that $A \subseteq C$, which means that we need to prove that for any choice of x , if $x \in A$, then $x \in C$.

To show this, consider any $x \in A$. Since $x \in A$ and we know that any $x \in A$ must also be an element of B , we see that $x \in B$. Similarly, since $x \in B$ and we know that any $x \in B$ must also be an element of C , we see that $x \in C$, which is what we needed to show. ■

Let's focus on a few of sentences. For starters, let's look at this sentence from the first paragraph:

This $[A \subseteq B]$ means that for any choice of x , if $x \in A$, then $x \in B$.

What, exactly, is the variable x here? It's not an arbitrarily-chosen x , since we didn't say something like “choose an arbitrary x .” Instead, it's a placeholder: it says that *if we find some* x where $x \in A$, then we can conclude that $x \in B$. All that we've done here is set up some possible confusion for later on in the case where we do define some variable named x .

Think back to Rule Three. Every sentence in a proof should set up a goal, introduce a variable, or combine results together into something new. This sentence doesn't set up a goal. It doesn't introduce a new variable. In a sense it *kinda* combines results together into something new, but really, it's not doing that. It's just restating the definition of what a subset is. As a result, this sentence probably fails Rule Three and should be cut.

This sentence actually does cause problems later in the proof, specifically in these sentences:

To show this, consider any $x \in A$. Since $x \in A$ and we know that any $x \in A$ must also be an element of B , we see that $x \in B$.

In the first sentence, we introduce a new variable x , which is chosen as an arbitrary element of the set A (which is fine by both Rule Three and Rule Four). You can imagine that the reader is going to look at this and say “okay, I'm going to pick some specific thing x .” In the next sentence, though, the proof talks about “any $x \in A$.” Now the reader is going to be confused: “hold on, are you talking about the x that you just asked me to pick in the preceding sentence, or are you talking about some other thing called x ?”

Think of it this way: the following code wouldn't be legal in C, C++, or Java:

```
int x = 137;
int x = 42; // Error!
```

The issue here is that x is already defined on the first line, so the second line is a variable redefinition error. If you want to talk about x going forward, just use its name, not its type:

```
int x = 137;
x = 42; // Okay!
```

The same is true of proofs. Phrases like “any x ,” “every x ,” or “any choice of x ” suggest that you’re introducing some new variable, rather than referring to an existing variable.

A better way to rewrite the above sentences would be to write something like this:

<i>Before</i>	<i>After</i>
To show this, consider any $x \in A$. Since $x \in A$ and we know that any $x \in A$ must also be an element of B , we see that $x \in B$.	To show this, consider any $x \in A$. Since $x \in A$ and $A \subseteq B$, we see that $x \in B$.

Something to specifically keep an eye out for arises when you switch between telling the reader *what you’re going to prove* and then *actually going and proving it*. For example, suppose that you want to prove this claim:

For any sets A and B , we have $A \cap B \subseteq A$.

Here’s a not-so-great way of proving this:

Incorrect! ***Proof:*** Let A and B be arbitrary sets. We will prove that $A \cap B \subseteq A$ by showing that every $x \in A \cap B$ satisfies $x \in A$. To see this, notice that since $x \in A \cap B$, we know that $x \in A$ and $x \in B$. In particular, this means that $x \in A$, as required. ■

There’s a subtle but important shift in the meaning of the variable x between the second and third sentences. In the second sentence (“We will prove that ...”), the variable x is a placeholder: it doesn’t actually stand for any specific value. In the third sentence (“To see this, ...”), the variable x is being used as though it’s an actual, concrete value. This is a problem, since we don’t know precisely what value x has. A better way to write this proof would be to explicitly pick x arbitrarily:

(Better) Proof: Let A and B be arbitrary sets. We will prove that $A \cap B \subseteq A$ by showing that every $x \in A \cap B$ satisfies $x \in A$. To see this, consider any $x \in A \cap B$. Notice that since $x \in A \cap B$, we know that $x \in A$ and $x \in B$. In particular, this means that $x \in A$, as required. ■

Why we enforce this rule: We tend to be fairly strict about this rule, and that can really catch people off-guard who aren’t expecting it. So why is that? There are two main reasons.

First, requiring that each variable have a clear, precise, unambiguous meaning tends to markedly improve the precision of the proof. Many mathematical errors arise when talking about how things work “in general” or by making overly broad statements about how certain classes of objects work. On the other hand, if you’ve singled out some specific object and given it a name, then there’s no need to make those broad claims. You just need to talk about the specific object that you have, referring to it by the specific name that you’ve chosen. From experience, proofs that do not pin things down at this level of detail tend to have more errors and to miss important but subtle details.

Second, this level of precision when speaking about variables requires that you, the writer, have a clear and unambiguous sense of what every term means. Many mistakes in proofs arise from swapping the meaning of one variable for another (for example, using a variable n to refer to two different natural numbers), or confusing a known and unknown quantity (for example, using a variable k that needs to be solved for rather than trying to deduce what it is). Articulating what each variable means makes it harder to make these sorts of mistakes and forces you to slow down as you’re writing to reflect on these details.

Make Specific Claims About Specific Variables

When you're first learning to write proofs, it's common to want to write proofs that make broad claims about how things work *in general* rather than pinning down the specifics. For example, consider this not-so-great proof that if $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.

Incorrect! **Proof:** Let A , B , and C be arbitrary sets where $A \subseteq B$ and $B \subseteq C$. We will prove that $A \subseteq C$.

Since $A \subseteq B$, we see that every element of A is an element of B . Similarly, since $B \subseteq C$, we see that every element of B is an element of C . Therefore, every element of A is an element of C , so by definition $A \subseteq C$, as required. ■

The intuition underlying this proof is good, but the way this is written is far too high-level. Specifically, remember that the definition of the statement $A \subseteq C$ is the following:

For every x , if $x \in A$, then $x \in C$.

In order to prove this claim by calling back to the definition, you'd need to show that if you chose an arbitrary element $x \in A$ that you'd find $x \in C$. The proof given above does not do this. The idea behind it – that anything in A is in B and anything in B is in C – is totally correct, but that's not how you'd phrase it in a proof. In proofwriting, if you want to make a claim that something is true in the general case, do so by using arbitrary choices or a proof by contradiction. For example:

Rewrite this...

Since $A \subseteq B$, every element of A is an element of B .

The function f maps different inputs to different outputs.

→

→

... like this

Consider any element $x \in A$. Since $A \subseteq B$ and $x \in A$, we see that $x \in B$.

Consider any arbitrary x and y where $x \neq y$. Then we know that $f(x) \neq f(y)$.

When you're reading over your proofs, take a minute to check whether you are making *specific, precise claims* about *named variables* or *broad, general claims* about *all objects of a certain type*. If you find yourself doing the latter, rewrite it to use the former. This will both clarify your reasoning and make it significantly harder to make mistakes. Plus, if you find that you *can't* pin down precisely what you mean about something, it might indicate that there's some concept you're having trouble with.

Why we enforce this rule: This rule – another one that we tend to be fairly strict about – is designed to make sure that you've properly justified each step of your reasoning by calling back to the appropriate definitions.

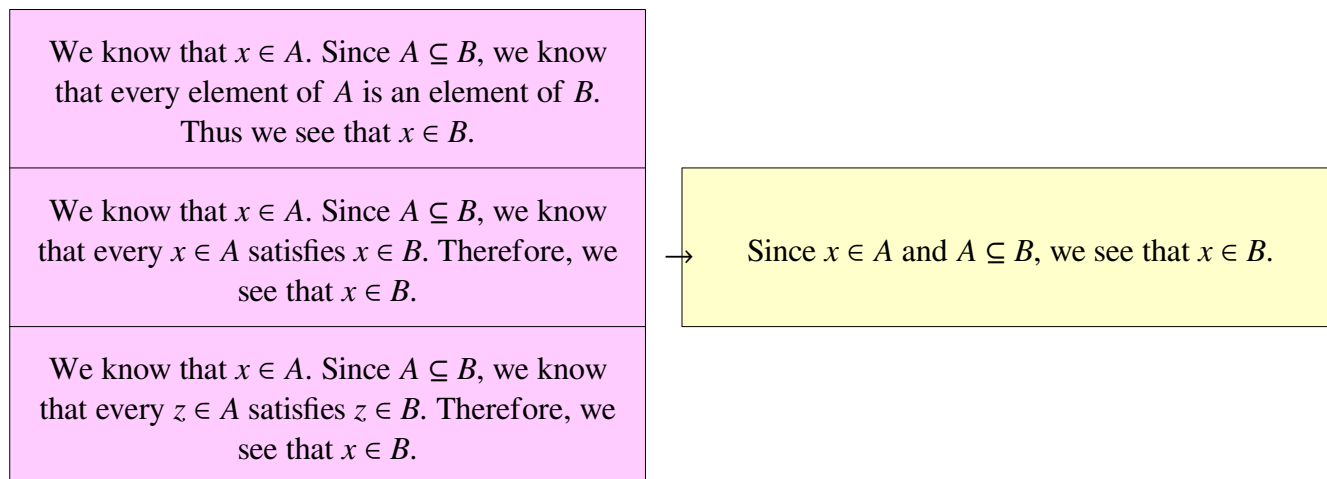
When you're first studying proof-based mathematics, you'll likely have a number of intuitions about how different types of objects behave. Some of these intuitions are great, and you should keep using them. Other intuitions, on the other hand, can be at odds with what the math says, and when that happens, you should refine those intuitions so that they guide you in the right direction.

The only way to know which of your intuitions are good and which need tuning is to explicitly validate those intuitions by attempting to formalize them mathematically. To do so, we ask that you speak with mathematical precision and to show how specific applications of definitions give you your result. If you're able to do this, great! It likely means your intuition is pointing you the right way. If not, that might indicate that your intuition might be suggesting something that the math says isn't true, in which case it's a good thing that you tried formalizing things! At that point, you should back up, pause, and see whether the result is still true for some other reason or whether you need to reshape your intuition for the future.

Don't Repeat Definitions; Use Them Instead

Mathematical definitions are wonderfully useful. They give us a way to take an intuitive idea like “even numbers” and to formalize them in a way that lets us manipulate them in proofs.

Most mathematical proofs will in some way, shape, or form touch on formal definitions. However, you should avoid restating definitions purely in the abstract and instead focus on how those definitions are specifically useful or relevant for what you're trying to do. For example, we recommended replacing statements like the ones on the left with one like what's on the right:



There are a few reasons why it's wise to avoid repeating definitions in the abstract. First, you can assume that the reader knows all of the relevant terms and definitions that are needed in your proofs. Your job as a proofwriter is not to convince the reader of what the definitions are, but to show how those definitions interact with one another to build into some result. In that sense, repeating a definition in the abstract, like what's done above and to the left, doesn't actually contribute anything to the argument you're laying out. The reader already knows the definition, so that sentence is fully redundant.

Second, restating definitions in the abstract risks violating other checklist items. Let's go one at a time through the three options on the left that we advise against. The first one is far too general (“every element of A is an element of B ”) and therefore breaks our advice of making specific claims about specific variables. The second one (“every $x \in A$ satisfies $x \in B$ ”) is a variable scoping error – is x the specific value referred to in the first sentence, or is it a placeholder? The third one is making specific claims about the variable z and doesn't have a scoping error, but in that case z is purely a placeholder – it doesn't refer to any value. In each of those cases, you can safely delete things.

And finally, restating definitions in the abstract just makes things longer. Compare the three options to the left to the one on the right. All three of those proof fragments are significantly longer than the more concise and direct version shown to the right.

Why we enforce this rule: Brevity. Simply put, this rule is a great way to reduce the amount of writing you do and to keep your proofs short.

The other reason we enforce this rule is that it reduces the space of possible errors you can make. As we've mentioned earlier, using placeholder variables is an easy way to run into trouble, either by confusing one variable for another or by thinking you've proved something that you actually haven't. Asking that you apply definitions rather than repeat them reduces the number of placeholder variables you have to work with in your proof, eliminating many potential opportunities for error.

Write In Complete Sentences and Complete Paragraphs

Although proofs exist to convey mathematical arguments, the expectation is that they should be written in grammatically-correct English sentences and in paragraph form.

A good test we recommend applying to your proofs is what we call the *mugga mugga test*. Take your proof and try reading it out loud, replacing all the mathematical content with the phrase “mugga mugga.” If what comes back is grammatically correct, then you’re on the right track! On the other hand, if what you write is hard to read aloud, or just plain doesn’t sound right, it means that you might need to go back and correct things. As an example, here’s a not-so-great proof that if n is even, then n^2 is even:

Incorrect! **Proof:** If n is even, $n = 2k$. $n^2 = 4k^2$, which is $2(2k^2)$. $2k^2 \in \mathbb{Z}$, so n^2 is even. ■

Let’s apply the mugga mugga test to this proof, one sentence at a time. Here’s the first sentence:

Original

If n is even, $n = 2k$.

Mugga Mugga Version

If n is even, *mugga mugga*.

The mugga-muggaified version of this sentence isn’t grammatically correct – it has no subject and no verb. The reason for this is that the subject of the original sentence is n and the verb is “equals,” but since we’ve written out the equality using the equals sign, it got mugga-muggified in the updated version of the sentence.

More generally:

Tip: Avoid writing sentences where mathematical notation must be treated as a verb.

So what should we do instead? Let’s begin with what you *shouldn’t* do. Don’t rewrite the sentence like this in order to pass the mugga mugga test:

If n is even, n equals $2k$.

This technically passes the mugga mugga test, but it’s doing so by taking a clear mathematical statement ($n = 2k$) and rendering the unambiguous, precise mathematical symbol $=$ in English. The whole reason for having mathematical symbols in the first place is so that we can be precise with our notation, and this is a step in the wrong direction.

Instead, consider rewriting the sentence in a way that introduces a new subject and a new verb. There are many ways that we can do this. Here are a few options to choose from:

Original

If n is even, then we can write $n = 2k$.

Since n is even, we see that there is some integer k such that $n = 2k$.

Because n is even, it can be expressed as $n = 2k$ for some integer k .

Mugga Mugga Version

If n is even, then we can write *mugga mugga*.

Since n is even, we see that there is some integer k such that *mugga mugga*.

Because n is even, it can be expressed as *mugga mugga* for some integer k .

Notice how in each sentence we’ve introduced an explicit subject and verb in a way that passes the mugga mugga test.

Let's look at this second sentence:

Original

$$n^2 = 4k^2, \text{ which is } 2(2k^2).$$

Mugga Mugga Version

Mugga mugga, which is mugga mugga.

Again, we're failing the mugga mugga test because the subject and verb of the sentence are expressed in mathematical notation. We'd be better off rewriting this sentence in one of the following ways:

Revision

We can rewrite the expression $n^2 = 4k^2$ as
 $n^2 = 2(2k^2)$.

Rewriting $4k^2$ as $2(2k^2)$, we see that
 $n^2 = 2(2k^2)$.

Mugga Mugga Version

We can rewrite the expression *mugga mugga* as *mugga mugga*.

Rewriting *mugga mugga* as *mugga mugga*,
we see that *mugga mugga*.

A common theme in the mugga mugga test is that you should avoid using math notation as the verb in a sentence. Similarly, you should avoid using mathematical notation or shorthands to abbreviate parts of sentences. There are a number of shorthands that have been developed over the years, primarily for use on blackboards where writing out longhand can take a while. For example, the word "therefore" is often abbreviated \therefore , and the word "because" is often abbreviated \because . These shorthands are just that – they're shorthands – and should not be used in mathematical proofs except if you're trying to write something up quickly and on a blackboard. For example, please, please, please don't write the following:

$$\because n \text{ is even, } n = 2k \text{ for some integer } k, \therefore n^2 = 4k^2 = 2(2k^2), \therefore n^2 \text{ is even } \because n^2 = 2m \text{ for } m = 2k^2.$$

This one really, really, really fails the mugga mugga test:

Original

$\because n$ is even, $n = 2k$ for some integer k ,
 $\therefore n^2 = 4k^2 = 2(2k^2)$, $\therefore n^2$ is even $\because n^2 = 2m$
for $m = 2k^2$.

Mugga Mugga Version

Mugga mugga n is even, *mugga mugga* for
some integer k , *mugga mugga mugga mugga*,
mugga mugga n² is even *mugga mugga*
mugga mugga for *mugga mugga*.

This almost reads like a parody of a terrible math lecture. So please don't write proofs like this. ☺

Just as you're expected to write in complete sentences, you're expected to write in complete *paragraphs*. This means that your proofs should not consist of bulleted or numbered lists of statements. For example, please don't write proofs like these:

- Let n be an even integer.
- Since n is even, we can write $n = 2k$ for some integer k .
- Then $n^2 = 4k^2$.
- So $n^2 = 2(2k^2)$.
- Let $m = 2k^2$.
- [etc.]

Although we can see what this proof is saying, this just isn't the format that's expected and so you shouldn't structure things this way.

Why we enforce this rule: We primarily enforce this rule because this is the standard convention in mathematical writing and we're hoping to train you to communicate mathematics effectively. Additionally, this rule makes proofs much easier to read by requiring you, the writer, to link your ideas together in a way that helps the argument flow.

Distinguish Between Proofs and Disproofs

The short version of this section goes as follows:

- A **proof** is an argument that explains why some **theorem** is true.
- A **disproof** is an argument that explains why some **claim** is false.
- Don't write a proof by contradiction when you mean to write a disproof.

Now, the longer version. ☺

If you are writing a proof of a result, that result is called a **theorem**. The term “theorem” specifically refers to a statement that is true under a specific set of assumptions. The general template for writing a proof looks like this:

Theorem: [*statement that you want to prove is true*]

Proof: [*some argument establishing why that statement is true*]

On the other hand, let's suppose that you have some statement that is *not* true, and you want to show that that statement is indeed false. This is called a **disproof**. Since you'll be showing that a given statement is not true, it is not appropriate to call that statement a “theorem.” Remember – the term “theorem” specifically refers to a statement that's true! When you're writing a disproof, you'd typically refer to the statement in question as a **claim** (something that's being proposed, but which isn't necessarily true) to indicate that the statement should be treated with some suspicion.

The general template for writing a disproof looks like this:

Claim: [*statement that you want to prove is false*]

Disproof: [*some argument establishing why that statement is false*]

Be very careful not to mix and match the terminology from proofs and disproofs. For example, suppose you want to disprove the claim that if A and B are sets, then $A \cap B = \emptyset$. (Here, this statement is false because it's implicitly a universally-quantified statement, and there indeed exist pairs of sets with a nonempty intersection). Here's how you *shouldn't* do this:

Incorrect! Theorem: If A and B are sets, then $A \cap B = \emptyset$.

Incorrect! Proof: We will show that this statement is not true. Consider the sets $A = \mathbb{N}$ and $B = \mathbb{N}$. Notice that $A \cap B = \mathbb{N} \cap \mathbb{N} = \mathbb{N}$, so $A \cap B \neq \emptyset$. ■

The problem with the above setup is that, to a quick glance, it seems like you're doing exactly the opposite of what you're actually doing. By labeling the statement as a theorem and the argument as a proof, you are signaling to your reader that you think that the statement is true and that you're going to provide a justification for it. If they then read your proof, they're going to be terribly confused, because you're starting your proof off by saying that you're going to show that your theorem – something that's supposed to be true – isn't actually true.

A better way to write this would be to do something like this:

Claim: If A and B are sets, then $A \cap B = \emptyset$.

Disproof: We will show that the negation of this statement is true, namely that there exist sets A and B where $A \cap B \neq \emptyset$.

Consider the sets $A = \mathbb{N}$ and $B = \mathbb{N}$. Notice that $A \cap B = \mathbb{N} \cap \mathbb{N} = \mathbb{N}$, so $A \cap B \neq \emptyset$. ■

Take a look at how this argument is laid out. First, the statement in question is marked as a claim, not a theorem, so someone reading over your work will get cued in that you're simply stating something rather than arguing that it's true. Next, the argument is explicitly labeled as a disproof, indicating to the reader that they're about to see why the claim isn't true. The specifics of that argument then outline a reason why the claim is false – specifically, it says that the negation of the claim is true, then explains why that's the case.

Another common error we see people make when writing out disproofs is to mix up two related but different concepts: disproofs (arguments that show why a claim isn't true) and proofs by contradiction (arguments that show that a claim *is* true by assuming for the sake of argument that it isn't). Although both a disproof and a proof by contradiction will involve working with the negation of a statement, they proceed very differently from one another. In a disproof, you take the negation of the statement in question, then **prove** that the negation is true. In a proof by contradiction, you **assume** that the negation is true, derive a contradiction, and then claim that, as a result, the statement must have been true all along. In other words, a disproof explains why something is **not** true, and a proof by contradiction explains why something **is** true. As a result, you have to be careful not to mix these concepts up.

For example, here's another example of how **not** to write a disproof:

Claim: If A and B are sets, then $A \cap B = \emptyset$.

Incorrect! Disproof: Assume for the sake of contradiction that there exist sets A and B where $A \cap B \neq \emptyset$.

Consider the sets $A = \mathbb{N}$ and $B = \mathbb{N}$. Notice that $A \cap B = \mathbb{N} \cap \mathbb{N} = \mathbb{N}$, so $A \cap B \neq \emptyset$. ■

This disproof says that we should start by assuming that the negation of the claim in question here is true. Remember that the whole point of a disproof is to explicitly **prove** that the negation of the claim is true, so if we start off by **assuming** the negation of the claim, there's nothing left to do!

Why we enforce this rule: This rule is designed to minimize confusion on the part of the person reading your proof. If you are writing a disproof of a result and structure it as though you're writing a proof of a theorem, the person reading your disproof will go in with completely incorrect expectations about what they're going to find. In the best case, a reader will quickly figure this out and begin rereading what you wrote from the top, which isn't the best use of their time. In the worst case, the reader will be totally lost and not understand what it is that you're trying to do. (There's an *even worse* case, and that's where a TA will look at what you wrote, say "well, you got the wrong answer, because you're trying to prove something false" and then give you zero points without reading further, but we'll ignore that for now.)