7 First proofs: Elementary set theory[8]

As mathematicians and students of mathematics, it is our job to separate those mathematical assertions that are true from those that are false. The method that mathematicians use to verify that a mathematical assertion is true is called *deductive proof*. Roughly speaking, a deductive proof is a *step by step* argument that starts from *initial assumptions* and applies *valid rules of deduction* to build to a desired conclusion.

Theorems, propositions, corollaries, lemmas, conjectures, and axioms We have been using the term *principle* to mean a mathematical assertion that has been proved to be true. Mathematicians usually label a principle as a *Theorem*, *Proposition*, *Lemma*, *Corollary* or *Claim*. These terms are not precisely defined but are used in the following way:

- A principle that is considered relatively important is called a *theorem*, otherwise it is called a *proposition*.
- A *corollary* refers to a a principle that is deduced as an easy consequence of a previously proved theorem.
- A *lemma* refers to a principle that may not be so interesting by itself, but is of interest because it is a step in the proof of a theorem or a proposition. A relatively easy lemma may be called a *claim*.

An assertion for which there is no proof may be informally called a *speculative assertion*. If a mathematician believes that a speculative assertion is true based on some evidence (such as many successful test cases, or a partially completed proof), he might propose the assertion as a *conjecture*. There are a number of very interesting conjectures in mathematics that have not been proved yet. Conjectures are sometimes named by the person who proposed them. For example, one of the most famous conjectures in mathematics is *Goldbach's conjecture*:

Every even number greater than 2 can be expressed as the sum of two primes.

It has been verified by computer that every even number between 4 and 4×10^{18} can be expressed as the sum of two primes. However, no one knows whether this is true of all even numbers.

The last word that refers to a universal principle is the word *axiom*. An axiom is a universal principle that we assume without proof. You need to have some universal principles to start with in order to prove anything. The axioms provide a starting point for mathematical reasoning.

One confusing aspect when you start doing proofs is: which universal principles do we treat as axioms (so we assume them without proof) and which ones require proof. This depends upon context. In this chapter we will prove lots of very basic facts including ones that you might think don't need a proof.

When we start doing proofs with numbers, our axioms will include things like the familiar commutative property of addition: For any two real numbers a,b, $a + b = b + a$. We'll list our axioms for workoing with numbers later.

[8]Version: 10-9-16.

First proof: An existential assertion Most theorems that mathematicians prove are universal principles, however sometimes they prove existential assertions. If a conjectured universal principle turns out to be false, then the counterexample yields an existential assertion. For example, suppose we had the following proposed universal principle:

Assertion 7.1. For any three sets A, B and C if $A \cap B \neq \emptyset$, $A \cap C \neq \emptyset$ and $B \cap C \neq \emptyset$ then $A \cap B \cap C \neq \emptyset$.

It's not hard to see that this assertion is false. To prove it false you need to give a counterexample. A counterexample provides a proof of the negation of the above assertion, which is an existential assertion.

Proposition 7.1. *There exist sets A, B and C such that $A \cap B \neq \emptyset$, $A \cap C \neq \emptyset$ and $B \cap C \neq \emptyset$ and $A \cap B \cap C = \emptyset$.*

Proof. Let $A = \{1, 2\}$, $B = \{1, 3\}$ and $C = \{2, 3\}$. Then $A \cap B = \{1\}$, $A \cap C = \{2\}$ and $B \cap C = \{3\}$, while $A \cap B \cap C = \emptyset$. □

An existential assertion specifies a particular scenario, and to prove it we just need to show that there's at least one solution. Usually, this is done as above: we simply describe an example, and demonstrate that the example fits the scenario. However, this is not always easy, for two reasons:

- It may be very hard to come up with an example, or
- even if we have a proposed example, it may be difficult to demonstrate that it satisfies the required properties.

Proving universal assertions: setting up the scenario Most of our efforts will be spent on proving *universal* principles. Such a principle has the form: "For all x of type T, if $A(x)$ then $C(x)$". As we've discussed, universal principles seem much more challenging to prove because the principle summarizes many assertions, one for every x of type T that satisfies $A(x)$. It is not enough to simply check a few examples of x's that satisfy $A(x)$. We need a way to simultaneously argue about all possible x's.

The proof of a universal principle *will always start* something like:

> Suppose x is an arbitrary object of type T. Assume $A(x)$. We must show $C(x)$.

The first two sentences *set up the hypothetical scenario* for the universal principle. Having set up the hypothesis our goal is to show that the objects in the scenario satisfy the conclusion. This scenario and goal comprise the initial *proof task* to be accomplished.

The main part of the proof requires us to argue convincingly that $C(x)$ holds. Our argument may use the information about x given by the scenario (that x of type T and satisfies $A(x)$) but makes no other assumptions about x.

The scenario given by the object x satisfying $A(x)$ is the *starting scenario* of the proof and $C(x)$ is the *starting goal*, and together these comprise the *starting proof task* to be accomplished.

One of the most important features of a proof is that as the proof proceeds, both the scenario and the goal of the proof task *may change.* New objects may be added and later disappear, new assumptions are made, and later dropped, and the goal can change as well. We'll refer to these changes as the *evolution of the proof task.* As the proof task changes, it is absolutely crucial to be aware of the current task:

- What are the active objects for the current scenario?
- What are the active assumptions?
- What is the current goal?

The idea of the evolution of the proof task may seem confusing, but becomes much less confusing once you understand that the evolution proceeds according to strict rules, which you will learn. You will need to master these rules, and also need to practice keeping track of the current proof task. As you develop mastery you'll find that both the rules, and the evolving proof task become natural and you will have the built a foundation that will enable you to study and write proofs in any area of mathematics.

7.1 Proving universal principles in elementary set theory

Proofs of universal principles all have the same basic structure regardless of the area of mathematics. To get started with these proofs we'll work in elementary set theory. We pick this area because:

- There are many examples of universal principles that are relatively easy to prove,
- The proofs illustrate many of the basic techniques of proof.
- The theorems, while easy, are useful in many areas of mathematics.

Before we actually start doing proofs, we need a starting point of mathematical knowledge. This starting point consists of the basic definitions and facts about set theory. We introduced these in Section 2.1. We now carefully review and expand on these definitions in a more careful way.

In most settings, we are considering sets whoses members come from a particular *universe.* We denote the universe set by U. In what follows, S and T are arbitrary subsets of U and $(A_j : j \in J)$ is an indexed collection of sets of subsets from U with index set J.

Set membership For every set S and every object x, either x is a member of S, denoted $x \in S$ or x is not a member of S, denoted $x \notin S$.

Subset and superset For two sets S and T we say that S *is a subset of* T, denoted $S \subseteq T$, or T *is a superset of* S, denoted $T \supseteq S$ provided that for every $x \in S$, we have $x \in T$.

Union The union of S and T, denoted $S \cup T$ is the set whose members are those x such that $x \in S$ or $x \in T$. The union of an indexed collection $(A_j : j \in J)$, denoted $\bigcup_{j \in J} A_j$, is the set consisting of those x such that there exists a $j \in J$ with $x \in A_j$.

Intersection The intersection of S and T, denoted $S \cap T$ is the set whose members are those x such that both $x \in S$ and $x \in T$. The intersection of indexed collection $(A_j : j \in J)$, denoted $\bigcap_{j \in J} A_j$, is the set consisting of those x such that for all $j \in J$, $x \in A_j$.

Set difference The difference between S and T, denoted $S \backslash T$ or $S - T$ is the set of those x that belong to S but not to T.

Complement The complement of A is the set $U \backslash A$. (Notice that the definition of complement assumes that there is a fixed universe U. If we change the universe then the notion of complement may change. For example if A is the set of even integers and the universe is the set of integers then A^c is the set of odd integers, but if the universe is the set of real numbers then A^c is the union of the set of odd integers and the set of real numbers that are not integers.)

Symmetric Difference The symmetric difference of two sets A and B, denoted $A \triangle B$ is the set $A \backslash B \cup B \backslash A$.

Equality of Sets We define what it means for two sets to be equal in two ways. These definitions mean the same thing.

Biconditional version Two sets S and T are equal provided that for every object x, $x \in S$ if and only if $x \in T$.

Containment version Two sets S and T are equal provided that $S \subseteq T$ and $T \subseteq S$.

Disjointness of sets Two sets A and B are *disjoint* provided that $A \cap B = \emptyset$. The collection $(A_j : j \in J)$ is *pairwise disjoint* if for any $i, j \in J$ with $i \neq j$ we have A_i and A_j are disjoint.

Here we are very careful to express these definitions precisely using the logical constructions (and, or, if-then, if and only if, there exists and for all) described in the earlier sections. For example, the condition that S is a subset of T is equivalent to the universal assertion that "for all $x \in S$ we have $x \in T$". (Note this is an indefinite assertion that depends on the free variables S and T, and that x is a dummy variable). Using this precise language sometimes makes the definitions seem more complicated, but is very useful as we start to prove things.

We will now state and prove some very simple universal principles. These principles are so simple as to be almost obvious, and they are easy to prove. This has the advantage that we can focus on the task at hand, learning what is required of an acceptable proof in mathematics.

Proposition 7.2. *For any indexed collection of sets* $(A_j : j \in J)$ *and any* $k \in J$ *we have ,* $A_k \subseteq \bigcup_{j \in J} A_j$.

In words this proposition says: if you take the union of a collection of sets, then each set in the collection is a subset of that union, which certainly seems quite obvious. Let's see what a careful proof of it looks like.

In some of our examples of proofs, particularly early on, we present proofs in two columns. The proof is contained in the left column. The right hand column contains commentary on the proof that explains how the proof is put together, and is intended to help beginner's learn how to write proofs. The commentary is not part of the proof, and when you write your own proofs you should not include a commentary section, unless specifically asked to do so. Also, the large spaces between lines in the proof column are there only to allow the lines of the proof to match up with the corresponding comment.

Proof	**Commentary**
Suppose $(A_j : j \in J)$ is an arbitrary indexed collection of sets and suppose k is an arbitrary member of J. We must show $A_k \subseteq \bigcup_{j\in J} A_j$.	Since we are proving a universal assertion, we begin by setting up the scenario, and stating the goal. The scenario consists of the index set J, the indexed family $(A_j : j \in J)$ and the index $k \in J$. The goal is to show $A_k \subseteq \bigcup_{j\in J} A_j$.
By definition of $\subseteq$, we must show that for every $x \in A_k$, we have $x \in \bigcup_{j\in J} A_j$.	We change the goal by applying the definition of $\subseteq$. Notice that the high-level structure of this new goal is a universal assertion.
So suppose x is an arbitrary member of A_k. We must show that $x \in \bigcup_{j\in J} A_j$.	This is a key point in the proof. Since the goal has been restated as another universal assertion, we introduce a new object into the scenario corresponding to the object in the $\forall$ quantifier. This leads to a new formulation of the goal.
By the definition of $\bigcup$, $x \in \bigcup_{j\in J} A_j$ means that there is an index i such that $x \in A_i$. So we must show that there is an $i \in J$ such that $x \in A_i$.	Here we use the definition of $\bigcup$ to restate the goal once more.
This is true since, by assumption, we have $k \in J$ and $x \in A_k$.	Using the two assumptions $k \in J$ and $x \in A_k$ from the scenario, we have met the final goal.
We therefore have $x \in \bigcup_{j\in J} A_j$ and since x was an arbitrary member of A_k we have $A_k \subseteq \bigcup_{j\in J} A_j$. $\square$	We now work backwards through the previously stated goals to get back to our original goal, completing the proof. The $\square$ symbol is used to indicate the end of a proof.

Remark 7.1. 1. This universal principle and many of the others in this section are easily visualized using the Venn Diagram picture. The picture is certainly helpful, but is unsat-

isfactory as a proof. We are learning the skill of writing proofs that will work in all areas of mathematics, and in most areas of mathematics you can't prove things by picture. Also, pictures are often misleading, and can lead you to think that something is true that isn't.

2. During the course of the proof, the proof task (scenario and goal) are modified several times. Each modification is driven by the current goal. This is what is meant by the *evolution of the proof task*.

3. The proof starts by *focusing on the task of simplifying the goal*. As we simplify the goal, the scenario changes and the goal change. We continued until the goal was so simple that there was no useful way to transform it further. At that point, we start trying to prove the goal. For this particular proof, achieving this final goal was simple. This will be the case in most of the early proofs we do, because we are focusing on mastering the process of modifying the scenario and goal. As we consider more interesting universal principles, we aren't so lucky; after modifying the scenario and the goal there will still be a lot of work to do.

Here's a related statement concerning the intersection of sets in an indexed family.

Proposition 7.3. *For all indexed collections of sets $(A_j : j \in J)$ and for all $k \in J$, $\cap_{j\in J} A_j \subseteq A_k$.*

Exercise 7.1. Prove Proposition 7.3

Corollary 7.4. *For any two sets A and B, $A \subseteq A \cup B$ and $A \cap B \subseteq A$.*

This is called a corollary because it follows easily by applying previously proved assertions.

Proof. Suppose that A and B are arbitrary sets. For the first conclusion, we want to apply Proposition 7.2. To do this we view (A, B) as an ordered pair of sets, so that it is an indexed collection. Then Proposition 7.2 implies $A \subseteq A \cup B$. Similarly Proposition 7.3 implies $A \cap B \subseteq A$. □

The next principle is called *transitivity of containment*.

Proposition 7.5. *For any three sets S,T, and U, if $S \subseteq T$ and $T \subseteq U$ then $S \subseteq U$.*

Proof	**Commentary**
Suppose that S, T and U are arbitrary sets. Assume $S \subseteq T$ and $T \subseteq U$. We must show $S \subseteq U$.	Setting up the initial proof task.
By definition of $\subseteq$, we must show that for all $x \in S$, we have $x \in U$.	Restating the goal using the definition of $\subseteq$.

Suppose that x is an arbitary member of S. We must show $x \in U$.	Since the new goal is a universal assertion, we introduce a new object into the scenario, and modify the goal accordingly.
Since $x \in S$ and $S \subseteq T$ we have $x \in T$. Since $x \in T$ and $T \subseteq U$ we have $x \in U$. Since x was an arbitrary member of S we conclude $S \subseteq U$, as required to prove the proposition. $\square$	Here we apply the definition of subset twice.

The next principle shows that union and intersection satisfy a "distributive law". The proof will introduce an important proof teachnique called *proof by cases*.

Proposition 7.6. *For any indexed collection $(A_j : j \in J)$ of sets and any set B we have:*

1. $(\bigcap_{j\in J} A_j) \cup B = \bigcap_{j\in J}(A_j \cup B)$
2. $(\bigcup_{j\in J} A_j) \cap B = \bigcup_{j\in J}(A_j \cap B)$

The conclusion has two parts. We will only prove the first conclusion, and leave the second conclusion as an exercise.

Proof	**Commentary**
Suppose that $(A_j : j \in J)$ is an indexed collection of sets and suppose that B is a set. We must show: $$(\bigcap_{j\in J} A_j) \cup B = \bigcap_{j\in J}(A_j \cup B).$$	Setting up the proof task.
Using the definition of set equality (containment version), it suffices to show: (1) The first set is a subset of the second, and (2) The second set is a subset of the first.	Here we have replaced our original goal by a combination of two goals. We will achieve each goal separately.
To prove (1), by the definition of $\subseteq$ we must show that for all $x \in (\bigcap_{j\in J} A_j) \cup B$ we have $x \in \bigcap_{j\in J}(A_j \cup B)$. Suppose x is an arbitrary member of $(\bigcap_{j\in J} A_j) \cup B$. We must show $x \in \bigcap_{j\in J}(A_j \cup B)$	Introduce x into the scenario and reformulate the goal.

By definition of $\bigcap$ this means we must show that for every $j \in J$, $x \in A_j \cup B$. Suppose k is an arbitrary member of J; we must show $x \in A_k \cup B$, and by the definition of $\cup$ this means we must show $x \in A_k$ or $x \in B$.

Since by hypothesis $x \in (\bigcap_{j\in J} A_j) \cup B$, using the definition of $\cup$, we have $x \in (\bigcap_{j\in J} A_j)$ or $x \in B$.

We now divide the proof into cases depending on whether $x \in \bigcap_{j\in J} A_j$ or $x \in B$.

Case 1. Assume $x \in \bigcap_{j\in J} A_j$. Then by definition of $\bigcap$, for all $j \in J$ we have $x \in A_j$. In particular, since $k \in J$, $x \in A_k$, and Thus $x \in A_k \cup B$.

Case 2. Assume $x \in B$. Then $x \in A_k \cup B$, by the definition of $\cup$.

In both cases we concluded $x \in A_k \cup B$, and since $k \in J$ was arbitrary we conclude that $x \in \bigcap_{j\in J}(A_k \cup B)$. Since x was an arbitrary member of $(\bigcap_{j\in J} A_k) \cup B$, we conclude that $(\bigcap_{j\in J} A_j) \cup B \subseteq \bigcap_{j\in J}(A_j \cup B)$. This completes the proof of (1).

Next we show (2) that $\bigcap_{j\in J}(A_j \cup B) \subseteq (\bigcap_{j\in J} A_j) \cup B$. For this we must show that for all $x \in \bigcap_{j\in J}(A_j \cup B)$ we have $x \in (\bigcap_{j\in J} A_j) \cup B$. Suppose x is an arbitrary member of $\bigcap_{j\in J}(A_j \cup B)$. We must show $x \in (\bigcap_{j\in J} A_j) \cup B$, so we must show that $x \in \bigcap_{j\in J} A_j$ or $x \in B$.

Applying the definition of $\bigcap$ changes the goal once again to the form of a universal principle. So we introduce an arbitrary $k \in J$ into the scenario and reformulate the goal. Applying the definition of $\cup$ changes the goal again.

Now that the goal is simplied, we focus on the assumptions we know about x to work towards the goal. The hypothesis on x tells us that x satisfies one of two conditions.

This is our first example of *proof by cases*. We have identified two possibilities. Our assumptions tell us that at least one of them is true, but we don't know which. The argument that we'll use to reach our goal depends on which of these possibilities is true. So we split the proof into two *cases*, in which we consider each possibility separately.

In the first case, we add the first possibility "$x \in \bigcap_{j\in J} A_j$" to the active assumptions of the scenario. This is called the *case assumption*. We then show the desired conclusion.

We now do the second case. Since the first case is over, the assumption " $x \in \bigcap_{j\in J} A_j$" is *removed* from the scenario. We then add the new case assumption "$x \in B$" to the scenario.

As we did in part (1), we start by modifying the goal. Because of the "for all" goal, we introduce x into the scenario, and modify the goal further.

We will divide into cases depending on whether $x \in B$ or $x \notin B$.	Again we use proof by cases. Here we know that $x \in B$ and $x \notin B$ cover all possibilities, since one is the negation of the other.
Case 1. Assume $x \in B$. Then $x \in (\bigcap_{j\in J} A_j) \cup B$ by the definition of $\cup$.	This is a very easy case.
Case 2. Assume $x \notin B$. In this case, we will show that $x \in \bigcap_{j\in J} A_j$. To achieve this goal we must show that for all $j \in J$ we have $x \in A_j$. So suppose that k is an arbitrary member of J. We must show $x \in A_k$.	Our goal is to prove $x \in \bigcap_{j\in J} A_j$ or $x \in B$. In this case $x \notin B$, so we'll show $x \in \bigcap_{j\in J} A_j$.
By our assumption on A, $x \in \bigcap_{j\in J}(A_j \cup B)$ so by the definition of $\bigcap$, since $k \in J$, we have $x \in A_k \cup B$, so $x \in A_k$ or $x \in B$. Since $x \notin B$, we must have $x \in A_k$. Since k was an arbitrary member of J we have $x \in \bigcap_{j\in J} A_j$.	Summing up Case 2.
In both cases we showed $x \in (\bigcap_{j\in J} A_j) \cup B$, and since x was an arbitrary member of $\bigcap_{j\in}(A_j \cup B)$ we conclude $\bigcap_{j\in J}(A_j \cup B) \subseteq (\bigcap_{j\in J} A_j) \cup B$, as required to complete part (2). $\square$	

Remark 7.2. 1. In both parts of the proof we split into cases. For a proof by cases to be valid, it is important that the different cases cover all of the possibilities. In the proof of the second part, this is obvious since x is either in B or not in B. In the proof of the first part, the two cases $x \in \bigcap_{j\in J} A_j$ and $x \in B$ cover all possibilities because within our scenario have deduced that $x \in \bigcap_{j\in J} A_j$ or $x \in B$. If the different cases cover all possibilities for the active scenario we say that the set of cases is *exhaustive*

2. In the proof of the second part, the two cases are *mutually exclusive*, which means that they can't both be true at the same time. However, in the proof of the second part, the two cases are not mutually exclusive. This is perfectly okay for a proof by cases. We require the cases to be exhaustive, but we don't require them to be mutually exclusive.

Exercise 7.2. When we reformulated our original goal in this proof, we used the containment version of the set equality definition. Provide an alternative proof using the biconditional version of set equality to write the initial goal.

Exercise 7.3. Prove the second part of Proposition 7.6.

Properties of Set complement Recall that in any given situation we assume that all the sets we're dealing with belong to some universe set U and the complement A^c of a set A is defined to be the set of objects $x \in U$ such that $x \notin A$.

Here we review and prove a few elementary properties of set complement.

Proposition 7.7. *For any universe U and subsets A and B of U we have:*

1. $(A^c)^c = A$
2. $(B \backslash A)^c = B^c \cup A$

Proof	**Commentary**
For the first part, we must prove two things (1a) $(A^c)^c \subseteq A$ and $A \subseteq (A^c)^c$.	
For (1a), suppose x is an arbitrary member of $(A^c)^c$. We must show $x \in A$.	
Assume for contradiction that $x \notin A$. We will derive a contradiction.	This introduces a new proof technique called "Proof by contradiction". Assume that what you are trying to prove is *false*; so you add its negation to the active assumptions. If this leads to an impossible conclusion, then the assumption you added must be impossible, and so what you're trying to show must be true.
Since $x \notin A$, then by definition of A^c, $x \in A^c$. Since $x \in A^c$, $x \notin (A^c)^c$. But this contradicts that $x \in (A^c)^c$. So the assumption $x \notin A$ must be incorrect and so $x \in A$.	We arrive at a contradiction and conclude that the assumption we added must be false, so its negation is true.
Since x was an arbitrary member of $(A^c)^c$ we conclude $(A^c)^c \subseteq A$, as required, completing (1a).	
The proofs of (1b) and of (2) are left as exercises.	

Exercise 7.4. Complete the proof of the first part of Proposition 7.7 by showing (1b) that for any set A, $A \subseteq (A^c)^c$.

Exercise 7.5. Prove the 2nd part of Proposition 7.7.

Proposition 7.8. *For any universe set U and any indexed family $(A_j : j \in J)$ of subsets of U we have:*

1. $(\bigcup_{j\in J} A_j)^c = \bigcap_{j\in J} A_j^c$.
2. $(\bigcap_{j\in J} A_j)^c = \bigcup_{j\in J} A_j^c$.

Proof

Suppose that U is an arbitrary set and $(A_j : j \in J)$ is an arbitrary indexed family of subsets of U. We must show that (1) and (2) are true.

We start with (1). For this we will show (1a) $(\bigcup_{j\in J} A_j)^c \subseteq \bigcap_{j\in J} A_j^c$ and (1b) $\bigcap_{j\in J} A_j^c \subseteq \bigcup_{j\in J} A_j)^c$.

For (1a), suppose x is an arbitrary member of $(\bigcup_{j\in J} A_j)^c$. We must show $x \in \bigcap_{j\in J} A_j^c$. This means that we must show that for all $j \in J$, $x \notin A_j$.

Suppose k is an arbtrary member of J. We must show $x \notin A_k$.

Assume, for contradiction, that $x \in A_k$. We will show that this leads to a contradiction of one of our assumptions.

Since $x \in A_k$ and $k \in J$ we have $x \in \bigcup_{j\in J} A_j$. But we already have that $x \notin \bigcup_{j\in J} A_j$, so we have arrived at a contradiction. Therefore $x \notin A_k$.

Since k was arbitrary member of J, we have $x \in \bigcap_{j\in J} A_j^c$, and since x was an arbitrary member of $(\bigcup_{j\in J} A_j)^c$ we conclude that $\bigcap_{j\in J} A_j^c \subseteq (\bigcup_{j\in J} A_j)^c$. This completes the proof of (1a).

The proof of (1b) is left as an exercise.

Commentary

Splitting the set equality goal into two goals.

As usual for a set containment proof, we introduce an arbitrary member of the set that we want to be contained in the other. We reformulate the goal further based on the definition of $\bigcap$.

Here we use proof by contradiction again.

Arriving at a contradiction we conclude that the assumption we added must be false.

For the proof of (2), we apply the result of (1). We need to show $(\bigcap_{j\in J} A_j)^c = \bigcup_{j\in J} A_j^c$. For each $j \in J$, define B_j to be A_j^c. Now apply (1) to the indexed family $(B_j : j \in J)$. This gives

$(\bigcup_{j\in J} B_j)^c = \bigcap_{j\in J} B_j^c$. Since $B_j = A_j^c$ for each $j \in J$ we can rewrite this as $(\bigcup_{j\in J} A_j^c)^c = \bigcap_{j\in J}(A_j^c)^c$, By Proposition 7.7 we have that for all $j \in J$, $(A_j^c)^c = A_j$, and so $(\bigcup_{j\in J} A_j^c)^c = \bigcap_{j\in J} A_j$, Taking the complement of both sides gives: $(\bigcup_{j\in J} A_j^c)^c)^c = (\bigcap_{j\in J} A_j)^c$. Using Proposition 7.7 again on the lefthand side gives: $(\bigcup_{j\in J} A_j^c) = (\bigcap_{j\in J} A_j)^c$, which is what we needed to prove.

This proof is different from previous ones. Rather than prove the claim by simplifying the goal, we prove it directly by showing that we can apply two previously proved universal assertions. When we apply a universal assertion we can replace the universally quantified objects in the assertion by any objects that meet the hypotheses of the assertion.

Permissible versus Useful We used proof by contradiction again in the first part previous proof. The reader wonder "When is it right to use proof by contradiction?" More generally, as you do more complex proofs, you will find yourself faced with many choices of how to proceed. How do you decide what choice to make?

First off, the choice you make must follow the rules of proof. As long as you make choices in your proofs that follow these rules, your argument is *valid*. This means that you'll never come to an incorrect conclusion. That's crucial; under no circumstances should you ever give an argument that is invalid.

Usually, as you try to construct your proof, there can be many valid choices. In that case, you want to make the choice that helps you reach your goal. The problem is, you won't always know in advance what choices are most useful in reaching your goal. This is where skill and experience comes into doing proofs. The art of doing proofs involves figuring out what choices will lead to the goal. If you apply a valid argument that fails to reach the proof goal, it's not *wrong*, it is *correct* but *unsuccessful*. If you make a correct, but unsuccessful proof attempt, you should try modifying your proof by making different *valid* choice in your proof.

For example, one place where you have a choice is whether to do a *direct proof*, where you work directly to conclude your goal, or *proof by contradiction*. Both are valid choices. If you can't get a successful direct proof, you can try proof by contradiction.

Properties of disjointness Two sets are said to be *disjoint* if they have no elements in common. Many proofs in mathematics boil down to proving that two (possibly very complex) sets are disjoint. Here we present some basic properties of disjointness.

We can restate the property of being disjoint as follows: Sets A and B are disjoint provided that for all $x \in A$ we have $x \notin B$. This definition is symmetric: we can switch the roles of A and B if we want: for all $x \in B$ we have $x \notin A$. This formulation is usually the most convenient

for formulating the goal in a proof that two sets are disjoint.

Proposition 7.9. *For all sets A,B,C and D, If A is disjioint from B and C is disjoint from D then $A \cap B$ is disjoint from $C \cup D$.*

Exercise 7.6. Prove Proposition 7.9.

Properties of set difference

Proposition 7.10. *For any four sets A,B, C, D we have:*

1. *$A \backslash B$ is disjoint from B.*
2. *$A \backslash B \subseteq A$.*
3. *$A \backslash B = A \cap B^c$.*
4. *$(A \backslash B) \backslash C \subseteq A \backslash (B \cup C)$.*
5. *$(A \backslash B) \cap C = (A \cap C) \backslash B = (A \cap C) \backslash (B \cap C)$.*
6. *$(A \backslash B) \cup C = (A \cup C) \backslash (B \cup C)$*
7. *If $A \subseteq C$ then $D \subseteq B$ then $A \backslash B \subseteq C \backslash D$.*

Exercise 7.7. Prove Proposition 7.10

Exercise 7.8. Here are some universal assertions that are similar to those in Proposition 7.10, but are actually false. Find counterexamples to theml.

- $(A \backslash B) \cup C = (A \cup C) \backslash B$.
- If $A \subseteq B$ and $C \subseteq D$ then $A \backslash C \subseteq B \backslash D$.

Properties of symmetric difference The symmetric difference $A \triangle B$ of sets A and B is defined to be the set A
$BcupB$
A. In words, this is the set of objects that belong to exactly one of the sets A and B. Here we do some basic proofs involving symmetric difference.

Proposition 7.11. *For any two sets A and B, we have:*

1. *$A \backslash B \subseteq A \triangle B$ and $B \backslash A \subseteq A \triangle B$.*
2. *$(A \cup B)^c$ and $A \cap B$ are both disjoint from $A \triangle B$.*

Remark 7.3. The second conclusion of the Proposition is an example where the formulation in English does not conform to the strict mathematical rules. In mathematics, the word "and" is used to combine two assertions. Here the word and is between two sets. The sentence is reinterpreted mathematically as "$(A \cup B)^c$ is disjoint from $A \triangle B$ and $A \cap B$ is disjoint from $A \triangle B$."

Proof. Suppose that A and B are arbitrary sets.

For the first part, we must show $A \backslash B \subseteq A \triangle B$ and $B \backslash A \subseteq A \triangle B$. Since $A \triangle B$ is the union of the two sets $A \backslash B$ and $B \backslash A$, corollary 7.4 tells us that $A \triangle B$ is a superset of both of these sets.

For the second part we must show two things (1) $(A \cup B)^c$ is disjoint from $A \triangle B$ and (2) $A \cap B$ is disjoint from $A \triangle B$.

Proof of (1). We need to show that for all $x \in (A \cup B)^c$ we have $x \not\in A \triangle B$. Suppose x is an arbitrary member of $(A \cup B)^c$. Then $x \notin A \cup B$ which means $x \notin A$ and $x \notin B$. Since $x \notin A$, we have $x \notin A \backslash B$ and since $x \notin B$ we have $x \notin B \backslash A$. So $x \notin (A \backslash B) \cup (B \backslash A)$ which means $x \notin A \triangle B$.

Proof of (2). We need to show that for all $x \in A \triangle B$ we have $x \notin A \cap B$. Suppose x is an arbitrary member of $A \triangle B$. We must show $x \notin A \cap B$. By definition of $A \triangle B$, we have $x \in A \backslash B$ or $x \in B \backslash A$. We split into cases depending on whether $x \in A$ or $x \notin A$.

Case 1. Assume $x \in A$. Then $x \notin B \backslash A$, so $x \in A \backslash B$. But then $x \notin B$ so $x \notin A \cap B$.

Case 2. Assume $x \notin A$. Then $x \notin A \cap B$.

In either case, we have $x \notin A \cap B$. Since x was an arbitrary member of $A \triangle B$, we conclude that $A \triangle B$ and $A \cap B$ are disjoint. □

Proposition 7.12. *For any two sets A and B, and any object x, we have:*

1. *If x belongs to exactly one of the sets A and B then $x \in A \triangle B$*
2. *If x belongs to neither of A and B, or to both of A and B, then $x \notin A \Delta B$.*

Proof. Suppose A and B are arbitrary sets, and x an arbitrary object.

For the first part, assume that x belongs to exactly one of the sets A and B. We must show $x \in A \triangle B$. We split into two cases depending on whether $x \in A$ or $x \in B$.

Case 1. Assume $x \in A$. Then $x \notin B$ (since x belongs to exactly one of the two sets). By definition of $A \backslash B$, $x \in A \backslash B$. By definition of $A \triangle B$, $A \triangle B = A \backslash B \cup B \backslash A$, and so by Corollary 7.4 $A \backslash B \subseteq A \triangle B$. Since $x \in A \backslash B$ we then have $x \in A \triangle B$.

Case 2. Assume $x \in B$. Then $x \notin A$ (since x belongs to exactly one of the two sets). By definition of $B \backslash A$, $x \in B \backslash A$. By definition of $A \triangle B$, $A \triangle B = A \backslash B \cup B \backslash A$, and so by Corollary 7.4 $B \backslash A \subseteq A \triangle B$. Since $x \in B \backslash A$ we then have $x \in A \triangle B$.

Since in either case, $x \in A \triangle B$, we conclude that $x \in A \triangle B$.

The proof of the second part is left as an exercise. □

Exercise 7.9. Prove the second part of Proposition 7.12.

Remark 7.4. In the proof of Part 1 of Proposition 7.12 we split into two cases depending on whether $x \in A$ or $x \in B$. The proofs of these cases are nearly identical; you can get the proof of case 2 by switching the roles of A and B in case 1. When this happens, we can use the following shortcut: Instead of writing out case 2 we can simply say: "The proof of case 2 is obtained by switching the roles of A and B in case 1."

Proposition 7.13. *For any sets A, B and C, we have:*

1. $A \triangle \emptyset = A$.
2. $(A \triangle B) \triangle C = A \triangle (B \triangle C)$.
3. $A \cap (B \triangle C) = (A \cap B) \triangle (A \cap C)$.

Proof. We'll only prove the third part; the other two are left as exercises.

Suppose A, B and C are sets. We must show that $A \cap (B \triangle C) = (A \cap B) \triangle (A \cap C)$, so we'll show (1) $A \cap (B \triangle C) \subseteq (A \cap B) \triangle (A \cap C)$ and (2) $(A \cap B) \triangle (A \cap C) \subseteq A \cap (B \triangle C)$.
Proof of (1). Assume $x \in A \cap (B \triangle C)$. We must show $x \in (A \cap B) \triangle (A \cap C)$. Then $x \in A$ and $x \in B \triangle C$. Since $x \in B \triangle C$, we have $x \in B \backslash C$ or $x \in C \backslash B$. We split into two cases depending on which of these is true.

Case 1. Suppose $x \in B \backslash C$. Then $x \in B$ and $x \notin C$. Since $x \in A$ we have $x \in A \cap B$, and since $x \notin C$ we have $x \notin A \cap C$. Therefore $x \in (A \cap B) \backslash (B \cap C)$, as required.

Case 2. Suppose $x \in C \backslash B$. This case is identical to the previous with B and C interchanged.

This completes the proof of (1).

Proof of (2) Assume $x \in (A \cap B) \triangle (A \cap C)$. We must show $x \in A \cap (B \triangle C)$. By definition of $\triangle$ we have $x \in (A \cap B) \backslash (A \cap C)$ or $x \in (A \cap C) \backslash (A \cap B)$. We split into two cases based on these conditions.

Case 2a. Assume $x \in (A \cap B) \backslash (A \cap C)$. Then $x \in A \cap B$ and $x \notin A \cap C$. Since $x \in A \cap B$ we have $x \in B$ and $x \in A$, and therefore since $x \notin A \cap C$, we must have $x \notin C$. Then $x \in B \backslash C \subseteq B \triangle C$, so $x \in A \cap (B \triangle C)$, as required.

Case 2b. Here we assume $x \in (A \cap C) \backslash (A \cap B)$. The proof is identical to Case 2a with B and C interchanged.

□

Exercise 7.10. 1. Prove Part 1 of Proposition 7.13.

2. Prove part 2 of Proposition 7.13.
3. Show that if we replace both occurrence of $\cap$ by $\cup$ in part 2 of Proposition 7.13, the resulting assertion is false.

Properties of the power set Recall that for a set A, $\mathcal{P}(A)$ is the set of all subsets of A.

Proposition 7.14. *For any two sets A and B, if $A \subseteq B$ then $\mathcal{P}(A) \subseteq \mathcal{P}(B)$.*

Proof. Suppose A and B are sets. Assume $A \subseteq B$. We must show $\mathcal{P}(A) \subseteq \mathcal{P}(B)$, which is equivalent to showing that for all $X \in \mathcal{P}(A)$ we must have $X \in \mathcal{P}(B)$. Suppose $X \in \mathcal{P}(A)$. We must show $X \in \mathcal{P}(B)$. Since $X \in \mathcal{P}(A)$, by definition of the power set, $X \subseteq A$. Since $X \subseteq A$ and $A \subseteq B$, then by transitivity of containment $X \subseteq B$, and so by the definition of $\mathcal{P}(B)$ we have $X \in \mathcal{P}(B)$. Since X was an arbitrary member of $\mathcal{P}A$ we conclude $\mathcal{P}(A) \subseteq \mathcal{P}(B)$. □

Remark 7.5. In this proof we used X instead of x as an arbitrary member of $\mathcal{P}(A)$. This is because members of $\mathcal{P}(A)$ are subsets of A, and it is conventional to represent subsets by capital letters.

Exercise 7.11. For each of the following two assertions, determine whether it's true. If it's true, prove it. If it's false, prove that its false.

1. For any two sets A and B, $\mathcal{P}(A \cap B) = \mathcal{P}(A) \cap \mathcal{P}(B)$.
2. For any two sets A and B, $\mathcal{P}(A \cup B) \subseteq \mathcal{P}(A) \cup \mathcal{P}(B)$.
3. For any two sets A and B, $\mathcal{P}(A \cup B) \supset \mathcal{P}(A) \cup \mathcal{P}(B)$.
4. For any two sets A and B, $\mathcal{P}(A) \backslash \mathcal{P}(B) \subseteq \mathcal{P}(A \backslash B)$.
5. For any two sets A and B, $\mathcal{P}(A) \backslash \mathcal{P}(B) \supseteq \mathcal{P}(A \backslash B)$.

Properties of set products Recall that if A and B are sets then $A \times B$ denotes the set of all ordered pairs with first entry in A and second entry in B.

Proposition 7.15. *For all sets A,B,C and D we have:*

1. $A \times (B \cup C) = (A \times B) \cup (A \times C)$.
2. $A \times (B \cap C) = (A \times B) \cap (A \times C)$.
3. $A \times \emptyset = \emptyset$.
4. $(A \times B) \cap (D \times C) = (A \cap D) \times (B \cap C)$.
5. $(A \times B) \cup (D \times C) \subseteq (A \cup D) \times (B \cup C)$.
6. $(A \times B) \cup (D \times C) \supseteq (A \cup D) \times (B \cap C)$.

Remark 7.6. When we prove things about product sets, we will sometimes choose an arbitrary member of $A \times B$. We could call this object x, but since its an ordered pair, it is usually more convenient to call it (a, b). When we say "Suppose (a, b) is an arbitrary member of $A \times B$" we really mean "Suppose a is an arbitrary member of A and b is an arbitrary member of B such that $(a, b) \in A \times B$."

Proof. Suppose that A, B, C and D are arbitrary sets. There are several things to prove. We wll only consider the first, and leave the others as exercises

For the proof of the first part, we must show (i) $A \times (B \cup C) \subseteq (A \times B) \cup (A \times C)$ and also (ii) $(A \times B) \cup (A \times C) \subseteq A \times (B \cup C)$. For part (i) we must show that every member of $A \times (B \cup C)$ is a member of $(A \times B) \cup (A \times C)$. Suppose that (a, b) is an arbitrary member of $A \times (B \cup C)$. We must show $(a, b) \in (A \times B) \cup (A \times C)$. Since $(a, b) \in A \times (B \cup C)$ we have $a \in A$ and $b \in B \cup C$. So we have $b \in B$ or $b \in C$. We consider these possibilities separately.
Case 1. Assume $b \in B$. Then $(a, b) \in A \times B$ and so $(a, b) \in (A \times B) \cup (A \times C)$, as required.
Case 2. Assume $b \in C$. Then $(a, b) \in A \times C$ and so $(a, b) \in (A \times B) \cup (A \times C)$.

Since (a, b) was an arbitrary member of $A \times (B \cup C)$ we conclude that $A \times (B \cup C) \subseteq (A \times B) \cup (A \times C)$ which concludes part (i). Part (ii) is left as an exercise. □

Exercise 7.12. 1. Complete the proof of the first part of Proposition 7.15.

2. Prove parts 2-6 of Proposition 7.15

7.2 Elementary principles for functions

Interesting mathematics starts with interesting questions. In many parts of these notes we'll study some aspect of the mathematical universe, and we'll focus on some basic questions.

In this subsection, we'll be looking at composition of functions. Recall that if $f : B \longrightarrow C$ and $g : A \longrightarrow B$ then $f \circ g$ is the function from A to C given by the rule that for all $a \in A$ $f \circ g(a) = f(g(a))$. Throughout this subsection we'll assume that when we form the composition $f \circ g$ that $\mathbf{Dom}(f) = \mathbf{Target}(g)$, which in this case is the set B.

Here are the questions we want to consider?

- Recall that in arithmetic, the addition and multiplication operations are associative, that is: for any three numbers a, b, c we have $(a+b)+c = a+(b+c)$ and $(a \times b) \times c = a \times (b \times c)$. Is the operation $\circ$ for combining functions associative? The is, if e, f, g are functions with $g : A \longrightarrow B$, $f : B \longrightarrow C$ and $e : C \longrightarrow D$ is it true that $(e \circ f) \circ g = e \circ (f \circ g)$?

- *When can a function be canceled from both sides of an equation?* One of the basic properties of addition of numbers is the cancellation property: For any three numbers a, b and c, if $a + b = a + c$ then $b = c$.

  Suppose three functions $f : B \longrightarrow C$, $g : A \longrightarrow B$ and $h : Alongrightarrow B$ satisfy $f \circ g = f \circ h$. Can we always cancel the f's and conclude $g = h$? If not, what properties of f allow us to do such a cancellation? A function f is called *left-cancellable* if this cancellation is always possible. We can ask a similar question for right cancellation: if $r \circ f = s \circ f$ what must be true about f to conclude that $r = s$?

- *Which functions have an inverse?* Recall that the identify function on set A, $\mathbf{id}_A$ is the function from A to A that maps each $x \in A$ to itself. A *left inverse* for $f : B \longrightarrow C$ is a function $p : C \longrightarrow B$ such that $p \circ f$ is the identify function on $\mathbf{Dom}(f) = B$. Which functions f have a left inverse? Can f have more than one left inverse? If so, which

functions have more than one left inverse and which functions have only one? Similarly, a right inverse for f is a function $q : C \longrightarrow B$ so that $f \circ q$ is the identity on $\mathbf{Target}(f) = C$. When does f have a right inverse? When does f has more than one right inverse? If p is a left inverse of f is it also a right inverse?

The answer to the first question is yes:

Proposition 7.16. *For any sets $A, B.C.D$ amd functions $g : A \longrightarrow B$, $f : B \longrightarrow C$ and $e : C \longrightarrow D$ we have $(e \circ f) \circ g = e \circ (f \circ g)$?*

Proof. Suppose A, B, C, D are arbitrary sets and $g : A \longrightarrow B$, $f : B \longrightarrow C$ and $e : C \longrightarrow D$ are arbitrary functions. We must show $(e \circ f) \circ g = e \circ (f \circ g)$. For this we need to show that they have the same domain and target. and that for every x in the domain the two functions have the same value.

By definition of function composition, $e \circ f$ has domain B and range D, and so $(e \circ f) \circ g$ has domain A and target D. Similarly, $f \circ g$ has domain A and target C and so $e \circ (f \circ g)$ has domain A and target D.

Now suppose that x is an arbitrary member of A. Then $(e \circ f) \circ g(x) = e \circ f(g(x) = e(f(g(x))$ and $e \circ (f \circ g)(x) = e(f \circ g(x)) = e(f(g(x)))$, and so the two functions are equal.

□

To investigate these and similar questions, we start with some definitions.

Definition 7.1. left cancellable A function $f : B \longrightarrow C$ is said to be *left-cancellable* provided that for every pair of functions g and h having the same domain, and having target set $B = \mathbf{Dom}(f)$, if $f \circ g = f \circ h$ then $g = h$.

right cancellable A function $f : B \longrightarrow C$ is said to be *right-cancellable* provided that for every pair of functions g and h both having domain equal to $\mathbf{Target}(f) = C$ and having the same target set T, if $g \circ f = h \circ f$ then $g = h$.

left inverse, left invertible A *left inverse* for function $f : B \longrightarrow C$ is a function $p : C \longrightarrow B$ such that $p \circ f = \mathbf{id}_B$. f is *left invertible* provided that it has at least one left inverse.

right inverse, right invertible A *right inverse* for function $f : B \longrightarrow C$ is a function $q : C \longrightarrow B$ such that $f \circ q = \mathbf{id}_C$. f is *right invertible* provided that it has at least one left inverse.

one-to-one,injective,injection A function f is *one-to-one* (or *injective*) provided that for any $x_1, x_2 \in \mathbf{Dom}(f)$ if $f(x_1) = f(x_2)$ then $x_1 = x_2$. In this case we also say that f is an *injection.*

onto,surjective A function $f : B \longrightarrow C$ is *onto* C (or *surjective*) provided that for all $y \in C$ there is an $x \in \mathbf{Dom}(f)$ such that $f(x) = y$. In this case we say that f is a *surjection.*

bijective, bijection A function $f : B \longrightarrow C$ is *bijective* if it is both one-to-one and onto C. In this case we say that f is a *bijection.*

The following easy proposition says that if you compose a function with the appropriate identify function, you get the function back.

Proposition 7.17. *Suppose $f : B \longrightarrow C$. Then $\mathbf{id}_C \circ f = f$ and $f \circ \mathbf{id}_B = f$.*

Exercise 7.13. Prove Proposition 7.17

Consider the three definitions of left-cancellability, left invertibility and one-to-one. These are all properties that a function might or might not have. The definitions seem very different. The following remarkable theorem tells us that, even though the definitions of these properties are quite different, they are equivalent! This means that if a function satisfies any one of them then it satisfies all three of them. are equivalent

Theorem 7.18. *For any function $f : B \longrightarrow C$ (with B nonempty and C nonempty), the following three properties of f are equivalent:*

1. *f is left-cancellable*
2. *f is left-invertible*
3. *f is one-to-one.*

Remark 7.7. When we want to prove equivalence of two conditions 1 and 2, we need to show that 1 implies 2, and also 2 implies 1. Here we need to show 6 things: 1 implies 2, 2 implies 1, 1 implies 3, 3 implies 1, 2 implies 3 and 3 implies 2. However, there is a short cut: if we order the conditions (in any order we like such as 3,2,1 then it is enough to prove that 1 implies 2, 2 implies 3 and 3 implies 1.

The other three implications follow "for free". For example once we know that 1 implies 2 and 2 implies 3, we can deduce that 1 implies 3.

Proof. Suppose $f : B \longrightarrow C$ with B nonempty. We must show that the three conditions in the conclusion are equivalent. We'll separately show 3 implies 2, 2 implies 1 and 1 implies 3.

Proof that 3 implies 2. Assume f is one-to-one. We must show that f is left-invertible, which means we need to construct a function $q : C \longrightarrow B$ such that $p \circ f = \mathbf{id}_B$. By the definition of $\mathbf{Range}(f)$, for each $r \in \mathbf{Range}(f)$ there is at least one member in $\mathbf{Dom}(f)$ that is mapped by f to r. For each $r \in \mathbf{Range}(f)$, let x_r be a member of $\mathbf{Dom}(f)$ such that $f(x_r) = r$. Let z be a member of B, which must exist since $B \neq \emptyset$. Now define the function $p : B \longrightarrow A$ by the rule $g(b) = x_b$ for $b \in \mathbf{Range}(f)$ and $g(b) = z$ for $b \in B - \mathbf{Range}(f)$. We now show that $q \circ f = \mathbf{id}_B$ which will show that q is a left-inverse for f.

The function $q \circ f$ and $\mathbf{id}_B$ have domain B and target B. We must now show that for all $b \in B$, $q(f(b)) = \mathbf{id}_B(b)$. Suppose b is an arbitrary member of B. Then $\mathbf{id}_B(b) = b$, so we must show $g(f(b)) = b$. Let $c = f(b)$. By the definition of q, since $c \in \mathbf{Range}(f)$, $q(c)$ is equal to x_c which is a member of $\mathbf{Dom}(f)$ that was chosen so that $f(x_c) = c$. We need to show that $x_c = b$. We have $f(b) = c$ and also $f(x_c) = c$. Since f is one-to-one we must have $x_c = b$. Therefore $g(f(b)) = b$, as required. Since b is an arbitrary member of B we conclude that $q \circ f = \mathbf{id}_B$ and so q is a left inverse of f.

Proof that 2 implies 1. Assume f has a left inverse. We must show that f is left-cancellable. To show that f is left-cancellable, we must show that for any two functions g and h having the same domain and having range equal to $\mathbf{Dom}(f) = B$, if $f \circ g = f \circ h$ we have $g = h$. Suppose A is an arbitrary set and $g : A \longrightarrow B$ and $h : A \longrightarrow B$. Assume $f \circ g = f \circ h$. We must show that $g = h$. Since f is left-invertible, we can choose a function $q : C \longrightarrow B$ such that $q \circ f = \mathbf{id}_B$. Since $f \circ g = f \circ h$, we have $q \circ f \circ g = q \circ f \circ h$. The first function is equal to $(q \circ f) \circ g = \mathbf{id}_B \circ g = g$ and the second function is equal to $(q \circ f) \circ h = \mathbf{id}_B \circ h = h$. Therefore $g = h$ as required.

Proof that 1 implies 3. The assertion that $1 \implies 3$ is logically equivalent to $(\neg 3) \implies (\neg 1)$. So we prove that if f is not one-to-one then f is not left-cancellable. Assume that f is not one-to-one. We'll find two different functions g and h with range A such that $f \circ g = f \circ h$, which will prove that f is not left cancellable. Since f is not one-to-one, there are two different domain elements a and a' such that $f(a) = f(a')$. Define g to be the function on domain $\{1\}$ such that $g(1) = a$ and h be the function on domain $\{1\}$ such that $h(1) = a'$. Then $f \circ g$ has domain $\{1\}$ and maps 1 to $f(a)$ and $f \circ h$ has domain $\{1\}$ and maps 1 to $f(a') = f(a)$. So $f \circ g = f \circ h$ but $g \neq h$ as required. □

There is a similar theorem for right-cancellability and right-invertibility.

Theorem 7.19. *For any function $f : B \longrightarrow C$ (with B and C nonempty), the following three properties of f are equivalent:*

1. *f is right-cancellable*
2. *f is right-invertible*
3. *f maps onto B.*

Exercise 7.14. Prove Theorem 7.19

Exercise 7.15.

1. Give an example of a function $f : A \longrightarrow B$ that has a left-inverse but no right-inverse. Show that your example has at least two different left-inverses.
2. Give an example of a function $f : A \longrightarrow B$ that has a right-inverse but no left inverse. Show that your example has at least two different right-inverses.

The previous exercise shows that a function can have a left-inverse without having a right-inverse and that it can have more than one left-inverse and more than one right-inverse.

If the function happens to have both a left inverse and a right inverse, then the picture simplifies. We need some definitions.

Definition 7.2. Invertible function A function that is both left-invertible and right-invertible is said to be *invertible*

Proposition 7.20. *Any invertible function $f : A \longrightarrow B$ has a unique left-inverse and a unique-right inverse, and they are equal to each other.*

Proof. Suppose $f : B \longrightarrow C$. Assume that f has at least one left-inverse and at least one right-inverse. Let $q : C \longrightarrow B$ be an arbitrary left inverse and $p : C \longrightarrow B$ be an arbitrary right inverse of f.

First we show that $q = p$. Since $q : C \longrightarrow B$ and $p : C \longrightarrow B$, we can form the composition $q \circ (f \circ p) : C \longrightarrow B$ and $(q \circ f) \circ p : C \longrightarrow B$, and these are equal by Proposition 7.16. Since p is a right-inverse of f, and q is a left-inverse we have the following chain of equalities:

$$(q \circ f) \circ p = \mathbf{id}_B \circ p = p,$$

and

$$q \circ (f \circ p) = q \circ \mathbf{id}_B = q.$$

Since $(q \circ f) \circ p = q \circ (f \circ p)$ we have $p = q$.

Now we show that if there is a right inverse there can be at most one left inverse. Suppose p is a right inverse and q_1, q_2 are left inverses. We claim $q_1 = q_2$. From the previous part applied to q_1 and p, we have $q_1 = p$ and from the previous part applied to q_2 and p we have $q_2 = p$ so $q_2 = q_1$.

By a similar argument we also have that if there is a left inverse there can be at most one right inverse.

We conclude that if f has both a left-inverse and a right-inverse then it has a unique left-inverse and a unique-right inverse and they must be equal. □

Definition 7.3. inverse function An *inverse* of a function f is a function that is both a left-inverse and a right-inverse of f. By Proposition 7.20, the inverse of an invertible function f is unique. We denote the unique inverse of an invertible function f by f^{-1}.

Theorem 7.21. *Let $f : A \longrightarrow B$. The following conditions are equivalent:*

1. *f is invertible.*
2. *f is bijective*
3. *f has a unique left-inverse*
4. *f has a unique right-inverse*
5. *f is left-cancellable and right-cancellable*

Exercise 7.16. Prove Theorem 7.21

Theorem 7.22. *For any functions $g : A \longrightarrow B$ and $f : B \longrightarrow C$, we have that $f \circ g : A \longrightarrow C$ is well-defined and*

1. *If g_0 is a left-inverse of g and f_0 is a left-inverse of f then $g_0 \circ f_0$ is a well-defined function from C to A, and is a left-inverse of $f \circ g$. Thus if g and f are left-invertible so is $f \circ g$.*
2. *If g_0 is a right-inverse of g and f_0 is a right-inverse of f, then $f_0 \circ g_0$ is a well-defined function from C to A, and is a right-inverse of $f \circ g$. Thus if g and f are right-invertible so is $f \circ g$.*

3. If g is invertible and f is invertible then so is $f \circ g$, and $(f \circ g)^{-1}$ is equal to $g^{-1} \circ f^{-1}$.

Proof. Suppose that $g : A \longrightarrow B$ and $f : B \longrightarrow C$ are functions with target. Then $f \circ g$ is a well-defined function from A to C since for any $a \in A$ $f(a)$ is defined and belongs to B and so $g(f(a)) \in C$.

For the first numbered part, suppose that g_0 is a left-inverse of g and that f_0 is a left-inverse of f. Then $g_0 : B \longrightarrow A$ and $f_0 : C \longrightarrow B$. Since $B = \mathbf{Dom}(g_0)$ is equal to the target of f_0, and $\mathbf{Dom}(f_0) = C$ and the function $g_0 \circ f_0$ is a well-defined function from C to A. We claim that it is a left-inverse of $f \circ g$. Using the fact that composition of functions is associative, we have:

$$(g_0 \circ f_0) \circ (f \circ g) = g_0 \circ ((f_0 \circ f) \circ g) = g_0 \circ (\mathbf{id}_A \circ g) = g_0 \circ g = \mathbf{id}_A,$$

as required.

The second part is similar to the first part and is left as an exercise.

For the third part, assume that g is invertible and f is invertible. Then by definition of invertible, g^{-1} and f^{-1} are left-inverses of g and f, respectively, and so $g^{-1} \circ f^{-1}$ is a left-inverse of $f \circ g$. Also, by definition of invertible, g^{-1} and f^{-1} are right-inverses of g and f, respectively, and so $g^{-1} \circ f^{-1}$ is a right-inverse of $f \circ g$. Since $g^{-1} \circ f^{-1}$ is both a left-inverse and a right-inverse of $f \circ g$, it is an inverse of $f \circ g$ and so $f \circ g$ is invertible and $(f \circ g)^{-1} = g^{-1} \circ f^{-1}$. □

Remark 7.8. It is important to notice that when we take the inverse of $f \circ g$, the inverses of f and g are composed in the opposite order $g^{-1} \circ f^{-1}$. Composing them in the order $f^{-1} \circ g^{-1}$ will not work. In fact $f^{-1} \circ g^{-1}$ is not well-defined in general since since the target A of g^{-1} doesn't agree with the domain C of f^{-1}.

Theorem 7.22 can be used to obtain the following:

Corollary 7.23. *For any functions $g : A \longrightarrow B$ and $f : B \longrightarrow C$,*

1. *If g is one-to-one and f is one-to-one then so is $f \circ g$.*
2. *If g is onto and f is onto then so is $f \circ g$.*
3. *If g is a bijection and f is a bijection then so is $f \circ g$.*

Proof. Suppose $g : A \longrightarrow B$ and $f : B \longrightarrow C$.

For the first part, assume that g is one-to-one and f is one-to-one. By Theorem 7.18 g and f are both left-invertible, and so by the first part of Theorem 7.22 $f \circ g$ is left-invertible and so by Theorem 7.18, $f \circ g$ is one-to-one.

The second and third parts are similar, and are left as exercises. □

Exercise 7.17. Prove the second and third parts of Corollary 7.23

Corollary 7.23 can be proved without using Theorems 7.22, 7.18, and 7.19. Instead you can directly apply the definition of one-to-one and onto. We'll show how to do this for part 2.

Proof. (Alternative proof of Corollary 7.23, part 2.) Suppose that $g : A \longrightarrow B$ and $f : B \longrightarrow C$ are arbitrary functions. Assume that g is onto and that f is onto. We must show that $f \circ g$ is onto, which means that we must show that for every $c \in C$ there is an $a \in A$ such that $f \circ g(a) = c$. Suppose that c is an arbitrary member of C. We must show that there is an $a \in A$ such that $f \circ g(a) = c$. To do this, we need to describe a way to obtain such an a. Since f is onto, there is a $b \in B$ such that $f(b) = c$. Let b be such a member of B. Since g is onto there is an $a \in A$ such that $g(a) = b$. Let a be such a member of A.

We claim that $f \circ g(a) = c$. We have $f \circ g(a) = f(g(a))$ which equals $f(b)$ by the choice of a, and $f(b) = c$ by the choice of c. Thus $f \circ g(a) = c$, as required.

Since $c \in C$ was arbitrary, and we found $a \in A$ so that $f \circ g(a) = c$ □

Remark 7.9. This is our first example of an extremely important type of proof. When we try to prove that $f \circ g$ (or any function) is onto, we need to show that "For any $c \in C$, there is an $a \in A$ such that $f \circ g(a) = c$. " This is an example of a $\forall\exists$ type of assertion.

After we introduced c our goal was to show that there is an $a \in A$ such that $f \circ g(a) = c$. The choice of a depends on various things in the scenario: c, f and g. In order to prove this, we have to do two things: provide specific instructions for finding such an a, and then verify that the a we found has the required properties.

Notice that the introduction of a to the scenario is very different from the way c was introduced. We introduced c because we had a goal which was a universal assertion. So we introduced c to stand for an arbitrary member of C.

When we introduced a our goal was an existential assertion, and we need a to satisfy some specific properties. So we don't want a to be an arbitrary member of A, but rather a member of a carefully chosen to satisfy the properties we needed. Once we gave instructions for a we introduced it to the scenario with the sentence "Let a be such a member of a".

Assertions with $\forall\exists$ structure are extremely important in mathematics, and the proofs typically follow the above pattern. We'll see this type of proof many times throughout these notes.

Exercise 7.18. Give a different proof of Corollary 7.23, part 1 that does not use Theorems 7.22, 7.18, and 7.19 but instead apply the definition of one-to-one.