Chapter A

Appendix: the basics of mathematical logic

The purpose of this appendix is to give a quick introduction to *mathematical logic*, which is the language one uses to conduct rigorous mathematical proofs. Knowing how mathematical logic works is also very helpful for understanding the mathematical way of thinking, which once mastered allows you to approach mathematical concepts and problems in a clear and confident way - including many of the proof-type questions in this text.

Writing logically is a very useful skill. It is somewhat related to, but not the same as, writing clearly, or efficiently, or convincingly, or informatively; ideally one would want to do all of these at once, but sometimes one has to make compromises, though with practice you'll be able to achieve more of your writing objectives concurrently. Thus a logical argument may sometimes look unwieldy, excessively complicated, or otherwise appear unconvincing. The big advantage of writing logically, however, is that one can be absolutely sure that your conclusion will be correct, as long as all your hypotheses were correct and your steps were logical; using other styles of writing one can be reasonably convinced that something is true, but there is a difference between being convinced and being *sure*.

Being logical is not the only desirable trait in writing, and in fact sometimes it gets in the way; mathematicians for instance often resort to short informal arguments which are not logically rigorous when they want to convince other mathematicians of a statement without going through all of the long details, and the same is true of course for non-mathematicians as well. So saying that a statement or argument is "not logical" is not necessarily a bad thing; there are often many situations when one has good reasons not to be emphatic about being logical. However, one should be aware of the distinction between logical reasoning

and more informal means of argument, and not try to pass off an illogical argument as being logically rigorous. In particular, if an exercise is asking for a proof, then it is expecting you to be logical in your answer.

Logic is a skill that needs to be learnt like any other, but this skill is also innate to all of you - indeed, you probably use the laws of logic unconsciously in your everyday speech and in your own internal (non-mathematical) reasoning. However, it does take a bit of training and practice to recognize this innate skill and to apply it to abstract situations such as those encountered in mathematical proofs. Because logic is innate, the laws of logic that you learn should *make sense* - if you find yourself having to memorize one of the principles or laws of logic here, without feeling a mental "click" or comprehending why that law should work, then you will probably *not* be able to use that law of logic correctly and effectively in practice. So, *please* don't study this appendix the way you might cram before a final - that is going to be useless. Instead, **put away your highlighter pen**, and *read* and *understand* this appendix rather than merely *studying* it!

A.1 Mathematical statements

Any mathematical argument proceeds in a sequence of *mathematical statements*. These are precise statements concerning various mathematical objects (numbers, vectors, functions, etc.) and relations between them (addition, equality, differentiation, etc.). These objects can either be constants or variables; more on this later. Statements[1] are either true or false.

Example A.1.1. $2+2=4$ is a true statement; $2+2=5$ is a false statement.

Not every combination of mathematical symbols is a statement. For instance,

$$= 2 + +4 = - = 2$$

is not a statement; we sometimes call it *ill-formed* or *ill-defined*. The statements in the previous example are *well-formed* or *well-defined*. Thus well-formed statements can be either true or false; ill-formed statements are considered to be neither true nor false (in fact, they are usually not considered statements at all). A more subtle example of an

[1]More precisely, statements with no free variables are either true or false. We shall discuss free variables later on in this appendix.

ill-formed statement is

$$0/0 = 1;$$

division by zero is undefined, and so the above statement is ill-formed. A logical argument should not contain any ill-formed statements, thus for instance if an argument uses a statement such as $x/y = z$, it needs to first ensure that y is not equal to zero. Many purported proofs of "0=1" or other false statements rely on overlooking this "statements must be well-formed" criterion.

Many of you have probably written ill-formed or otherwise inaccurate statements in your mathematical work, while intending to mean some other, well-formed and accurate statement. To a certain extent this is permissible - it is similar to misspelling some words in a sentence, or using a slightly inaccurate or ungrammatical word in place of a correct one ("She ran good" instead of "She ran well"). In many cases, the reader (or grader) can detect this mis-step and correct for it. However, it looks unprofessional and suggests that you may not know what you are talking about. And if indeed you actually do not know what you are talking about, and are applying mathematical or logical rules blindly, then writing an ill-formed statement can quickly confuse you into writing more and more nonsense - usually of the sort which receives no credit in grading. So it is important, especially when just learning a subject, to take care in keeping statements well-formed and precise. Once you have more skill and confidence, of course you can afford once again to speak loosely, because you will know what you are doing and won't be in as much danger of veering off into nonsense.

One of the basic axioms of mathematical logic is that every well-formed statement is either true or false, but not both. (Though if there are free variables, the truth of a statement may depend on the values of these variables. More on this later.) Furthermore, the truth or falsity of a statement is intrinsic to the statement, and does not depend on the opinion of the person viewing the statement (as long as all the definitions and notations are agreed upon, of course). So to prove that a statement is true, it suffices to show that it is not false, while to show that a statement is false, it suffices to show that it is not true; this is the principle underlying the powerful technique of *proof by contradiction*, which we discuss later. This axiom is viable as long as one is working with precise concepts, for which the truth or falsity can be determined (at least in principle) in an objective and consistent manner. However, if one is working in very non-mathematical situations, then this axiom becomes

much more dubious, and so it can be a mistake to apply mathematical logic to non-mathematical situations. (For instance, a statement such as "this rock weighs 52 pounds" is reasonably precise and objective, and so it is fairly safe to use mathematical reasoning to manipulate it, whereas vague statements such as "this rock is heavy", "this piece of music is beautiful" or "God exists" are much more problematic. So while mathematical logic is a very useful and powerful tool, it still does have some limitations of applicability.) One can still attempt to apply logic (or principles similar to logic) in these cases (for instance, by creating a *mathematical model* of a real-life phenomenon), but this is now science or philosophy, not mathematics, and we will not discuss it further here.

Remark A.1.2. There are other models of logic which attempt to deal with statements that are not definitely true or definitely false, such as modal logic, intuitionist logic, or fuzzy logic, but these are well beyond the scope of this text.

Being true is different from being *useful* or *efficient*. For instance, the statement

$$2 = 2$$

is true but unlikely to be very useful. The statement

$$4 \leq 4$$

is also true, but not very efficient (the statement $4 = 4$ is more precise). It may also be that a statement may be false yet still be useful, for instance

$$\pi = 22/7$$

is false, but is still useful as a first approximation. In mathematical reasoning, we only concern ourselves with truth rather than usefulness or efficiency; the reason is that truth is objective (everybody can agree on it) and we can deduce true statements from precise rules, whereas usefulness and efficiency are to some extent matters of opinion, and do not follow precise rules. Also, even if some of the individual steps in an argument may not seem very useful or efficient, it is still possible (indeed, quite common) for the final conclusion to be quite non-trivial (i.e., not obviously true) and useful.

Statements are different from *expressions*. Statements are true or false; expressions are a sequence of mathematical symbols which produces some mathematical object (a number, matrix, function, set, etc.)

as its value. For instance

$$2 + 3 * 5$$

is an expression, not a statement; it produces a number as its value. Meanwhile,

$$2 + 3 * 5 = 17$$

is a statement, not an expression. Thus it does not make any sense to ask whether $2 + 3 * 5$ is true or false. As with statements, expressions can be well-defined or ill-defined; $2 + 3/0$, for instance, is ill-defined. More subtle examples of ill-defined expressions arise when, for instance, attempting to add a vector to a matrix, or evaluating a function outside of its domain, e.g., $\sin^{-1}(2)$.

One can make statements out of expressions by using *relations* such as $=$, $<$, $\geq$, $\in$, $\subset$, etc. or by using *properties* (such as "is prime", "is continuous", "is invertible", etc.) For instance, "30 + 5 is prime" is a statement, as is "$30 + 5 \leq 42 - 7$". Note that mathematical statements are allowed to contain English words.

One can make a *compound statement* from more primitive statements by using *logical connectives* such as and, or, not, if-then, if-and-only-if, and so forth. We give some examples below, in decreasing order of intuitiveness.

Conjunction. If X is a statement and Y is a statement, the statement "X and Y" is true if X and Y are both true, and is false otherwise. For instance, "$2 + 2 = 4$ and $3 + 3 = 6$" is true, while "$2 + 2 = 4$ and $3 + 3 = 5$" is not. Another example: "$2 + 2 = 4$ and $2 + 2 = 4$" is true, even if it is a bit redundant; logic is concerned with truth, not efficiency.

Due to the expressiveness of the English language, one can reword the statement "X and Y" in many ways, e.g., "X and also Y", or "Both X and Y are true", etc. Interestingly, the statement "X, but Y" is logically the same statement as "X and Y", but they have different connotations (both statements affirm that X and Y are both true, but the first version suggests that X and Y are in contrast to each other, while the second version suggests that X and Y support each other). Again, logic is about truth, not about connotations or suggestions.

Disjunction. If X is a statement and Y is a statement, the statement "X or Y" is true if either X or Y is true, or both. For instance, "$2+2 = 4$ or $3+3 = 5$" is true, but "$2+2 = 5$ or $3+3 = 5$" is not. Also "$2+2 = 4$ or $3+3 = 6$" is true (even if it is a bit inefficient; it would be a stronger statement to say "$2 + 2 = 4$ and $3 + 3 = 6$"). Thus by default, the word

"or" in mathematical logic defaults to *inclusive or*. The reason we do this is that with inclusive or, to verify "X or Y", it suffices to verify that just one of X or Y is true; we don't need to show that the other one is false. So we know, for instance, that "$2+2 = 4$ or $2353+5931 = 7284$" is true without having to look at the second equation. As in the previous discussion, the statement "$2 + 2 = 4$ or $2 + 2 = 4$" is true, even if it is highly inefficient.

If one really does want to use exclusive or, use a statement such as "Either X or Y is true, but not both" or "Exactly one of X or Y is true". Exclusive or does come up in mathematics, but nowhere near as often as inclusive or.

Negation. The statement "X is not true" or "X is false", or "It is not the case that X", is called the *negation* of X, and is true if and only if X is false, and is false if and only if X is true. For instance, the statement "It is not the case that $2 + 2 = 5$" is a true statement. Of course we could abbreviate this statement to "$2 + 2 \neq 5$".

Negations convert "and" into "or". For instance, the negation of "Jane Doe has black hair and Jane Doe has blue eyes" is "Jane Doe doesn't have black hair or doesn't have blue eyes", *not* "Jane Doe doesn't have black hair and doesn't have blue eyes" (can you see why?). Similarly, if x is an integer, the negation of "x is even and non-negative" is "x is odd or negative", not "x is odd and negative". (Note how it is important here that or is inclusive rather than exclusive.) Or the negation of "$x \geq 2$ and $x \leq 6$" (i.e., "$2 \leq x \leq 6$") is "$x < 2$ or $x > 6$", not "$x < 2$ and $x > 6$" or "$2 < x > 6$.".

Similarly, negations convert "or" into "and". The negation of "John Doe has brown hair or black hair" is "John Doe does not have brown hair and does not have black hair", or equivalently "John Doe has neither brown nor black hair". If x is a real number, the negation of "$x \geq 1$ or $x \leq -1$" is "$x < 1$ and $x > -1$" (i.e., $-1 < x < 1$).

It is quite possible that a negation of a statement will produce a statement which could not possibly be true. For instance, if x is an integer, the negation of "x is either even or odd" is "x is neither even nor odd", which cannot possibly be true. Remember, though, that even if a statement is false, it is still a statement, and it is definitely possible to arrive at a true statement using an argument which at times involves false statements. (Proofs by contradiction, for instance, fall into this category. Another example is proof by dividing into cases. If one divides into three mutually exclusive cases, Case 1, Case 2, and Case 3, then at any given time two of the cases will be false and only one will be true, however this does not necessarily mean that the proof as a whole is incorrect or that the conclusion is false.)

Negations are sometimes unintuitive to work with, especially if there are multiple negations; a statement such as "It is not the case that either x is not odd, or x is not larger than or equal to 3, but not both" is not particularly pleasant to use. Fortunately, one rarely has to work with more than one or two negations at a time, since often negations cancel each other. For instance, the negation of "X is not true" is just "X is true", or more succinctly just "X". Of course one should be careful when negating more complicated expressions because of the switching of "and" and "or", and similar issues.

If and only if (iff). If X is a statement, and Y is a statement, we say that "X is true if and only if Y is true", whenever X is true, Y has to be also, and whenever Y is true, X has to be also (i.e., X and Y are "equally true"). Other ways of saying the same thing are "X and Y are logically equivalent statements", or "X is true iff Y is true", or "$X \leftrightarrow Y$". Thus for instance, if x is a real number, then the statement "$x = 3$ if and only if $2x = 6$" is true: this means that whenever $x = 3$ is true, then $2x = 6$ is true, and whenever $2x = 6$ is true, then $x = 3$ is true. On the other hand, the statement "$x = 3$ if and only if $x^2 = 9$" is false; while it is true that whenever $x = 3$ is true, $x^2 = 9$ is also true, it is not the case that whenever $x^2 = 9$ is true, that $x = 3$ is also automatically true (think of what happens when $x = -3$).

Statements that are equally true, are also equally false: if X and Y are logically equivalent, and X is false, then Y has to be false also (because if Y were true, then X would also have to be true). Conversely, any two statements which are equally false will also be logically equivalent. Thus for instance $2 + 2 = 5$ if and only if $4 + 4 = 10$.

Sometimes it is of interest to show that more than two statements are logically equivalent; for instance, one might want to assert that three statements X, Y, and Z are all logically equivalent. This means whenever one of the statements is true, then all of the statements are true; and it also means that if one of the statements is false, then all of the statements are false. This may seem like a lot of logical implications to prove, but in practice, once one demonstrates enough logical implications between X, Y, and Z, one can often conclude all the others and conclude that they are all logicallly equivalent. See for instance Exercises A.1.5, A.1.6.

— Exercises —

Exercise A.1.1. What is the negation of the statement "either X is true, or Y is true, but not both"?

Exercise A.1.2. What is the negation of the statement "X is true if and only if Y is true"? (There may be multiple ways to phrase this negation).

Exercise A.1.3. Suppose that you have shown that whenever X is true, then Y is true, and whenever X is false, then Y is false. Have you now demonstrated that X and Y are logically equivalent? Explain.

Exercise A.1.4. Suppose that you have shown that whenever X is true, then Y is true, and whenever Y is false, then X is false. Have you now demonstrated that X is true if and only if Y is true? Explain.

Exercise A.1.5. Suppose you know that X is true if and only if Y is true, and you know that Y is true if and only if Z is true. Is this enough to show that X, Y, Z are all logically equivalent? Explain.

Exercise A.1.6. Suppose you know that whenever X is true, then Y is true; that whenever Y is true, then Z is true; and whenever Z is true, then X is true. Is this enough to show that X, Y, Z are all logically equivalent? Explain.

A.2 Implication

Now we come to the least intuitive of the commonly used logical connectives - implication. If X is a statement, and Y is a statement, then "if X, then Y" is the implication from X to Y; it is also written "when X is true, Y is true", or "X implies Y" or "Y is true when X is" or "X is true only if Y is true" (this last one takes a bit of mental effort to see). What this statement "if X, then Y" means depends on whether X is true or false. If X is true, then "if X, then Y" is true when Y is true, and false when Y is false. If however X is false, then "if X, then Y" is *always* true, regardless of whether Y is true or false! To put it another way, when X is true, the statement "if X, then Y" implies that Y is true. But when X is false, the statement "if X, then Y" offers no information about whether Y is true or not; the statement is true, but *vacuous* (i.e., does not convey any new information beyond the fact that the hypothesis is false).

Examples A.2.1. If x is an integer, then the statement "If $x = 2$, then $x^2 = 4$" is true, regardless of whether x is actually equal to 2 or not (though this statement is only likely to be useful when x *is* equal to 2). This statement does not assert that x is equal to 2, and does not assert that x^2 is equal to 4, but it does assert that when and if x is equal to 2, then x^2 is equal to 4. If x is not equal to 2, the statement is still true but offers no conclusion on x or x^2.

Some special cases of the above implication: the implication "If $2 = 2$, then $2^2 = 4$" is true (true implies true). The implication "If $3 = 2$, then $3^2 = 4$" is true (false implies false). The implication "If $-2 = 2$, then $(-2)^2 = 4$" is true (false implies true). The latter two implications are considered vacuous - they do not offer any new information since their hypothesis is false. (Nevertheless, it is still possible to employ vacuous implications to good effect in a proof - a vacously true statement is still true. We shall see one such example shortly.)

As we see, the falsity of the hypothesis does not destroy the truth of an implication, in fact it is just the opposite! (When a hypothesis is false, the implication is automatically true.) The only way to disprove an implication is to show that the hypothesis is true while the conclusion is false. Thus "If $2 + 2 = 4$, then $4 + 4 = 2$" is a false implication. (True does not imply false.)

One can also think of the statement "if X, then Y" as "Y is *at least as true as* X" - if X is true, then Y also has to be true, but if X is false, Y could be as false as X, but it could also be true. This should be compared with "X if and only if Y", which asserts that X and Y are *equally true*.

Vacuously true implications are often used in ordinary speech, sometimes without knowing that the implication is vacuous. A somewhat frivolous example is "If wishes were wings, then pigs would fly". (The statement "hell freezes over" is also a popular choice for a false hypothesis.) A more serious one is "If John had left work at 5pm, then he would be here by now." This kind of statement is often used in a situation in which the conclusion and hypothesis are both false; but the implication is still true regardless. This statement, by the way, can be used to illustrate the technique of proof by contradiction: if you believe that "If John had left work at 5pm, then he would be here by now", and you also know that "John is not here by now", then you can conclude that "John did not leave work at 5pm", because John leaving work at 5pm would lead to a contradiction. Note how a vacuous implication can be used to derive a useful truth.

To summarize, implications are sometimes vacuous, but this is not actually a problem in logic, since these implications are still true, and vacuous implications can still be useful in logical arguments. In particular, one can safely use statements like "If X, then Y" without necessarily having to worry about whether the hypothesis X is actually true or not (i.e., whether the implication is vacuous or not).

Implications can also be true even when there is no causal link between the hypothesis and conclusion. The statement "If $1 + 1 = 2$, then Washington D.C. is the capital of the United States" is true (true implies true), although rather odd; the statement "If $2 + 2 = 3$, then New York is the capital of the United States" is similarly true (false implies false). Of course, such a statement may be unstable (the capital of the United States may one day change, while $1 + 1$ will always remain equal to 2) but it is true, at least for the moment. While it is possible to use acausal implications in a logical argument, it is not recommended as it can cause unneeded confusion. (Thus, for instance, while it is true that a false statement can be used to imply any other statement, true or false, doing so arbitrarily would probably not be helpful to the reader.)

To prove an implication "If X, then Y", the usual way to do this is to first assume that X is true, and use this (together with whatever other facts and hypotheses you have) to deduce Y. This is still a valid procedure even if X later turns out to be false; the implication does not guarantee anything about the truth of X, and only guarantees the truth of Y conditionally on X first being true. For instance, the following is a valid proof of a true proposition, even though both hypothesis and conclusion of the proposition are false:

Proposition A.2.2. *If* $2 + 2 = 5$, *then* $4 = 10 - 4$.

Proof. Assume $2+2 = 5$. Multiplying both sides by 2, we obtain $4+4 = 10$. Subtracting 4 from both sides, we obtain $4 = 10 - 4$ as desired. □

On the other hand, a common error is to prove an implication by first assuming the *conclusion* and then arriving at the hypothesis. For instance, the following Proposition is correct, but the proof is not:

Proposition A.2.3. *Suppose that* $2x + 3 = 7$. *Show that* $x = 2$.

Proof. (Incorrect) $x = 2$; so $2x = 4$; so $2x + 3 = 7$. □

When doing proofs, it is important that you are able to distinguish the hypothesis from the conclusion; there is a danger of getting hopelessly confused if this distinction is not clear.

Here is a short proof which uses implications which are possibly vacuous.

Theorem A.2.4. *Suppose that n is an integer. Then $n(n+1)$ is an even integer.*

Proof. Since n is an integer, n is even or odd. If n is even, then $n(n+1)$ is also even, since any multiple of an even number is even. If n is odd, then $n+1$ is even, which again implies that $n(n+1)$ is even. Thus in either case $n(n+1)$ is even, and we are done. □

Note that this proof relied on two implications: "if n is even, then $n(n+1)$ is even", and "if n is odd, then $n(n+1)$ is even". Since n cannot be both odd and even, at least one of these implications has a false hypothesis and is therefore vacuous. Nevertheless, both these implications are true, and one needs *both* of them in order to prove the theorem, because we don't know in advance whether n is even or odd. And even if we did, it might not be worth the trouble to check it. For instance, as a special case of this theorem we immediately know

Corollary A.2.5. *Let $n = (253+142)*123-(423+198)^{342}+538-213$. Then $n(n+1)$ is an even integer.*

In this particular case, one can work out exactly which parity n is - even or odd - and then use only one of the two implications in the above Theorem, discarding the vacuous one. This may seem like it is more efficient, but it is a false economy, because one then has to determine what parity n is, and this requires a bit of effort - more effort than it would take if we had just left both implications, including the vacuous one, in the argument. So, somewhat paradoxically, the inclusion of vacuous, false, or otherwise "useless" statements in an argument can actually save you effort in the long run! (I'm not suggesting, of course, that you ought to pack your proofs with lots of time-wasting and irrelevant statements; all I'm saying here is that you need not be unduly concerned that some hypotheses in your argument might not be correct, as long as your argument is still structured to give the correct conclusion regardless of whether those hypotheses were true or false.)

The statement "If X, then Y" is not the same as "If Y, then X": for instance, while "If $x = 2$, then $x^2 = 4$" is true, "If $x^2 = 4$, then $x = 2$" can be false if x is equal to -2. These two statements are called *converses* of each other; thus the converse of a true implication is not necessarily another true implication. We use the statement "X if and only if Y" to denote the statement that "If X, then Y; and if Y, then X". Thus for instance, we can say that $x = 2$ if and only if $2x = 4$,

because if $x = 2$ then $2x = 4$, while if $2x = 4$ then $x = 2$. One way of thinking about an if-and-only-if statement is to view "X if and only if Y" as saying that X is just as true as Y; if one is true then so is the other, and if one is false, then so is the other. For instance, the statement "If $3 = 2$, then $6 = 4$" is true, since both hypothesis and conclusion are false. (Under this view, "If X, then Y" can be viewed as a statement that Y is at least as true as X.) Thus one could say "X and Y are equally true" instead of "X if and only if Y".

Similarly, the statement "If X is true, then Y is true" is **not** the same as "If X is false, then Y is false". Saying that "if $x = 2$, then $x^2 = 4$" does not imply that "if $x \neq 2$, then $x^2 \neq 4$", and indeed we have $x = -2$ as a counterexample in this case. If-then statements are not the same as if-and-only-if statements. (If we knew that "X is true if and only if Y is true", then we would also know that "X is false if and only if Y is false".) The statement "If X is false, then Y is false" is sometimes called the *inverse* of "If X is true, then Y is true"; thus the inverse of a true implication is not necessarily a true implication.

If you know that "If X is true, then Y is true", then it is also true that "If Y is false, then X is false" (because if Y is false, then X can't be true, since that would imply Y is true, a contradiction). For instance, if we knew that "If $x = 2$, then $x^2 = 4$", then we also know that "If $x^2 \neq 4$, then $x \neq 2$". Or if we knew "If John had left work at 5pm, he would be here by now", then we also know "If John isn't here now, then he could not have left work at 5pm". The statement "If Y is false, then X is false" is known as the *contrapositive* of "If X, then Y" and both statements are equally true.

In particular, if you know that X implies something which is known to be false, then X itself must be false. This is the idea behind *proof by contradiction* or *reductio ad absurdum*: to show something must be false, assume first that it is true, and show that this implies something which you know to be false (e.g., that a statement is simultaneously true and not true). For instance:

Proposition A.2.6. *Suppose that x be a positive number such that $\sin(x) = 1$. Then $x \geq \pi/2$.*

Proof. Suppose for sake of contradiction that $x < \pi/2$. Since x is positive, we thus have $0 < x < \pi/2$. Since $\sin(x)$ is increasing for $0 < x < \pi/2$, and $\sin(0) = 0$ and $\sin(\pi/2) = 1$, we thus have $0 < \sin(x) < 1$. But this contradicts the hypothesis that $\sin(x) = 1$. Hence $x \geq \pi/2$. □

Note that one feature of proof by contradiction is that at some point in the proof you assume a hypothesis (in this case, that $x < \pi/2$) which later turns out to be false. Note however that this does not alter the fact that the argument remains valid, and that the conclusion is true; this is because the ultimate conclusion does not rely on that hypothesis being true (indeed, it relies instead on it being false!).

Proof by contradiction is particularly useful for showing "negative" statements - that X is false, that a is not equal to b, that kind of thing. But the line between positive and negative statements is sort of blurry. (Is the statement $x \geq 2$ a positive or negative statement? What about its negation, that $x < 2$?) So this is not a hard and fast rule.

Logicians often use special symbols to denote logical connectives; for instance "X implies Y" can be written "$X \implies Y$", "X is not true" can be written "$\sim X$", "$!X$", or "$\neg X$", "X and Y" can be written "$X \wedge Y$" or "$X \& Y$", and so forth. But for general-purpose mathematics, these symbols are not often used; English words are often more readable, and don't take up much more space. Also, using these symbols tends to blur the line between expressions and statements; it's not as easy to understand "$((x = 3) \wedge (y = 5)) \implies (x + y = 8)$" as "If $x = 3$ and $y = 5$, then $x + y = 8$". So in general I would not recommend using these symbols (except possibly for $\implies$, which is a very intuitive symbol).

A.3 The structure of proofs

To prove a statement, one often starts by assuming the hypothesis and working one's way toward a conclusion; this is the *direct* approach to proving a statement. Such a proof might look something like the following:

Proposition A.3.1. *A implies B.*

Proof. Assume A is true. Since A is true, C is true. Since C is true, D is true. Since D is true, B is true, as desired. $\square$

An example of such a direct approach is

Proposition A.3.2. *If $x = \pi$, then $\sin(x/2) + 1 = 2$.*

Proof. Let $x = \pi$. Since $x = \pi$, we have $x/2 = \pi/2$. Since $x/2 = \pi/2$ we have $\sin(x/2) = 1$. Since $\sin(x/2) = 1$, we have $\sin(x/2) + 1 = 2$. $\square$

In the above proof, we started at the hypothesis and moved steadily from there toward a conclusion. It is also possible to work backwards from the conclusion and seeing what it would take to imply it. For instance, a typical proof of Proposition A.3.1 of this sort might look like the following:

Proof. To show B, it would suffice to show D. Since C implies D, we just need to show C. But C follows from A. $\square$

As an example of this, we give another proof of Proposition A.3.2:

Proof. To show $\sin(x/2)+1 = 2$, it would suffice to show that $\sin(x/2) = 1$. Since $x/2 = \pi/2$ would imply $\sin(x/2) = 1$, we just need to show that $x/2 = \pi/2$. But this follows since $x = \pi$. $\square$

Logically speaking, the above two proofs of Proposition A.3.2 are the same, just arranged differently. Note how this proof style is different from the (incorrect) approach of starting with the conclusion and seeing what it would imply (as in Proposition A.2.3); instead, we start with the conclusion and see what would imply it.

Another example of a proof written in this backwards style is the following:

Proposition A.3.3. *Let $0 < r < 1$ be a real number. Then the series $\sum_{n=1}^{\infty} nr^n$ is convergent.*

Proof. To show this series is convergent, it suffices by the ratio test to show that the ratio

$$\left|\frac{r^{n+1}(n+1)}{r^n n}\right| = r\frac{n+1}{n}$$

converges to something less than 1 as $n \to \infty$. Since r is already less than 1, it will be enough to show that $\frac{n+1}{n}$ converges to 1. But since $\frac{n+1}{n} = 1 + \frac{1}{n}$, it suffices to show that $\frac{1}{n} \to 0$. But this is clear since $n \to \infty$. $\square$

One could also do any combination of moving forwards from the hypothesis and backwards from the conclusion. For instance, the following would be a valid proof of Proposition A.3.1:

Proof. To show B, it would suffice to show D. So now let us show D. Since we have A by hypothesis, we have C. Since C implies D, we thus have D as desired. $\square$

Again, from a logical point of view this is exactly the same proof as before. Thus there are many ways to write the same proof down; how you do so is up to you, but certain ways of writing proofs are more readable and natural than others, and different arrangements tend to emphasize different parts of the argument. (Of course, when you are just starting out doing mathematical proofs, you're generally happy to get *some* proof of a result, and don't care so much about getting the "best" arrangement of that proof; but the point here is that a proof can take many different forms.)

The above proofs were pretty simple because there was just one hypothesis and one conclusion. When there are multiple hypotheses and conclusions, and the proof splits into cases, then proofs can get more complicated. For instance a proof might look as tortuous as this:

Proposition A.3.4. *Suppose that A and B are true. Then C and D are true.*

Proof. Since A is true, E is true. From E and B we know that F is true. Also, in light of A, to show D it suffices to show G. There are now two cases: H and I. If H is true, then from F and H we obtain C, and from A and H we obtain G. If instead I is true, then from I we have G, and from I and G we obtain C. Thus in both cases we obtain both C and G, and hence C and D. □

Incidentally, the above proof could be rearranged into a much tidier manner, but you at least get the idea of how complicated a proof could become. To show an implication there are several ways to proceed: you can work forward from the hypothesis; you can work backward from the conclusion; or you can divide into cases in the hope to split the problem into several easier sub-problems. Another is to argue by contradiction, for instance you can have an argument of the form

Proposition A.3.5. *Suppose that A is true. Then B is false.*

Proof. Suppose for sake of contradiction that B is true. This would imply that C is true. But since A is true, this implies that D is true; which contradicts C. Thus B must be false. □

As you can see, there are several things to try when attempting a proof. With experience, it will become clearer which approaches are likely to work easily, which ones will probably work but require much effort, and which ones are probably going to fail. In many cases there is

really only one obvious way to proceed. Of course, there may definitely be multiple ways to approach a problem, so if you see more than one way to begin a problem, you can just try whichever one looks the easiest, but be prepared to switch to another approach if it begins to look hopeless.

Also, it helps when doing a proof to keep track of which statements are *known* (either as hypotheses, or deduced from the hypotheses, or coming from other theorems and results), and which statements are *desired* (either the conclusion, or something which would imply the conclusion, or some intermediate claim or lemma which will be useful in eventually obtaining the conclusion). Mixing the two up is almost always a bad idea, and can lead to one getting hopelessly lost in a proof.

A.4 Variables and quantifiers

One can get quite far in logic just by starting with primitive statements (such as "$2+2=4$" or "John has black hair"), then forming compound statements using logical connectives, and then using various laws of logic to pass from one's hypotheses to one's conclusions; this is known as *propositional logic* or *Boolean logic*. (It is possible to list a dozen or so such laws of propositional logic, which are sufficient to do everything one wants to do, but I have deliberately chosen not to do so here, because you might then be tempted to memorize that list, and that is **not** how one should learn how to do logic, unless one happens to be a computer or some other non-thinking device. However, if you really are curious as to what the formal laws of logic are, look up "laws of propositional logic" or something similar in the library or on the internet.)

However, to do mathematics, this level of logic is insufficient, because it does not incorporate the fundamental concept of *variables* - those familiar symbols such as x or n which denote various quantities which are unknown, or set to some value, or assumed to obey some property. Indeed we have already sneaked in some of these variables in order to illustrate some of the concepts in propositional logic (mainly because it gets boring after a while to talk endlessly about variable-free statements such as $2+2=4$ or "Jane has black hair"). *Mathematical logic* is thus the same as propositional logic but with the additional ingredient of variables added.

A *variable* is a symbol, such as n or x, which denotes a certain type of mathematical object - an integer, a vector, a matrix, that kind of thing. In almost all circumstances, the type of object that the variable represents should be declared, otherwise it will be difficult to make well-formed statements using it. (There are very few true statements that one can make about variables without knowing the type of variables involved. For instance, given a variable x of any type whatsoever, it is true that $x = x$, and if we also know that $x = y$, then we can conclude that $y = x$. But one cannot say, for instance, that $x + y = y + x$, until we know what type of objects x and y are and whether they support the operation of addition; for instance, the above statement is ill-formed if x is a matrix and y is a vector. Thus if one actually wants to do some useful mathematics, then every variable should have an explicit type.)

One can form expressions and statements involving variables, for instance, if x is a real variable (i.e., a variable which is a real number), $x+3$ is an expression, and $x+3=5$ is a statement. But now the truth of a statement may depend on the value of the variables involved; for instance the statement $x+3=5$ is true if x is equal to 2, but is false if x is not equal to 2. Thus the truth of a statement involving a variable may depend on the *context* of the statement - in this case, it depends on what x is supposed to be. (This is a modification of the rule for propositional logic, in which all statements have a definite truth value.)

Sometimes we do not set a variable to be anything (other than specifying its type). Thus, we could consider the statement $x+3=5$ where x is an unspecified real number. In such a case we call this variable a *free variable*; thus we are considering $x+3=5$ with x a free variable. Statements with free variables might not have a definite truth value, as they depend on an unspecified variable. For instance, we have already remarked that $x+3=5$ does not have a definite truth value if x is a free real variable, though of course for each given value of x the statement is either true or false. On the other hand, the statement $(x+1)^2 = x^2+2x+1$ is true for every real number x, and so we can regard this as a true statement even when x is a free variable.

At other times, we *set* a variable to equal a fixed value, by using a statement such as "Let $x = 2$" or "Set x equal to 2". In this case, the variable is known as a *bound variable*, and statements involving only bound variables and no free variables do have a definite truth value. For instance, if we set $x = 342$, then the statement "$x + 135 = 477$" now has a definite truth value, whereas if x is a free real variable then "$x + 135 = 477$" could be either true or false, depending on what x is. Thus, as we have said before, the truth of a statement such as "$x + 135 = 477$" depends on the context - whether x is free or bound, and if it is bound, what it is bound to.

One can also turn a free variable into a bound variable by using the quantifiers "for all" or "for some". For instance, the statement

$$(x+1)^2 = x^2 + 2x + 1$$

is a statement with one free variable x, and need not have a definite truth value, but the statement

$$(x+1)^2 = x^2 + 2x + 1 \text{ for all real numbers } x$$

is a statement with one bound variable x, and now has a definite truth value (in this case, the statement is true). Similarly, the statement

$$x + 3 = 5$$

has one free variable, and does not have a definite truth value, but the statement

$$x + 3 = 5 \text{ for some real number } x$$

is true, since it is true for $x = 2$. On the other hand, the statement

$$x + 3 = 5 \text{ for all real numbers } x$$

is false, because there are some (indeed, there are many) real numbers x for which $x + 3$ is not equal to 5.

Universal quantifiers. Let $P(x)$ be some statement depending on a free variable x. The statement "$P(x)$ is true for all x of type T" means that given any x of type T, the statement $P(x)$ is true regardless of what the exact value of x is. In other words, the statement is the same as saying "if x is of type T, then $P(x)$ is true". Thus the usual way to prove such a statement is to let x be a free variable of type T (by saying something like "Let x be any object of type T"), and then proving $P(x)$ for that object. The statement becomes false if one can produce even a single counterexample, i.e., an element x which lies in T but for which $P(x)$ is false. For instance, the statement "x^2 is greater than x for all positive x" can be shown to be false by producing a single example, such as $x = 1$ or $x = 1/2$, where x^2 is not greater than x.

On the other hand, producing a single example where $P(x)$ *is* true will not show that $P(x)$ is true for *all* x. For instance, just because the equation $x + 3 = 5$ has a solution when $x = 2$ does not imply that $x + 3 = 5$ for all real numbers x; it only shows that $x + 3 = 5$ is true for some real number x. (This is the source of the often-quoted, though somewhat inaccurate, slogan "One cannot prove a statement just by giving an example". The more precise statement is that one cannot prove a "for all" statement by examples, though one can certainly prove "for some" statements this way, and one can also *disprove* "for all" statements by a single counterexample.)

It occasionally happens that there are in fact no variables x of type T. In that case the statement "$P(x)$ is true for all x of type T" is *vacuously true* - it is true but has no content, similar to a vacuous implication. For instance, the statement

$$6 < 2x < 4 \text{ for all } 3 < x < 2$$

is true, and easily proven, but is vacuous. (Such a vacuously true statement can still be useful in an argument, though this doesn't happen very often.)

One can use phrases such as "For every" or "For each" instead of "For all", e.g., one can rephrase "$(x+1)^2 = x^2 + 2x + 1$ for all real numbers x" as "For each real number x, $(x+1)^2$ is equal to x^2+2x+1". For the purposes of logic these rephrasings are equivalent. The symbol $\forall$ can be used instead of "For all", thus for instance "$\forall x \in X : P(x)$ is true" or "$P(x)$ is true $\forall x \in X$" is synonymous with "$P(x)$ is true for all $x \in X$".

Existential quantifiers. The statement "$P(x)$ is true for some x of type T" means that there exists at least one x of type T for which $P(x)$ is true, although it may be that there is more than one such x. (One would use a quantifier such as "for exactly one x" instead of "for some x" if one wanted both existence and uniqueness of such an x.) To prove such a statement it suffices to provide a single example of such an x. For instance, to show that

$$x^2 + 2x - 8 = 0 \text{ for some real number } x$$

all one needs to do is find a single real number x for which $x^2+2x-8 = 0$, for instance $x = 2$ will do. (One could also use $x = -4$, but one doesn't need to use both.) Note that one has the freedom to select x to be anything one wants when proving a for-some statement; this is in contrast to proving a for-all statement, where one has to let x be arbitrary. (One can compare the two statements by thinking of two games between you and an opponent. In the first game, the opponent gets to pick what x is, and then you have to prove $P(x)$; if you can

always win this game, then you have proven that $P(x)$ is true for *all* x. In the second game, *you* get to choose what x is, and then you prove $P(x)$; if you can win this game, you have proven that $P(x)$ is true for *some* x.)

Usually, saying something is true for *all* x is much stronger than just saying it is true for *some* x. There is one exception though, if the condition on x is impossible to satisfy, then the for-all statement is vacuously true, but the for-some statement is false. For instance

$$6 < 2x < 4 \text{ for all } 3 < x < 2$$

is true, but

$$6 < 2x < 4 \text{ for some } 3 < x < 2$$

is false.

One can use phrases such as "For at least one" or "There exists ...such that" instead of "For some". For instance, one can rephrase "$x^2 + 2x - 8 = 0$ for some real number x" as "There exists a real number x such that $x^2 + 2x - 8 = 0$". The symbol $\exists$ can be used instead of "There exists ...such that", thus for instance "$\exists x \in X : P(x)$ is true" is synonymous with "$P(x)$ is true for some $x \in X$".

A.5 Nested quantifiers

One can nest two or more quantifiers together. For instance, consider the statement

For every positive number x, there exists a positive number y such that $y^2 = x$.

What does this statement mean? It means that for each positive number x, the statement

There exists a positive number y such that $y^2 = x$

is true. In other words, one can find a positive square root of x for each positive number x. So the statement is saying that every positive number has a positive square root.

To continue the gaming metaphor, suppose you play a game where your opponent first picks a positive number x, and then you pick a positive number y. You win the game if $y^2 = x$. If you can always win the game regardless of what your opponent does, then you have proven that for every positive x, there exists a positive y such that $y^2 = x$.

Negating a universal statement produces an existential statement. The negation of "All swans are white" is not "All swans are not white", but rather "There is some swan which is not white". Similarly, the negation of "For every $0 < x < \pi/2$, we have $\cos(x) \geq 0$" is "For some $0 < x < \pi/2$, we have $\cos(x) < 0$, **not** "For every $0 < x < \pi/2$, we have $\cos(x) < 0$".

Negating an existential statement produces a universal statement. The negation of "There exists a black swan" is not "There exists a swan which is non-black", but rather "All swans are non-black". Similarly, the negation of "There exists a real number x such that $x^2 + x + 1 = 0$" is "For every real number x, $x^2 + x + 1 \neq 0$", **not** "There exists a real number x such that $x^2 + x + 1 \neq 0$". (The situation here is very similar to how "and" and "or" behave with respect to negations.)

If you know that a statement $P(x)$ is true for all x, then you can set x to be anything you want, and $P(x)$ will be true for that value of x; this is what "for all" means. Thus for instance if you know that

$$(x+1)^2 = x^2 + 2x + 1 \text{ for all real numbers } x,$$

then you can conclude for instance that

$$(\pi + 1)^2 = \pi^2 + 2\pi + 1,$$

or for instance that

$$(\cos(y) + 1)^2 = \cos(y)^2 + 2\cos(y) + 1 \text{ for all real numbers } y$$

(because if y is real, then $\cos(y)$ is also real), and so forth. Thus universal statements are very versatile in their applicability - you can get $P(x)$ to hold for whatever x you wish. Existential statements, by contrast, are more limited; if you know that

$$x^2 + 2x - 8 = 0 \text{ for some real number } x$$

then you cannot simply substitute in any real number you wish, e.g., π, and conclude that $\pi^2 + 2\pi - 8 = 0$. However, you can of course still conclude that $x^2 + 2x - 8 = 0$ for some real number x, it's just that you don't get to pick which x it is. (To continue the gaming metaphor, you can make $P(x)$ hold, but your opponent gets to pick x for you, you don't get to choose for yourself.)

Remark A.5.1. In the history of logic, quantifiers were formally studied thousands of years before Boolean logic was. Indeed, *Aristotlean logic*, developed of course by Aristotle (384BC – 322BC) and his school, deals with objects, their properties, and quantifiers such as "for all" and "for some". A typical line of reasoning (or *syllogism*) in Aristotlean logic goes like this: "All men are mortal. Socrates is a man. Hence, Socrates is mortal". Aristotlean logic is a subset of mathematical logic, but is not as expressive because it lacks the concept of logical connectives such as and, or, or if-then (although "not" is allowed), and also lacks the concept of a binary relation such as = or <.

Swapping the order of two quantifiers may or may not make a difference to the truth of a statement. Swapping two "for all" quantifiers is harmless: a statement such as

For all real numbers a, and for all real numbers b,
we have $(a+b)^2 = a^2 + 2ab + b^2$

is logically equivalent to the statement

For all real numbers b, and for all real numbers a,
we have $(a+b)^2 = a^2 + 2ab + b^2$

(why? The reason has nothing to do with whether the identity $(a+b)^2 = a^2 + 2ab + b^2$ is actually true or not). Similarly, swapping two "there exists" quantifiers has no effect:

There exists a real number a, and there exists a real number b,
such that $a^2 + b^2 = 0$

is logically equivalent to

There exists a real number b, and there exists a real number a,
such that $a^2 + b^2 = 0$.

However, swapping a "for all" with a "there exists" makes a lot of difference. Consider the following two statements:

(a) For every integer n, there exists an integer m which is larger than n.

(b) There exists an integer m such that m is larger than n for every integer n.

Statement (a) is obviously true: if your opponent hands you an integer n, you can always find an integer m which is larger than n. But Statement (b) is false: if you choose m first, then you cannot ensure that m is larger than every integer n; your opponent can easily pick a number n bigger than m to defeat that. The crucial difference between the two statements is that in Statement (a), the integer n was chosen *first*, and integer m could then be chosen in a manner depending on n; but in Statement (b), one was forced to choose m first, without knowing in advance what n is going to be. In short, the reason why the order of quantifiers is important is that the inner variables may possibly depend on the outer variables, but not vice versa.

— Exercises —

Exercise A.5.1. What does each of the following statements mean, and which of them are true? Can you find gaming metaphors for each of these statements?

(a) For every positive number x, and every positive number y, we have $y^2 = x$.

(b) There exists a positive number x such that for every positive number y, we have $y^2 = x$.

(c) There exists a positive number x, and there exists a positive number y, such that $y^2 = x$.

(d) For every positive number y, there exists a positive number x such that $y^2 = x$.

(e) There exists a positive number y such that for every positive number x, we have $y^2 = x$.

A.6 Some examples of proofs and quantifiers

Here we give some simple examples of proofs involving the "for all" and "there exists" quantifiers. The results themselves are simple, but you should pay attention instead to how the quantifiers are arranged and how the proofs are structured.

Proposition A.6.1. *For every $\varepsilon > 0$ there exists a $\delta > 0$ such that $2\delta < \varepsilon$.*

Proof. Let $\varepsilon > 0$ be arbitrary. We have to show that there exists a $\delta > 0$ such that $2\delta < \varepsilon$. We only need to pick one such δ; choosing $\delta := \varepsilon/3$ will work, since one then has $2\delta = 2\varepsilon/3 < \varepsilon$. □

Notice how ε has to be arbitrary, because we are proving something for *every* ε; on the other hand, δ can be chosen as you wish, because you only need to show that there *exists* a δ which does what you want. Note also that δ can depend on ε, because the δ-quantifier is nested inside the ε-quantifier. If the quantifiers were reversed, i.e., if you were asked to prove "There exists a $\delta > 0$ such that for every $\varepsilon > 0$, $2\delta < \varepsilon$", then you would have to select δ *first* before being given ε. In this case it is impossible to prove the statement, because it is false (why?).

Normally, when one has to prove a "There exists..." statement, e.g., "Prove that there exists an $\varepsilon > 0$ such that X is true", one proceeds by selecting ε carefully, and then showing that X is true for that ε. However, this sometimes requires a lot of foresight, and it is legitimate to defer the selection of ε until later in the argument, when it becomes clearer what properties ε needs to satisfy. The only thing to watch out for is to make sure that ε does not depend on any of the bound variables nested inside X. For instance:

Proposition A.6.2. *There exists an $\varepsilon > 0$ such that $\sin(x) > x/2$ for all $0 < x < \varepsilon$.*

Proof. We pick $\varepsilon > 0$ to be chosen later, and let $0 < x < \varepsilon$. Since the derivative of $\sin(x)$ is $\cos(x)$, we see from the mean-value theorem we have

$$\frac{\sin(x)}{x} = \frac{\sin(x) - \sin(0)}{x - 0} = \cos(y)$$

for some $0 \leq y \leq x$. Thus in order to ensure that $\sin(x) > x/2$, it would suffice to ensure that $\cos(y) > 1/2$. To do this, it would suffice to ensure that $0 \leq y < \pi/3$ (since the cosine function takes the value of 1 at 0, takes the value of 1/2 at $\pi/3$, and is decreasing in between). Since $0 \leq y \leq x$ and $0 < x < \varepsilon$, we see that $0 \leq y < \varepsilon$. Thus if we pick $\varepsilon := \pi/3$, then we have $0 \leq y < \pi/3$ as desired, and so we can ensure that $\sin(x) > x/2$ for all $0 < x < \varepsilon$. □

Note that the value of ε that we picked at the end did not depend on the nested variables x and y. This makes the above argument legitimate. Indeed, we can rearrange it so that we don't have to postpone anything:

Proof. We choose $\varepsilon := \pi/3$; clearly $\varepsilon > 0$. Now we have to show that for all $0 < x < \pi/3$, we have $\sin(x) > x/2$. So let $0 < x < \pi/3$ be arbitrary. By the mean-value theorem we have

$$\frac{\sin(x)}{x} = \frac{\sin(x) - \sin(0)}{x - 0} = \cos(y)$$

for some $0 \leq y \leq x$. Since $0 \leq y \leq x$ and $0 < x < \pi/3$, we have $0 \leq y < \pi/3$. Thus $\cos(y) > \cos(\pi/3) = 1/2$, since cos is decreasing on the interval $[0, \pi/3]$. Thus we have $\sin(x)/x > 1/2$ and hence $\sin(x) > x/2$ as desired. □

If we had chosen ε to depend on x and y then the argument would not be valid, because ε is the outer variable and x, y are nested inside it.

A.7 Equality

As mentioned before, one can create statements by starting with expressions (such as $2 \times 3 + 5$) and then asking whether an expression obeys a certain property, or whether two expressions are related by some sort of relation ($=$, $\leq$, $\in$, etc.). There are many relations, but the most important one is *equality*, and it is worth spending a little time reviewing this concept.

Equality is a relation linking two objects x, y of the same type T (e.g., two integers, or two matrices, or two vectors, etc.). Given two such objects x and y, the statement $x = y$ may or may not be true; it depends on the value of x and y and also on how equality is defined for the class of objects under consideration. For instance, as real numbers, the two numbers $0.9999\ldots$ and 1 are equal. In modular arithmetic with modulus 10 (in which numbers are considered equal to their remainders modulo 10), the numbers 12 and 2 are considered equal, $12 = 2$, even though this is not the case in ordinary arithmetic.

How equality is defined depends on the class T of objects under consideration, and to some extent is just a matter of definition. However, for the purposes of logic we require that equality obeys the following four *axioms of equality*:

- (Reflexive axiom). Given any object x, we have $x = x$.
- (Symmetry axiom). Given any two objects x and y of the same type, if $x = y$, then $y = x$.
- (Transitive axiom). Given any three objects x, y, z of the same type, if $x = y$ and $y = z$, then $x = z$.
- (Substitution axiom). Given any two objects x and y of the same type, if $x = y$, then $f(x) = f(y)$ for all functions or operations f.

Similarly, for any property $P(x)$ depending on x, if $x = y$, then $P(x)$ and $P(y)$ are equivalent statements.

The first three axioms are clear, together, they assert that equality is an *equivalence relation*. To illustrate the substitution we give some examples.

Example A.7.1. Let x and y be real numbers. If $x = y$, then $2x = 2y$, and $\sin(x) = \sin(y)$. Furthermore, $x + z = y + z$ for any real number z.

Example A.7.2. Let n and m be integers. If n is odd and $n = m$, then m must also be odd. If we have a third integer k, and we know that $n > k$ and $n = m$, then we also know that $m > k$.

Example A.7.3. Let x, y, z be real numbers. If we know that $x = \sin(y)$ and $y = z^2$, then (by the substitution axiom) we have $\sin(y) = \sin(z^2)$, and hence (by the transitive axiom) we have $x = \sin(z^2)$.

Thus, from the point of view of logic, we can define equality on a class of objects however we please, so long as it obeys the reflexive, symmetry, and transitive axioms, and is consistent with all other operations on the class of objects under discussion in the sense that the substitution axiom was true for all of those operations. For instance, if we decided one day to modify the integers so that 12 was now equal to 2, one could only do so if one also made sure that 2 was now equal to 12, and that $f(2) = f(12)$ for any operation f on these modified integers. For instance, we now need $2 + 5$ to be equal to $12 + 5$. (In this case, pursuing this line of reasoning will eventually lead to modular arithmetic with modulus 10.)

— Exercises —

Exercise A.7.1. Suppose you have four real numbers a, b, c, d and you know that $a = b$ and $c = d$. Use the above four axioms to deduce that $a + d = b + c$.