1 A quick introduction to logic

Mathematics deals in *statements* — assertions that are either unambiguously true or unambiguously false — and *proofs* — watertight arguments, based on fundamental rules of logic, that establish irrefutably the truth or falsity of statements. It's a language, and as such has a vocabulary and a grammar. The vocabulary includes technical words, that we will learn as we need them, but it also includes many everyday words and phrases, such as "if", "and", "implies", "it follows that", et cetera, whose mathematical meanings sometimes differ slightly from their ordinary meanings[1]. The grammar consists of the basic rules of logic or inference that allow us to understand the truth or falsity of complex statements from our knowledge of the truth or falsity of simpler statements.

This section of the notes introduces the basic vocabulary and grammar of mathematics. In Section 2, we will begin to talk about how the vocabulary and grammar are used to prove statements. In writing these two section I have drawn heavily on notes by John Bryant & Penelope Kirby [1] (Florida State), and Tom Hutchings [2] (Berkeley).

Underlying many of the examples in these first two sections (and all of the rest of the notes) is the basic language of sets. A quick introduction to the language of sets is given in Appendix A.

1.1 Statements

A basic object in mathematics is the *statement*: an assertion that is either true or false:

- "3 is a prime number." A true statement — we say that it has *truth value True* or simply T.
- "November 26, 1971 was a Friday." This is also true (you could look it up).
- "If I go out in the rain without a coat on, I will get wet."
- "There is no whole number between 8 and 10". A fine statement, albeit a false one — we say that it has *truth value False* or simply F.
- "There is life on Mars." Even though we don't know (yet) whether this is a true statement or a false one, everyone would agree that it has a definite truth value — there either is or there isn't life on Mars. So this is a statement.
- "All positive whole numbers are even". A false statement.
- "At least one positive whole number is even". A true statement.

[1] For example: if I tell you that tonight I'll either see Lion King at the movie theater, or go to a concert at DeBartolo, you might be surprised when I end up going to the movie *and* the concert: in ordinary language, "or" is most usually exclusive. But when we say mathematically that "either p or q is true", that always leaves open the possibility that *both* p *and* q are true: in mathematical language, "or" is always inclusive.

- "If you draw a closed curve on a piece of paper, you can find four points on the curve that form the four corners of a square". This is a statement, but is it a true one or a false one? Surprisingly, we don't know. The assertion was conjectured to be true in 1911 (by Otto Toeplitz)[2], but it has resisted all efforts at either a proof or a disproof.

Here are examples of things which are *not* statements:

- "Do you like ice cream?" A question, not an assertion.
- "Turn in the first homework assignment by Friday." An imperative, or a command, not an assertion.
- "$3x^2 - 2x + 1 = 0$." This is an assertion, but it does not have a definite truth-value — there are some x's for which it is true, and some for which it is false. So this is not a statement.
- "This statement is false." This is certainly an assertion. Is it true? If so, then it must be false. But if it is false, then it must be true. We can't assign a definite truth value to this assertion, so it is not a statement. This, and many other assertions like it, are referred to as *paradoxes*, and we try to avoid them as much as possible!

Some of our statement examples were quite simple ("There is life on Mars"), while others were more complicated beasts built up from simpler statements ("If I go out in the rain without a coat on, I will get wet"). Here we review the ways in which we build more complicated statements from simpler ones[3].

- **Negation**: If p is a statement, then the negation of p, which we sometimes call "not p" and sometimes write symbolically as $\neg p$[4], is a statement that has the opposite truth value to p. So:
  - $\neg p$ is false when p is true, and
  - $\neg p$ is true when p is false.

  Here are two clarifying examples:
  - If p is "There is life on Mars" then the negation of p is "There is no life on Mars". (It could also be "It is not the case that there is life on Mars".)
  - If p is

[2] O. Toeplitz, Über einige Aufgaben der Analysis situs, *Verhandlungen der Schweizerischen Naturforschenden Gesellschaft in Solothurn* **94** (1911), p. 197; see also https://en.wikipedia.org/wiki/Inscribed_square_problem

[3] A good analogy for what we are about to do, is how we combine numbers to form more complicated expressions that are still numbers, via the operations of addition, subtraction, multiplication, division, taking square roots, et cetera.

[4] Read this symbol as "not p".

"For every whole number n, there is a field with n elements"

(it doesn't matter what "field" might mean), then the negation of p is

"There is *some* whole number n for which there is not a field with n elements".

The negation is **NOT**: "There is *no* whole number n for which there is a field on n elements". Between "For every ..." and "There is no ..." there is a huge gap that is not covered by either statement — what if there are some n for which there is a field with n elements, and others for which there isn't? Then both of "For every ..." and "There is no ..." are false. But there's no such gap between "For every ..." and "There is some that is not ..." — whatever the possible sizes of fields, either one or other statement is true and the other is false. We'll come back to this idea when we talk about quantifiers.

We can use a *truth table* to summarize the effect of negation on a statement:

p	not p
T	F
F	T

We read this as: if p is true then $\neg p$ is false (first row), and if p is false then $\neg p$ is true (second row).

- **Conjunction**: If p and q are statements, then the conjunction of p and q is a statement that is true whenever both p *and* q are true, and false otherwise. We will typically write "p and q" for the conjunction. The symbolic notation is $p \wedge q$[5].

  If p is "There is life on Mars" and q is "There is water on Mars", then the conjunction $p \wedge q$ is "There is both life and water on Mars", and would only be true if we found there to be *both* life *and* water on Mars; finding that there is only one of these, or none of them, would not be good enough to make the conjunction true.[6]

  Here is the truth table for conjunction:

p	q	p and q
T	T	T
T	F	F
F	T	F
F	F	F

  Notice that since there are two options for the truth values of each of p, q, the truth table needs $2 \times 2 = 4$ rows.

[5]Read this as "p and q"

[6]Instead of saying "the conjunction $p \wedge q$" we will sometimes say "the conjunction of p and q", or even "the 'and' of p and q", or simply "p and q".

- **Disjunction**: If p and q are statements, then the disjunction of p and q is the statement that is true whenever at least one of p, q is true, and false otherwise. We write "p or q" for this compound statement, and sometimes denote it symbolically by $p \vee q$[7]. It is very important to remember that, by universal convention among mathematicians, "or" is *inclusive*: $p \vee q$ is true if p is true, or if q is true, *or if both are true*. This is a different from the ordinary language use of "or", which tends to be exclusive (see the footnote in the introductory text to Section 1).

  If p is "There is life on Mars" and q is "There is water on Mars", then the disjunction of p and q is "There is either life or water on Mars", and would only be false if we found there to be *neither* life *nor* water on Mars; finding any one of these (or both) would be good enough to make the disjunction true.[8]

  Here is the truth table for disjunction:

p	q	p or q
T	T	T
T	F	T
F	T	T
F	F	F

  As a specific, important, example, here's the truth table of "p or (not p)":

p	not p	p or (not p)
T	F	T
F	T	T

  (notice only two lines were needed here, because only one basic statement — p — is involved in "p or (not p)", and it can either be True or False).

  Observe: regardless of the truth value of p, the truth value of "p or (not p)" is True. This makes "p or (not p)" something called a *tautology* — a complex statement, made up of a combination of simpler statements, that turns out to be True regardless of the truth values of the simpler statements. In other words, a tautology is an *absolutely, always, eternally* True statement. Much of mathematics is engaged in the hunt for tautologies (usually going by the name "theorems"). The particular tautology we have just seen is an especially simple, but important, one. It's called *the law of the excluded middle*, because it says that

  > every statement is either True, or False

[7]Read this as "p or q"

[8]As with conjunction, instead of saying "the disjunction $p \vee q$" we will sometimes say "the disjunction of p and q", or even "the 'or' of p and q", or simply "p or q".

(there is no middle).[9]

From these basic operations we can build up much more complicated compound statements. For example, if we have three statements p, q and r we can consider the compound statement

$$\text{not } (p \text{ and } q) \text{ or not } (p \text{ and } r) \text{ or not } (q \text{ and } r).$$

or

$$\neg(p \wedge q) \vee \neg(p \wedge r) \vee \neg(q \wedge r)$$

(Notice that the symbolic formulation is typographically a lot nicer in this case; I'll stick with that formulation throughout this example.) If p and q are true and if r is false, then $p \wedge q$ is true, so $\neg(p \wedge q)$ is false. By similar reasoning $\neg(p \wedge r)$ and $\neg(q \wedge r)$ are both true. So we are looking at the disjunction of three statements, one of which is false and the other two of which are true. We haven't defined the disjunction of three statements, but it's obvious what it must be: the disjunction is true as long as at least one of the three statements is true. That means that in the particular case under consideration (p, q true, r false), the compound statement $\neg(p \wedge q) \vee \neg(p \wedge r) \vee \neg(q \wedge r)$ is true.

We can do this for all $2 \times 2 \times 2 = 8$ possible assignments of truth values to p, q and r, to form a truth table for the compound statement:

p	q	r	$\neg(p \wedge q) \vee \neg(p \wedge r) \vee \neg(q \wedge r)$
T	T	T	F
T	T	F	T
T	F	T	T
T	F	F	T
F	T	T	T
F	T	F	T
F	F	T	T
F	F	F	T

It appears that the statement $\neg(p \wedge q) \vee \neg(p \wedge r) \vee \neg(q \wedge r)$ is false only when all three of p, q, r are true, so in words it is the statement "At least one of p, q, r is false".

As another example, consider $\neg(p \wedge q \wedge r)$ (where again we haven't defined the conjunction of three statements, but it's obvious what it must be: the conjunction is true only if all three

[9]Given how we defined a statement, the law of the excluded middle must be true. It is a "tautology" in the ordinary language sense.

of the three statements are True). Here's the truth table for this compound statement:

p	q	r	$\neg(p \wedge q \wedge r)$
T	T	T	F
T	T	F	T
T	F	T	T
T	F	F	T
F	T	T	T
F	T	F	T
F	F	T	T
F	F	F	T

It's exactly the same as the truth table for $\neg(p \wedge q) \vee \neg(p \wedge r) \vee \neg(q \wedge r)$, which of course it should be: even without writing the full truth table, it should have been evident that the statement $\neg(p \wedge q \wedge r)$ is the same as "At least one of p, q, r is false". This illustrates that two apparently different compound statements may have the same truth tables, and so may be considered "the same" statement.

Formally, if A is one statement built from the simpler statements p, q and r, using combinations of $\neg$, $\wedge$ and $\vee$, and B is another one, then A and B are *equivalent* (though we will often somewhat sloppily say *the same*) if: for each possible assignment of truth values to p, q and r, the resulting truth value of A is the same as the resulting truth value of B. Effectively, this means that if you use a single truth table to figure out what A and B look like, then the column corresponding to A is the same as the column corresponding to B. Of course, this can be extended to pairs of statements built from any number of simpler statements.

Here are a few pairs of equivalent statements; the equivalence of each pair is quickly verified by comparing truth tables.

- $(p \wedge q) \wedge r$ and $p \wedge (q \wedge r)$
- $(p \vee q) \vee r$ and $p \vee (q \vee r)$
- $\neg(p \wedge q)$ and $(\neg p) \vee (\neg q)$
- $\neg(p \vee q)$ and $(\neg p) \wedge (\neg q)$

If you were uncomfortable with line "We haven't defined the conjunction of three statements, but it's obvious what it must be ...", then you will be happy with the equivalence of the first pair: it shows that whatever pair-by-pair order we choose to deal with the conjunction of three statements, the resulting truth table is the same (and is the same as the truth table of "All of p, q, r are true"), so it is really ok to slightly sloppily talk about the conjunction of three statements. The equivalence of the second pair does the same job for the disjunction of three statements. With some (actually a lot) more work we could show that if $p_1, p_2, \ldots, p_n$

are n statements, then whatever pair-by-pair order we choose to deal with the conjunction $p_1 \wedge p_2 \wedge \ldots \wedge p_n$ the resulting truth table is the same, and is the same as the truth table of "All of p_1, p_2, ..., p_n are true"); and there is an analogous statement for $p_1 \vee p_2 \vee \ldots \vee p_n$. (We will return to this, in the slightly different but essentially equivalent context of "associativity of addition", when we come to discuss proofs by induction.)

The third and fourth pairs of equivalences above are called *De Morgan's laws*, which we will return to in more generality later.

To illustrate what was meant earlier by "use a single truth table to figure out what A and B look like", here's the single truth table that shows the validity of the second of De Morgan's laws, that $\neg(p \vee q)$ and $(\neg p) \wedge (\neg q)$ are equivalent:

p	q	$p \vee q$	$\neg(\mathbf{p} \vee \mathbf{q})$	$\neg p$	$\neg q$	$(\neg \mathbf{p}) \wedge (\neg \mathbf{q})$
T	T	T	**F**	F	F	**F**
T	F	T	**F**	F	T	**F**
F	T	T	**F**	T	F	**F**
F	F	F	**T**	T	T	**T**

Some comments:

- We've introduced some auxiliary columns into the truth-table, mostly for bookkeeping purposes[10]; some people find this helpful, others don't, it is entirely a matter of personal taste.

- The columns corresponding to $\neg(p \vee q)$ and $(\neg p) \wedge (\neg q)$ are identical. Since for each particular row, the truth values assigned to p and q for the purposes of determining the truth value of $\neg(p \vee q)$, are the same as the truth values assigned to p and q for the purposes of determining the truth value of $(\neg p) \wedge (\neg q)$, this allows us to conclude from the truth table that $\neg(p \vee q)$ and $(\neg p) \wedge (\neg q)$ are equivalent statements.

1.2 An note on parentheses

There's an inherent ambiguity in any reasonably complicated combination of numbers via the basic arithmetic operations of $+$, $-$, et cetera. For example, if we perform the addition first and then the multiplication we get

$$1 + 2 \times 3 = 3 \times 3 = 9$$

while if we perform the multiplication first and then the addition we get

$$1 + 2 \times 3 = 1 + 6 = 7.$$

[10] I believe that "bookkeeping" (and variants "bookkeeper", "bookkeepers") is the only unhyphenated English word with three double letters in a row ("woolly" doesn't count). There's also exactly one with four doubles — subbookkeeper. This will not be on the exam.

Similarly, there's an inherent ambiguity in any reasonably complex statement, related to the order in which to perform operations such as $\neg$, $\wedge$ and $\vee$ mentioned in the statement. Different choices of order may lead to different truth tables. For example, consider the statement "$notporq$". This could mean "take the disjunction of the following two statements:

- q, and
- the negation of p".

Or it could mean "take the negation of the following statement:

- the disjunction of q and p".

This are unfortunately different statements: if p and q are both true, then the first is true while the second is false.

One way to avoid both these ambiguities (arithmetical and logical) is to decide, once and for all time, on an order of precedence among arithmetic operations, and an order of precedence among logical operations. There *is* a fairly standard such order[11] for arithmetic operations. There are also established orders of precedence among logical operations. But there Tare two problems with taking this approach to eliminate ambiguity:

- first, there are *competing* orders of precedence, none universal, so that order-of-precedence alone does not eliminates all ambiguity, and
- second, an order of precedence among logical operators is something that one must remember, and there is no obvious motivation behind it to act as a memory aid. Who wants that?

For these reasons, I prefer to avoid ambiguity by using parentheses to indicate order of operation, with the convention being that you work from the inside out. So for example, to indicate "take the disjunction of the two statements 'q' and 'the negation of p'" I would write

"(not p) or q"

(indicating that the negation should be performed before the disjunction), while to indicate "take the negation of the statement 'the disjunction of q and p'" I would write

"not (p or q)"

(indicating that the disjunction should be performed before the negation).

I encourage you to carefully consider any statement that you write for ambiguity, and make every effort to de-ambiguize!

[11]see, for example, https://en.wikipedia.org/wiki/Logical_connective#Order_of_precedence

1.3 Implication

Implication is by far the most important operation that builds more complicated statements from simpler ones — it lies at the heart or virtually all logical arguments — and so (unsurprisingly) is probably the most subtle.

If p and q are two statements, then we can form the *implication* assertion "p implies q", symbolically $p \Rightarrow q$[12]. We *define* "p implies q" to be (a statement that is logically equivalent to) the statement

"(not p) or (p and q)".

To illustrate: we have already seen the example "If I go out in the rain without a coat on, I will get wet." This is the implication "p implies q" where p is "I go out in the rain without a coat" and q is "I get wet". Viewed as an ordinary language sentence it conveys precisely the same meaning as "either

- I don't go out in the rain without a coat,

or

- I do go out in the rain without a coat, and (as a result), I get wet"

which is indeed "(not p) or (p and q)".

The implication "p implies q" can be rendered into ordinary language in many other ways, such as:

- If p (happens) then (so does) q
- (The occurrence of) p is (a) sufficient (condition) for q (to happen)
- q (happens) whenever p (happens).

Mirroring the list above, the implication "If I go out in the rain without a coat on, I will get wet" can be expressed in ordinary language as

- My going out in the rain without a coat leads to (implies) my getting wet
- If I go out in the rain without a coat on, then I get wet
- My going out in the rain without a coat is a sufficient condition for me to get wet
- I get wet whenever I go out in the rain without a coat.

Some more notation related to implication:

- p is referred to as the *premise* or *hypothesis* of the implication

[12] Read this as "p implies q".

- q is referred to as the *conclusion*.

We can easily form the truth table of "p implies q", by forming the truth table of "(not p) or (p and q)":

p	q	not p	p and q	(not p) or (p and q)	p implies q ($p \Rightarrow q$)	(not p) or q
T	T	F	T	T	T	T
T	F	F	F	F	F	F
F	T	T	F	T	T	T
F	F	T	F	T	T	T

For good measure, I've thrown in the truth table of "(not p) or q" above; notice that it is exactly the same as that of "(not p) or (p and q)", that is, the same as "p implies q", and many texts take this simpler statement to be the definition of "p implies q".

> **Convention**: From here on, we will take "p implies q" to mean simply "(not p) or q".

Two lines of the truth table for "p implies q" should be obvious: If p is true, and q is true, then surely we want "p implies q" to be true; while if p is true but q is false, then surely we want "p implies q" to be false. The other two lines of the truth table, corresponding to when p is false, are more subtle. To justify them, think of "p implies q" as a promise, a contract, that says that **IF** something (p) happens, **THEN** something else (q) happens. The contract is a good one in the case when p and q both happen, and a bad one when p happens but q doesn't (that justifies the first two lines). If p *doesn't* happen (the last two lines of the table) then the contract is never invoked, so there is no basis on which to declare it bad, so we declare it good.

In terms of our example ("If I go out in the rain without a coat on, I will get wet"): suppose the TV weather forecaster tells me that if I go out without my coat in today's rain, I will get wet. From this, I would expect to get wet if I did go out in the rain without my coat; if that happened I would say "true promise" about the forecaster's statement, whereas if I went out in the rain without my coat and didn't get wet, I would say "false promise". But what if I didn't go out in the rain without a coat? The forecaster said nothing about what happens then, so whether I stay dry (by going out with a coat, or by staying home), or get wet (by taking a bath, or because my house has a leaking roof), she would not be breaking any kind of promise. If either of these last two things occur, I should still say that the implication stated was true because she did not break her promise.

If this isn't convincing, another justification for the "implies" truth table is given in the first homework.

The negation of an implication

We have observed via a truth table that "p implies q" is equivalent to "(not p) or q" (and, to repeat a remark from earlier, in many texts this is the definition of "implies"). By De

Morgan's law (and the easy fact that we have not yet explicitly stated, that "not (not p)" always has the same truth value as p), the negation "not (p implies q)" of an implication is equivalent to "p and (not q)". Think of it in these terms: to "falsify" the weather forecasters contract, to demonstrate that it was not a sound contract, you would need to both

- go out in the rain without a coat on

and

- *not* get wet.

The contrapositive of an implication

Here is another statement that is equivalent to "p implies q": the statement

"(not q) implies (not p)".

This reformulation is called the *contrapositive* of the original implication; we will see many times as the year progresses.

There are a few ways to check that the contrapositive of an implication is equivalent to the implication:

Via a truth table The easiest way, but not too informative.

Via DE Morgan's law By definition, "(not q) implies (not p)" is equivalent to "(not (not q)) or (not p)", which is in turn equivalent to "q or (not p)", which is equivalent to "(not p) or q", which by definition is equivalent to "p implies q".[13]

Via "reasoning" (probably the best way to initially get an understand of what's going on) After thinking through some examples (both non-mathematical and mathematical) you should be fairly comfortable with the fact that "p implies q" carries exactly the same content and meaning as "(not q) implies (not p)". Some examples:

"If I go out in the rain without a coat, then I will get wet"

carries just the same meaning as

"If I didn't get wet, then I didn't go out in the rain without a coat",

and

"If $b^2 - 4ac < 0$ then the equation $ax^2 + bx + c = 0$ doesn't have real solutions"

carries just the same meaning as

"If $ax^2 + bx + c = 0$ has real solutions, then $b^2 - 4ac \geq 0$".

[13] This can be thought of as our first example of a "proof" in the way mathematics understands the word: we have carefully put together some simpler truths ("not (not p)" is the same as p; "a or b" is the same as "b or a") in just the right order to deduce a non-obvious truth, that two superficially different statements are in fact the same.

The converse of an implication

A statement related to "p implies q" that is **NOT** equivalent to it (**NOT**, **NOT**, **NOT**, really **NOT**), is the *converse* "q implies p", which has truth table

p	q	q implies p
T	T	T
T	F	T
F	T	F
F	F	T

Note that this is *not* the same as the truth table of "p implies q". One strategy we will **NEVER** employ to show that p implies q is to show q implies p and then say that that's enough to deduce that p implies q — as the truth table shows, it **ISN'T**!!! Convincing me that

if I get wet, then I go out in the rain without my coat on

does not help towards convincing me that

if I go out in the rain without my coat on, then I get wet.

Indeed, I don't think you could convince me of the former, since it's not a true statement: if on a warm, dry morning I take a bath, then I get wet (so invoke the contract in the statement above), but I don't go out in the rain without my coat on, so the contract fails. The latter statement, on the other hand ("if I go out in the rain without my coat on, then I get wet") is, I think, true.

One reason that the converse of an implication gets confused with the implication itself, is that sometimes when we use implicative language in ordinary-speak, we actual mean the converse. A classic example is: "If you don't eat your vegetables, you won't get dessert". Formally this is the implication "p implies q" where p:"you don't eat your vegetables" and q:"you don't get dessert". And if it really is this implication, then you would not be upset when, upon eating your vegetables, you were *not* given dessert: you didn't "not eat your vegetables", so the contract *wasn't* invoked, and in this case there is no promise regarding dessert!

But in fact, you would be justified in being very peeved if, after diligently eating your vegetables, you were denied dessert. This is because you, like most sensible people, would have interpreted "If you don't eat your vegetables, you won't get dessert" as a contract whose meaning is that if you eat your vegetables, then you get rewarded with dessert. In other words, "(not p) implies (not q)", or, contrapositively (and equivalently), "q implies p." So although the ordinary language wording of the statement formally parsed as an implication in one direction, its meaning was really an implication in the converse direction.

These kinds of ambiguities occur a lot in ordinary language, and for this reason I will try to keep my examples in the mathematical realm, where there should be no ambiguity.

The "if and only if" ("iff") statement

Related to implication is *bidirectional implication*, or the *if and only if* statement. The statement "p if and only if q" is shorthand for "(p implies q) and (q implies p)", and is denoted symbolically by $p \Leftrightarrow q$[14] (an explanation for the terminology "if and only of" is given in Section 1.5). The sense of this statement is that p and q sink or swim together; for the bidirectional implication to be true, it must be that either p and q are simultaneously true or simultaneously false. The truth table for bidirectional implication is:

p	q	p if and only if q
T	T	T
T	F	F
F	T	F
F	F	T

The phrase "if and only if" is often abbreviated to "iff".

The bidirectional implication statement can be rendered into ordinary english as:

- p (happens) if and only if q (happens)
- (The occurrence of) p is (a) necessary and sufficient (condition) for q (to happen)
- (The occurrence of) q is (a) necessary and sufficient (condition) for p (to happen).

Here's an example from the world of sports:

"A team wins the World Series
if and only if .
they win the last MLB game of the calendar year"

Indeed, to the win the World Series, it is *necessary* to win the last game of the year; it is also *sufficient*.

1.4 An note on symbols versus words

The field of logic is filled with precisely defined symbolic notations, such as $\neg$, $\wedge$, $\vee$, $\Rightarrow$, $\Leftrightarrow$, $\therefore$, $\because$, $\exists$, $\forall$ (we'll see these last two in a while), et cetera. Many of these appear in this section, usually to keep complex propositions manageable. But if you pick up a mathematical paper, you will notice an

almost complete absence

of these symbols. The convention that is almost universally adhered to by mathematicians today is to

[14] Read this as "p if and only if q".

write mathematics in prose.

For example you will frequently see things like "A implies B, which implies C, which in turn implies D" in a mathematical exposition (note that this is a complete English sentence), and almost never see the symbolically equivalent

$$\text{“}A \Rightarrow B \Rightarrow C \Rightarrow D''.$$

And often a proof will be presented in the following way in a paper: "Since A and B are true, and we have argued that A and B together imply C, we deduce that C is true", but you will (almost) never see the symbolically equivalent

$$\begin{array}{rll} & & \text{“}A \text{ is true} \\ & & B \text{ is true} \\ \therefore & & A \wedge B \\ & \text{also} & (A \wedge B) \Rightarrow C \text{ is true} \\ \therefore & & C \text{ is true”}. \end{array}$$

Although the symbolic notation is sometimes a nice shorthand, when we come to write proofs of propositions I will be **strongly encouraging** you to follow the standard convention, and present proofs in (essentially) complete English sentences, avoiding logical symbols as much as possible. There will be much more on this later, particularly in discussions of homework problems.

1.5 An note on language: "if and only if", and "necessary and sufficient"

The logical implication "p implies q" is often rendered into ordinary language as "p is sufficient for q" (this should make sense: "p is sufficient for q" says that if p happens, then q happens, but allows for the possibility that q happens even when p doesn't, and this is exactly what "p implies q" means). Another phrase that you will encounter a lot is

> "p is necessary for q".

"p is necessary for q" means the same thing as "if p doesn't happen, then q doesn't happen", which is the same as "(not p) implies (not q)", which is the contrapositive of (and so equivalent to) "q implies p".

It is for this reason that the bidirectional equivalence "p if and only if q" is often rendered into ordinary language as

> "p is necessary and sufficient for q".

The "sufficient" part captures "p implies q", and as we have just seen, the "necessary" part captures "q implies p".

What about the phrase "if and only if" itself? Remember that the bidirectional implication between p and q is shorthand for "(p implies q) *and* (q implies p)". The logical implication "q implies p" is often rendered into English as "p if q" (this should make sense: "p if q" says that if q happens, then p happens, and this is exactly what "q implies p" means). Another phrase that you will encounter a lot is

"p only if q".

"p only if q" clearly means the same thing as "q is necessary for p", and so (by the discussion earlier in this section) means the same as "p implies q".

It is for this reason that the bidirectional equivalence is often rendered into ordinary language as

"p if and only if q".

The "if" part captures "q implies p", and as we have just seen, the "only if" part captures "p implies q".

1.6 A collection of useful equivalences

We have seen that some pairs of superficially different looking statements are in fact the same, in the sense that their truth tables are identical. One example was

$$\text{"}\neg(p \wedge q)\text{" and "}(\neg p) \vee (\neg q)\text{".}$$

This says that whenever we encounter the negation of the conjunction of two things (p and q, in this case) in the middle of a proof, we can replace it with the disjunction of the negations (if that would be helpful).

This example might seem somewhat silly — once "$\neg(p \wedge q)$" and "$(\neg p) \vee (\neg q)$" are translated into ordinary language, it is hard to tell them apart! Indeed, "$\neg(p \wedge q)$" translates to "it is not the case that both p and q are true", while "$(\neg p) \vee (\neg q)$" translates to "it is the case that at least one of p and q is false". So it's unclear how much mileage we might get from replacing one with the other. (We will in fact see that this substitution is sometimes quite useful.)

A more substantial example is the pair, discussed earlier in the implication section,

$$\text{"}\neg(p \Rightarrow q)\text{" and "}p \wedge (\neg q)\text{".}$$

Suppose we are are in the middle of a proof, and we find ourselves working with the statement "p doesn't imply q". What can we do with this? Using the above equivalent pair, we can replace it with the statement "p and not q". This formally doesn't change anything, but the new statement is so different-*looking* from the old that we might well be able to use this new viewpoint to move the argument forward, in a way that we couldn't have done sticking with the old statement.

Much of the business of formulating proofs involves this kind of manipulation — substituting one expression for another, equivalent, one, and leveraging the change of viewpoint to make progress — and so it is useful to have an arsenal of pairs of equivalent statements at one's disposal. Here is a list of the most common pairs of equivalent statements, together with their common names. They can all be verified by constructing truth tables for each of the two pairs of statements, and checking that the truth tables have identical final columns. For the most part, it's not important to remember the specific names of these pairs.

Name	Pair of equivalent statements
Identity law	$p \wedge T$ and p $p \vee F$ and p
Domination law	$p \vee T$ and T $p \wedge F$ and F
Idempotent law	$p \vee p$ and p $p \wedge p$ and p
Double negation law	$\neg(\neg p)$ and p
Commutative law	$p \vee q$ and $q \vee p$ $p \wedge q$ and $q \wedge p$
Associative law	$(p \vee q) \vee r$ and $p \vee (q \vee r)$ $(p \wedge q) \wedge r$ and $(p \wedge q) \wedge r$
Distributive law	$p \vee (q \wedge r)$ and $(p \vee q) \wedge (p \vee r)$ $p \wedge (q \vee r)$ and $(p \wedge q) \vee (p \wedge r)$
De Morgan's law	$\neg(p \wedge q)$ and $(\neg p) \vee (\neg q)$ $\neg(p \vee q)$ and $(\neg p) \wedge (\neg q)$
De Morgan's law for n terms	$\neg(p_1 \wedge p_2 \wedge \cdots \wedge p_n)$ and $(\neg p_1) \vee (\neg p_2) \vee \cdots \vee (\neg p_n)$ $\neg(p_1 \vee p_2 \vee \cdots \vee p_n)$ and $(\neg p_1) \wedge (\neg p_2) \wedge \cdots \wedge (\neg p_n)$
Absorption law	$p \wedge (p \vee q)$ and p $p \vee (p \wedge q)$ and p
Tautology law	$p \vee (\neg p)$ and T
Contradiction law	$p \wedge (\neg p)$ and F
Equivalence law	$p \Leftrightarrow q$ and $(p \Rightarrow q) \wedge (q \Rightarrow p)$
Implication law	$p \Rightarrow q$ and $(\neg p) \vee q$
Implication Negation law	$\neg(p \Rightarrow q)$ and $p \wedge (\neg q)$
Contrapositive law	$p \Rightarrow q$ and $(\neg q) \Rightarrow (\neg p)$

Table 0: Common pairs of equivalent statements.

These pairs of equivalent statements can be used to form "proofs" that various other, more complex, pairs of statements are in fact equivalent. Indeed, one way to establish that statement X is equivalent to statement Y is to create a chain of statements, all of which are equivalent, with X at the start of the chain and Y at the end, each one obtained from the previous one by substituting out one expression for another, equivalent expression (either using a pair from the table above, or a pair that has already been established as being equivalent, either via a truth table, or by the method described here).

Here's an example. Suppose you want to show that $(p \wedge (p \Rightarrow q)) \Rightarrow q$ is equivalent simply to T (so, in the language we will introduce shortly, is a *tautology*), without using a truth table. In other words, you want to "reason out" that modus ponens is a valid logical inference, rather than "brute force" it. Here would be one possible approach:

$(p \wedge (p \Rightarrow q)) \Rightarrow q$	is equivalent to	$\neg(p \wedge (p \Rightarrow q)) \vee q$ (Implication law (or, definition of implication))
	which is equivalent to	$((\neg p) \vee \neg(p \Rightarrow q)) \vee q$ (De Morgan's law)
	which is equivalent to	$((\neg p) \vee (p \wedge (\neg q))) \vee q$ (Negation of implication)
	which is equivalent to	$(((\neg p) \vee p) \wedge ((\neg p) \vee (\neg q))) \vee q$ (Distributive law)
	which is equivalent to	$((p \vee (\neg p)) \wedge ((\neg p) \vee (\neg q))) \vee q$ (Commutative law)
	which is equivalent to	$(T \wedge ((\neg p) \vee (\neg q))) \vee q$ (Tautology law)
	which is equivalent to	$(((\neg p) \vee (\neg q)) \wedge T) \vee q$ (Commutative law)
	which is equivalent to	$((\neg p) \vee (\neg q)) \vee q$ (Identity law)
	which is equivalent to	$(\neg p) \vee ((\neg q) \vee q)$ (Associative law)
	which is equivalent to	$(\neg p) \vee (q \vee (\neg q))$ (Commutative law)
	which is equivalent to	$(\neg p) \vee T$ (Tautology law)
	which is equivalent to	T (Domination law).

This seems like overkill: it would have been much faster, *in this particular case*, to use a truth table. But as the number of simpler propositions involved grows, the truth table approach becomes less and less desirable. For example, with two propositions (as we have here) the truth table has only $2^2 = 4$ rows; but with 20 propositions, the truth table has $2^{2}0 \approx 1,000,000$ rows! At that point, it is completely impractical to use a truth table to verify an equivalence, and it is absolutely necessary to the pure reasoning illustrated in the example above.

1.7 Predicates

We've defined a statement to be an assertion that is either unambiguously true or unambiguously false. So "A human has walked on the moon" is a (true) statement, while "A Martian has walked in South Bend" is a (false) statement.

Returning to the mathematical world, what about something like "$x^2 + y^2 = 4$"? This is neither true nor false, because no specification has been made as to the values of x and y. If (x, y) is on the circle of radius 2 centered at the origin — say, $x = \sqrt{3}$, $y = -1$ — then the assertion is true, otherwise it is false.

A *predicate* is an assertion involving a variable or variables, whose truth or falsity is not absolute, but instead depends on the particular values the variables take on. So "$x^2 + y^2 = 4$"

is a predicate. Predicates abound in mathematics; we frequently are studying objects that depend on some parameter, and want to know for which values of the parameter some various assertions are true.

There are three ways in which predicates might be built up to become statements. One is by asserting an implication between predicates involving the same variables, of the form "if the first predicate is true, then the second must be also". Here's an example:

$$\text{"if } x - 2 = 1 \text{ then } x^2 - 9 = 0\text{".}$$

This is "p implies q" where p: "$x - 2 = 1$" and q: "$x^2 - 9 = 0$" are both predicates, not statements. It happens to be a true statement: either

- $x - 2 \neq 1$, in which case "not p" is true, so "(not p) or q" is true,

or

- $x - 2 = 1$, in which case $x = 3$, so $x^2 - 9 = 0$, so q is true, making "(not p) or q" true.

This way of turning a predicate into a statement is somewhat boring: "if $x - 2 = 1$ then $x^2 - 9 = 0$" amounts to saying "consider the predicate '$x^2 - 9 = 0$' in the specific situation where $x - 2 = 1$". Two (much) more interesting ways of turning a predicate into a statement involve the notion of quantifying over some universal set, which we now discuss.

Quantifiers

Predicates may also become statements by adding *quantifiers*. One quantifier, "for all", says that the predicate holds for all possible values. As an example consider the (false) statement

for every number n, n is a prime number.

We notate this statement as

$$(\forall n)p(n)$$

(read: "For all n, $p(n)$ (holds)") where $p(n)$ is the predicate "n is a prime number" — the "(n)" is added to p to indicate that p depends on the variable n. The formal reason the statement is false lies in the precise meaning that we assign to it: for any predicate $p(n)$, the statement "$(\forall n)p(n)$" is declared to be

- true, if $p(n)$ is true for *every* possible choice of n, and

- false, if there is even a single n for which $p(n)$ is not true.

The quantifier $\forall$ is referred to as the *universal quantifier*.

The *existential quantifier*, symbolically $\exists$, says that the predicate holds for *some* choice of the variable (but not necessarily for all of them). So with $p(n)$ as above, the true statement

$$(\exists n)p(n)$$

(read: "There exists n such that $p(n)$ is true") asserts that *some* number is prime. Formally, for any predicate $p(n)$, the statement "$(\exists n)p(n)$" is declared to be:

- true, if there is *at least one* n for which $p(n)$ is true, and
- false, if $p(n)$ is false for *every* n.

The universe of discourse

When one hears things like "for every n", or "there is an n", one should immediately ask "Where is one looking for n?" — the truth or otherwise of the associated statement may depend crucially on what the pool of possible n is. For example, consider the statement "There exists an x such that $x^2 = 2$" (or: "$(\exists x)r(x)$" where $r(x)$ is the predicate "$x^2 = 2$"). This is a true statement, if one is searching among real numbers — the value $x = \sqrt{2}$ witnesses the truth. On the other hand, if one is searching only among positive integers, then the statement becomes false — there is clearly no positive integer x with $x^2 = 2$. (Later, we'll talk about what happens if one is searching among rational numbers).

For this reason it is imperative, when using quantifiers, to know exactly what is the universe of possible values for the variable (or variables) of the predicate (or predicates) involved in the quantification. This is referred to as the *universe of discourse* of the variable. Usually, it is abundantly clear, from the context, what the universe of discourse is; if it is not clear, it needs to be made explicit in the quantification.

One way to make the universe of discourse explicit is to simply say what it is:

"With x running over positive real numbers, there exists x such that $x^2 = 2$"

or

"With the universe of discourse for x being positive real numbers,
there exists x such that $x^2 = 2$".

Another, more common way, is to build the universe of discourse into the quantification: e.g.,

"There exists a positive real x such that $x^2 = 2$".

Symbolically, this last statement could be written

$$\text{“}(\exists x \in \mathbb{R}^+)(x^2 = 2)\text{”}.$$

Here, "$\mathbb{R}^+$" is a standard notation for the positive real numbers (we'll see this later), and the symbol "$\in$" is the set theory notation for "is an element of" (so "$x \in \mathbb{R}^+$" conveys that x lives inside the set of positive real numbers). We will have more to say on basic set theory later.

This last method of building the universe of discourse into quantification is especially useful when a statement involves multiple variables, each with a different universe of discourse. Consider, for example, the following statement, which essentially says that one can find a rational number as close as one wants to any real number:

"for every real number x, for every positive real number ε, there is a rational number r that is within ε of x".

There are three variables — x, ε and r — each with a different universe of discourse. The above rendering of the statement is much cleaner than the (equivalent):

"With the universe of discourse for x being real numbers, for ε being positive real numbers, and for r being rational numbers, for every x and ε there is r such that r is within ε of x".

Symbolically, the statement we are discussing might be succinctly expressed as:

$$\text{``}(\forall x \in \mathbb{R})(\forall \varepsilon \in \mathbb{R}^+)(\exists r \in \mathbb{Q})(-\varepsilon < x - r < \varepsilon)\text{''}.$$

Here, "$\mathbb{R}$" is a standard notation for the real numbers, and "$\mathbb{Q}$" is a standard notation for the rational numbers. If it is absolutely clear from the context (as it will be throughout most of this course) that all variables are real numbers (that is, that all universes of discourse are either the set of real numbers, or subsets thereof), then we could also write

$$\text{``}(\forall x)(\forall \varepsilon > 0)(\exists r \in \mathbb{Q})(-\varepsilon < x - r < \varepsilon)\text{''}.$$

The statement we are discussing happens to be true, although proving it will involve a great deal of machinery. We will come to it towards the middle of the semester. Notice that we read quantifiers, as in ordinary reading, from left to right.

A predicate needs a quantifier for *every* variable to turn into a statement. Consider, for example,

$$\text{``for all } x,\ xy = 0\text{''}.$$

This is not a valid statement. It is true if y happens to be 0, and it is false if y is not 0, so its truth or falsity depends on the choice of y. On the other hand, both

$$\text{``for all } x, \text{ for all } y,\ xy = 0\text{'' and ``for all } x, \text{ there is a } y \text{ such that } xy = 0\text{''}$$

are statements (the first is false, the second is true).

Order of quantifiers

When a predicate with multiple variables gets turned into a statement by the addition of variables, the order in which we list the quantifiers is very important — typically a statement will change its meaning quite dramatically if we flip the order. For example, consider the predicate $p(m, n)$: "m is greater than n", or, more succinctly, "$m > n$". With the universe of discourse for all variables being the set of real numbers, the (true) statement

$$(\forall n)(\exists m)p(m, n), \quad \text{or} \quad (\forall n)(\exists m)(m > n)$$

says "For every number n, there is a number m such that m is greater than n". Flipping the order of the quantifiers leads to the false statement

$$(\exists m)(\forall n)(m > n),$$

or, "there is some number m such that every number is smaller than m".

For another example, let the variable x range over Major League baseball players, and the variable y range over Major league baseball teams. Let $p(x, y)$ be the predicate "player x is a shortstop for team y". Consider the following four statements which formally are similar-looking, but that translate into four *very* different statements in ordinary language:

- $(\forall x)(\exists y)p(x, y)$: for every player x, there is a team y such that x is the shortstop for y; in other words, every baseball player is *some* team's shortstop — false.
- $(\exists y)(\forall x)p(x, y)$: there is a team y such that every player x is the shortstop for y; in other words, there is some particular team such that *every* baseball player is that team's shortstop — false.
- $(\forall y)(\exists x)p(x, y)$: for every team y, there is a player x such that x is the shortstop for y; in other words, every team has a shortstop — true.
- $(\exists x)(\forall y)p(x, y)$: there is a player x such that every team y has x as its shortstop; in other words, there is some particular player who is *every* team's shortstop — false.

You should keep these absurd examples in mind as you work with quantifiers, and remember that it is very important to keep careful track of the order in which you introduce them — at least if you care about the meaning of the statements that you get in the end!

For a slightly more complicated example, here is what is called the *Archimedean principle* of positive real numbers:

"If N and s are positive numbers, there's a positive number t with $ts > N$."

(This is true no matter how big N is or how small s is.) We can take the predicate $p(N, s, t)$ to be "$ts > N$", and then encode the statement as $(\forall N)(\forall s)(\exists t)p(N, s, t)$. Note that this is implicitly assuming that we have agreed that we are working in the world of positive real numbers; if we were instead working in the world of all real numbers, we could write something like:

$$(\forall N)(\forall s)\left[((N > 0) \wedge (s > 0)) \Rightarrow (\exists t)((t > 0) \wedge p(N, s, t))\right]$$

(which we might read as, "For every N and s, if N and s are positive, then there is a t such that both of the following hold: t is positive and $ts > N$").

Negation of predicates

The operation of negation has a very simple effect on quantifiers: it turns $\forall$ into $\exists$ and vice versa, while bringing the negation inside the predicate. Think about $(\forall x)p(x)$ and $(\exists x)(\neg p(x))$, for example. If the first is true then $p(x)$ holds for every x, so $\neg p(x)$ holds for no x, so the second is false, while if the first is false then there is an x for which $p(x)$ doesn't hold, and so the second is true. This argument shows that

$$\neg((\forall x)p(x)) \quad \text{is equivalent to} \quad (\exists x)(\neg p(x)),$$

and similarly we can argue that

$$\neg((\exists x)p(x)) \quad \text{is equivalent to} \quad (\forall x)(\neg p(x)).$$

If the universe of discourse for the variable x is finite, then the two equivalences above are just DeMorgan's laws, rewritten. Indeed, if the possible choices for x are $x_1, x_2, \ldots, x_n$, then $(\forall x)p(x)$ is the same as $p(x_1) \wedge p(x_2) \wedge \cdots \wedge p(x_n)$, and so by DeMorgan's law the statement $\neg((\forall x)p(x))$ is the same as $(\neg p(x_1)) \vee (\neg p(x_2)) \vee \cdots \vee (\neg p(x_n))$, which is just another way of saying $(\exists x)(\neg p(x))$. Similarly, one can argue that $\neg((\exists x)p(x))$ means the same as $(\forall x)(\neg p(x))$. So negation of quantifiers can be thought of as DeMorgan's law generalized to the situation where the number of predicates being and-ed or or-ed is not necessarily finite.

What about a statement with more quantifiers? Well, we can just repeatedly apply what we have just established, working through the quantifiers one by one. For example, what is the negation of the statement $(\exists x)(\exists y)(\forall z)p(x, y, z)$?

$$\begin{array}{rcl} \neg((\exists x)(\exists y)(\forall z)p(x,y,z)) & \text{is equivalent to} & (\forall x)(\neg((\exists y)(\forall z)p(x,y,z))) \\ & \text{which is equivalent to} & (\forall x)(\forall y)(\neg((\forall z)p(x,y,z))) \\ & \text{which is equivalent to} & (\forall x)(\forall y)(\exists z)(\neg p(x,y,z)). \end{array}$$

The homework will included some more involved examples, such as negating the Archimedean principle (and interpreting the result).

1.8 Tautologies

A *tautology* is a statement that is built up from various shorter statements or quantified predicates p, q, r, …, that has the property that no matter what truth value is assigned to each of the shorter statements, the compound statement is true.

A simple example is "There either is life on Mars, or there is not", which can be expressed as $p \vee (\neg p)$ where p is the statement "There is life on Mars". If p is true then so is $p \vee (\neg p)$, while if p is false then $\neg p$ is true, so again $p \vee (\neg p)$ is true. In general the tautology $p \vee (\neg p)$ is referred to as the *law of the excluded middle* (there is no middle ground in logic: either a statement is true or it is false).

Looking at the truth table of the bidirectional implication $\Leftrightarrow$, it should be evident that if p and q are any two compound statements that are build up from the same collection of shorter statements and that have the same truth tables, then $p \Leftrightarrow q$ is a tautology; so for example,

$$\neg(p \vee q \vee r) \Leftrightarrow (\neg p) \wedge (\neg q) \wedge (\neg r)$$

is a tautology. A tautology can be thought of as indicating that a certain statement is true, in an unqualified way; the above tautology indicates the truth of one of De Morgan's laws.

An important tautology is

$$(p \wedge (p \Rightarrow q)) \Rightarrow q$$

(easy check: this could only be false if q is false and both p and $p \Rightarrow q$ are true; but if q is false and p is true, $p \Rightarrow q$ must be false; so there is no assignment of truth values to p and q that makes the compound statement false). This tautology indicates the truth of the most basic rule of logical deduction, that if a statement p is true, and it is also true that p implies another statement q, then it is correct to infer that q is true. This is called *modus ponens*.

For example, if you know (or believe) the truth of the implication

"If I go out in the rain without a coat on, I will get wet"

and I also give you the information that I go out in the rain without a coat on, then it is legitimate for you to reach the conclusion that I get wet.

The ideas raised in this short section are fundamental to mathematics. A major goal of mathematics is to discover *theorems*, or statements that are true. In other words, a theorem is essentially a tautology. Most complex theorems are obtained by

- starting with a collection of simpler statements that are already known to be true (maybe because they have already been shown to be true, or maybe because they are assumed to be true, because they are axioms of the particular mathematical system under discussion), and then
- deducing the truth of the more complex statement via a series of applications of rules of inference, such as modus ponens.

This process is referred to as *proving* the theorem. Almost every result that we use in this class, we will prove; this is something that sets Math 10850/60 apart from courses such as Math 10550/60 (Calculus 1/2), which explain and apply the techniques of calculus, without laying a rigorous foundation. The notion of proof will be explored in more detail in Section 2.