1 Introduction

Propositional logic is only the *beginning* of the logic we need for mathematics and science. Propositional logic is quite poor in its expressive capacity, since the only way to fit statements to a specific area is to provide an interpretation to the propositional symbols, much like we did in Lecture 19 in proving statements about graph coloring and König's lemma. The key ingredient missing is that we have no way to make general statements. For example, using our coding trick into propositional symbols, we can prove for any *specific* planar graph that it is four-colorable. However, we have no way of even stating that all planar graphs are four-colorable. (A planar graph is one which can be drawn on a sheet of paper so that edges only cross each other at vertices. Kenneth Appel and Wolfgang Haken proved in 1976 that the vertices of any planar graph can be colored by using at most four colors so that no adjacent vertices have the same color.)

Being able to generalize and reason about objects generally is crucially important in mathematics and science. For example, in arithmetic we can make statements like $1 + 2 = 2 + 1$ and $2 + 3 = 3 + 2$ and $27 + 56 = 56 + 27$, but we can more succinctly simple state the commutative law: $x + y = y + x$ for all numbers, and then apply this statement when we need it:

- $x + y = y + x$ for all numbers.
- 27 and 56 are two numbers, so set $x = 27$ and $y = 56$.
- Therefore, $27 + 56 = 56 + 27$.

The science of human mortality has general laws, for example, that all men are mortal, generalized from aeons of human experience. This discipline is organized around such laws, which can then be applied to particular circumstances:

- All men are mortal.
- Socrates is a man.
- Therefore, Socrates is mortal.

The last bit of reasoning was captured in an even more general form by Aristotle two thousand years ago, by means of the syllogism:

- All A are B.
- s is an A.
- Therefore, s is a B.

A and B are symbols which stands for a property, A=man and B=mortal in the previous argument, and s is a symbol which stands for a person, s=Socrates in the previous argument. It is harder to see how to generalize the first argument using the commutative law, but it seems to have the same form: reasoning from a general truth to a particular instance. First-order logic will provide us with the facility to express general statements and identify the above arguments as valid forms of inference.

2 The Island of Knights and Knaves

Recall that on the Island of Knights and Knaves, the following three laws hold:

1. Everyone is a knight or knave.
2. Knights always speak truly.

3. Knaves always speak falsely.

Up till now our natives were very limited in what they could say. For example, if Grogg was a native of the island, he could say things like

- Nogg and Phogg are of the same type.
- Nogg, Phogg and Trogg are of the same type.
- Nogg and Phogg and Trogg and Yipp and Yapp and etc. are all of the same type.

Since there are only finitely many inhabitants of the island, Grogg could continue naming them off in this way (if he knew everyone!), but the inconvenience is probably not worth the effort. If this were the limits of the expressive capabilities of the language of the Island, I doubt mathematics and science would arise, since the efforts to express general truths would not be worth the cost of writing them down.

Example 19.2.1 No native of the island can say of themselves, "I am a knave." If a knight said this, he would be lying and if a knave said this he would be telling the truth. Could anyone say

- "All of us are knaves."

Such a native would be calling himself a knave, along with every other inhabitant. If there were only one native, this statement means the same thing as "I am a knave.", which no native could say. But if there are two or more natives, the statement is more inclusive. No knight could say this, since it is not true. But a knave could say this, provided there was at least one knight on the island. So, it would not be true that every inhabitant was a knave.

Example 19.2.2 Suppose a native says

- "Some of us are knaves."

A knave could not say this, since he would then be speaking truly. But a knight could say this, if there were at least one knave on the island.

A native could express the statements "All of us are knaves." and "Some of us are knaves." using propositional logic, provided like Grogg they knew all the inhabitants:

- "All of us are knaves." : Nogg is a knave AND Phogg is a knave AND Trogg is a knave AND Yipp is a knave AND Yapp is a knave etc.
- "All of us are knaves." : Nogg is a knave OR Phogg is a knave OR Trogg is a knave OR Yipp is a knave OR Yapp is a knave etc.

We could a introduce special proposition A which means "All of us are knaves." and another proposition E which means "Some of us are knaves.". This strategy is only of very limited utility.

Suppose you come to the island with the intent of studying the inhabitants and to determine whether there is a correlation between lying and various kinds of behavior on the island. Recall that if an inhabitant A makes a statement $\mathcal{S}$, then the propositions $K(A) \leftrightarrow \mathcal{S}$ is true, where $K(A)$ is the proposition that A is a knight. That is the true value of A's statement is the same as the proposition that they are a knight. The propositional symbol $K(A)$ looks awkward, but it is useful to record the speaker with their statement, since you may gather the statements of many speakers.

Example 19.2.3 Suppose you survey each of the inhabitants to learn something about their more general views on the proportion of knights and knaves on the Island. Curiously, each gave the same response:

- "All of us are of the same type."

What can you deduce about the composition of the inhabitants of the island.

Everyone is making the same statement, so they must all be of the same type. Since the statement made is true, everyone must be a knight.

In these examples is you will need to record the information that everyone made the same statement, and try to deduce the consequences of this fact. Let A_1 be the statement that "All of us are of the same type." . The problem you would face, as a propositional logician, is that drawing inferences would be cumbersome. You might record in your notebook the following statement:

$$(K(\text{Grogg}) \leftrightarrow A_1) \wedge (K(\text{Nogg}) \leftrightarrow A_1) \wedge (K(\text{Phogg}) \leftrightarrow A_1) \wedge \ldots,$$

and if the number of inhabitants is finite, you could express the statement that all inhabitants said the same thing. But it does not follow, based on propositional logic, that every inhabitant is a knight, which we denoted by the propositional symbol A.

First-order logic gives us the expressive capability to record more finely what each inhabitant tells us, while allowing us to reason about them based on their responses. The key is that first-order logic dissects statements such as "All of us are knights." into the parts "All of us" + "...are knights." and "Some of us" + "...are knights." The expressions "All of us" and "Some of us" are called *quantifiers*, which is the new logical component added by first-order logic. The expression "...are knights." is called a *predicate* is a second component added to first-order logic. The meaning of the predicates will depend on the context, much as propositional symbols in propositional logic. The meaning of the quantifiers is (largely) fixed by first-order logic to allow us general principles of reasoning, much as we fixed the meaning of the connectives $\{\bot, \top, \neg, \wedge, \vee, \rightarrow, \leftrightarrow\}$ in propositional logic.

The reason for the qualification the meaning of quantifiers "All of us" and "Some of us" are *largely* fixed by first-order logic can be seen in the following example.

Example 19.2.4 The island of knights and knaves has a sister island called Island II (the island of knights and knaves is called Island I). You travel to Island II to study its inhabitants. On this island every inhabitant gives the response:

- "Some of us are knights and some of us are knaves."

What can you deduce about the composition of the inhabitants of this island.

Again, everyone is making the same statement, so must be of the same type. Since the statement made is false, everyone must be a knave.

The inhabitants of our two islands may make identical statements, "All of us are of the same type." and "Some of us are knights and some of us are knaves.", but they are not saying the same thing, if the first group is talking exclusively about themselves and the second group is talking exclusively about themselves as well. The statements are superficially the same, but have different meanings. Still, as researchers studying each group, we reason in exactly the same way regardless of which group is making the assertion. The solution to this difficulty in first-order logic is to recognize that while the *logical properties* of the quantifiers "All of us" and "Some of us" should be fixed, we should be able to vary the *domain* of the quantifiers from context to context. On one island "All of us" means "All of us on this Island I", while on the sister island "All of us" means "All of us on Island II".

So, first-order logic dissects propositional symbols into *quantifiers* (the logical component) and *predicates* (the nonlogical component). The glue which binds quantifiers to predicates are played by *variables*. We have been using variables all along, not in the language of propositional logic itself, but in talking about propositional logic. You have come to understand α, β, and γ to be variables representing "arbitrary" propositions and Γ to be a variable representing sets of propositions. We have also used P, Q and R as variables for propositional symbols. We have also introduced variables like τ for tableaux (and sometimes natural deduction proofs) and π for paths through tableaux. I have tried to be consistent about the different shapes of variables so you can come to identify what *domain* they refer to. So α is a variable which ranges over propositions, Γ is a variable which ranges over sets of propositions and τ is a variable which ranges over tableaux.

What a variable "means" over and above the fact that it comes from a domain, depends upon the context in which it is used. This is actually intimately tied to the quantifier. To make this more concrete, lets replay the reasoning from 19.2.3. You record the responses of each inhabitant and note that

- All of the inhabitants of Island I said "All of us are of the same type."

Let x be an arbitrary inhabitant of Island I. Then I know the following must be true

(*) x is a knight if and only if what x said was true: that "All of us are of the same type."

Since every inhabitant said the same thing, they must be of the same type. Let's slow this bit of reasoning down: let y be another arbitrary inhabitant of Island I, then I know the following must be true

(**) y is a knight if and only if what y said was true: that "All of us are of the same type."

It now follows from (*) and (**) by *propositional logic* that

x is a knight if and only if y is a knight.

Since x and y are arbitrary inhabitants of Island (and am not assuming x is Grogg and y is Phogg), it follows that every inhabitant is of the same type. Our reasoning is based on the following valid inferences

1. All of the inhabitants of Island I said "All of us are of the same type.", and since x is an inhabitant of Island I, x also said "All of us are of the same type.". (And the same for y.)
2. Since we have shown that : x is a knight if and only if y is a knight, and we have only assumed that x and y were inhabitants of Island I, it follows that ALL of the inhabitants of Island I are of the same type.

We can make the same argument on Island II, except that the inhabitants made the statement "Some of us are knights and some of us are knaves.", and we may even use the same variables x and y, but we understand that the role of x and y has changed as the *domain* of our quantifier has changed from Island I to Island II.

The last paragraph introduces deep and subtle features about our use of variables. When we wish to reason generally about the inhabitants of Island I we introduce a variable such as x, which represents an arbitrary inhabitant of Island I. If we wish to draw a universal truth about all inhabitants on Island I, we must assume no additional property about x other than those that hold of all inhabitants of the island. In our previous example we may infer that x said "All of us are of the same type.", but we may not infer that x is a smoker if we wish our conclusion to hold universally for all inhabitants of Island I. Contrast this with our use of proper names: if Grogg and Phogg both make the statement that "All of us are of the same type.", we cannot infer that everyone on Island I made the same statement, without introducing further facts, such as Grogg and Phogg are the only inhabitants of Island I. Furthermore, there may be specific properties of Grogg and Phogg that are relevant in drawing inference about them, specifically, but which do not hold universally, such as that Grogg smokes but Phogg does not.

Sometimes we want to infer general statements about a subpopulation.

Example 19.2.5 In this example, the aim is to discover whether there is any correlation between lying and smoking on Island I. All smokers on the island have been interviewed about their views, and all have made the same claim: "All smokers are of the same type." What can we deduce about the smoking population.

Since every *smoker* has made the same claim, we can deduce that all *smokers* really are of the same type. The reason for this as follows: let x and y be smokers on Island I, then

1. x is a knight if and only if the statement "All smokers are of the same type." is true and y is a knight if and only if the statement ' 'All smokers are of the same type." is true.
2. So, x is a knight if and only if y is a knight, that is they are of the same type. But since x and y are arbitrary smokers on Island I (and nothing more is assumed about them), it follows that ALL smokers on Island I are of the same type.

Now that we know all smokers are of the same type, they must all be knights, because the statement that each made is true.

The inference that all smokers on Island I are of the same type holds only for smokers on Island I, nonsmokers may be of different types. The reason the inference must be restricted to smokers is that by introducing the variables x and y to make the inference we assumed they were smokers, so any general inferences we draw are restricted to the subpopulation of smokers.

First-order logic will split propositional atoms by introducing *properties* and allows us to reason generally about the properties of members of a population by introducing quantifiers and variables. We now turn to look at how first-order logic expresses general statements.

3 Predicate and Quantifiers

In first-order logic we will use lower case letters x, y and z to stand for arbitrary objects of the domain under discussion. What the domain is depends on the application. For example, if we are studying the inhabitants of Island I, the domain will be the people on the Island and the letters x, y and z will stand for arbitrary inhabitants; if we are doing number theory, the domain will be numbers and x, y and z will stand for arbitrary numbers. It is this generality in application of first-order that makes it so useful.

Given any property of objects in the domain, we may introduce a *predicate symbol* P and for any object x from the domain we write Px to symbolize the proposition that x has the property which P refers to. For example, let S be the property of being a smoker, then Sx is the proposition that x is a smoker. We will also have names of objects in our domain, and these will be symbolized by lower case letters from the beginning of the alphabet. If c names Grogg, the Pc is the proposition that Grogg smokes.

Suppose we wish to say that *every* object x has property P; first-order logic has the symbol $\forall$–called the *universal quantifier*–which stands for "all" or "every". We write $(\forall x)$, which you can read as "for all x" or "for every x", so the statement $(\forall x)Px$ is rendered in English as "Everything has the property P.", or in mathematical English (where variables are commonplace), "For all x, Px."

There is a second quantifier in first-order logic which corresponds to the English word "some" and which is symbolized by $\exists$–the *existential quantifier*. We write $(\exists x)$, which you can read as "some x" or "there exists an x", so the statement $(\exists x)Px$ is rendered in English as "Something has property P." or in mathematical English as "There exists an x such that Px."

In addition to predicate, variable, names (or *constants*) and quantifiers, first-order logic has the propositional connectives $\neg, \wedge, \vee, \rightarrow$. These connectives interact with quantifiers in new ways. For example, let x, y and z be variable which stand for unspecified people, and S the property of smoking. Since Sx is a proposition (a statement which can be true or false), we can write $\neg Sx$ for x does not smoke, that is, x is a *nonsmoker*. So, the propositional connectives provide a way of constructing new predicates which represent properties of the members of the domain.

We have seen $(\forall x)Sx$ says "Everyone smokes." and $(\exists x)Sx$ says that "Someone smokes." We can also form $\neg(\exists x)Sx$, which says that "No one smokes." (it is *not* there case there there exists a smoker.) An alternative way of saying this is $(\forall x)\neg Sx$ (for every person x, x is not a smoker.)

Let K be the predicate that represents the property of being a knight, so that Kx says that x is a knight. How can we say "All knights are smokers."? We can restate this as: "For all x, if x is a knight, then x is a smoker.", and hence symbolize it by $(\forall x)(Kx \rightarrow Sx)$. Note that although I have been using the variable x, the statement $(\forall y)(Ky \rightarrow Sy)$ means the same thing: all knights are smokers. How can we say that "*Only* knights are smokers."? One way of rendering this is $(\forall x)(Sx \rightarrow Kx)$ (for all people x, if x is a smoker, then x is a knight.) Another way of rendering this is as $(\forall x)(\neg Kx \rightarrow \neg Sx)$ (for all x, if x is not a knight, then x is not a smoker – or on the island of knights and knaves, no knave is a smoker: for all x, if x is a knave then x is a nonsmoker.) You may recognize that the second form is the contrapositive, in a sense of the first form. The tautological truths of propositional logic will also hold in first-order logic. Since Kx and Sx are propositions (capable of being true or false), $Sx \rightarrow Kx$ is equivalent to $Kx \rightarrow Sx$. The technique introduced here also provides a way of making general statements about a subpopulation. If we want to say something about all smokers, such as all smokers are knights, then we use the conditional to restrict the domain of the quantifier $\forall$: $(\forall x)(Sx \rightarrow Kx)$. If P represents some arbitrary property, then we render the general statement about smokers "All smokers also have the property P." as $(\forall x)(Sx \rightarrow Px)$.

How would we say "Some smokers are knights."? We use the existential quantifier: there is a person x who is a smoker and is also a knight, $(\exists x)(Sx \wedge Kx)$. This technique allows us to make statements about the existence of members of a subpopulation. If we want to say "Some smokers are knaves" we use conjunction with the predicate S to restrict the general domain to that of smokers, $(\exists x)(Sx \wedge \neg Kx)$.

How would we say "No smoker is a knight."? We can do this with the universal quantifier: for all x, if x is a smoker, then x is not a knight, $(\forall x)(Sx \rightarrow \neg Kx)$. On the island of knights and knaves, we could also translate this as "*Only* knaves are smokers." Alternatively we can use the existential quantifier, it is not the case that there is a smoker who is not a knight, $\neg(\exists x)(Sx \wedge Kx)$. If you compare the two renderings, you notice that inside each quantifier we have the statements $(Sx \rightarrow \neg Kx)$ and $(Sx \wedge Kx)$, and these are related, since $\neg(Sx \rightarrow \neg Kx)$ is tautologically equivalent to $(Sx \wedge Kx)$. In first-order logic the two statements $(\forall x)(Sx \rightarrow \neg Kx)$ and $\neg(\exists x)(Sx \wedge Kx)$ are equivalent, although this equivalence is not quite captured by *tautological equivalence*. The connection is as follows

- $\neg(\exists x)Px$, there is no x such that Px, is equivalent to $(\forall x)\neg Px$, for every x, it is not the case that Px; and
- $\neg(\forall x)Px$, it is not the case that for every x, Px, is equivalent to $(\exists x)\neg Px$, there is some x such that Px.

We summarize the techniques we have observed. They are very useful in rendering English statements into first-order logic. Let P and Q be any predicates.

1. $(\forall x)(Px \rightarrow Qx)$, for every x, if x has P, then x has Q. This allows to restrict the domain to a subdomain, as in "All P's are Q's."
2. $(\forall x)(Px \rightarrow \neg Qx)$, for every x, if x has P, then x does not have Q. Equivalently, "*No* P is a Q."
3. $(\forall x)(Qx \rightarrow Px)$, for every x, if x has Q, then x has P. Equivalently, "*Only* P's are Q's."
4. $(\exists x)(Px \wedge Qx)$, there exists an x such that Px and Qx. Equivalently, "Some P's are Q's.
5. $\neg(\exists x)(Px \wedge Qx)$, it is not the case that there is an x such that x has P and x has Q. Equivalently, "No P is a Q." This is also equivalent to (2).