

CHAPTER 3

First(-order) things first

With the ghost of George Boole finally satisfied, you're now ready to try your luck at sleeping once again. As you close your eyes, a wild amount of noise starts coming in through the window of your room. You get up, look outside, and you see Tarski's tank going up and down the street. Confused and choosing to ignore the (logical) consequences, you decide to think a bit more about what you just learned. All the limitations of the Boolean world become suddenly apparent...

1. Syntax still means how we write things down

1.1. First steps. Let's start with our basic symbols. We already learned about some of them in the previous chapter. But now, as we saw previously, we will introduce more, namely, we will introduce **quantifiers**:

- \exists (syntactically “rotated E ”, semantically “exists”);
- \forall (syntactically “rotated A ”, semantically “for all”);

But we also need things to quantify *over*, things like the property of mortality or of humanity or whatever else. Since we're trying, at the end of the day, to also do some mathematics, we will also think about functions.

Definition 1.1.1. A *first-order language* is a set of symbols \mathcal{L} which consists of two disjoint subsets:

- (1) The *logical symbols* of \mathcal{L} :¹
 - A countable set of *variables* Var , which we will usually denote by $x, x_1, x_2, \dots, y, y_1, y_2, \dots, z, z_1, z_2, \dots$, etc.
 - Our good old logical symbols: $\wedge, \vee, \neg, \rightarrow, \leftrightarrow, \forall, \exists$.
 - A symbol for equality \doteq .

¹This part is appears unchanged in EVERY first-order language \mathcal{L} .

(2) The *signature* of \mathcal{L} (aka the *non-logical symbols*) of \mathcal{L} , these are the following three disjoint sets:

- A set of *constant symbols* $\text{Const}(\mathcal{L})$. We will usually denote (abstract) constant symbols as $\underline{c}, \underline{c}_1, \underline{c}_2, \dots$
- A set of *relation symbols* $\text{Rel}(\mathcal{L})$. We will usually denote (abstract) relation symbols as $\underline{R}, \underline{R}_1, \underline{R}_2, \dots$
- A set of *function symbols* $\text{Fun}(\mathcal{L})$. We will usually denote (abstract) function symbols as $\underline{f}, \underline{f}_1, \underline{f}_2, \dots$

As part of the signature, we also have an *arity* function

$$\text{arity}_{\mathcal{L}} : \text{Rel}(\mathcal{L}) \sqcup \text{Fun}(\mathcal{L}) \rightarrow \mathbb{N}_{\geq 1}.$$

Let $\underline{R} \in \text{Rel}(\mathcal{L})$. If $\text{arity}(\underline{R}) = n$ then we say that \underline{R} is an *n-ary relation symbol*. Similarly, if $\underline{f} \in \text{Fun}(\mathcal{L})$. and $\text{arity}(\underline{f}) = n$ then we say that \underline{f} is an *n-ary function symbol*.

To make our lives a bit simpler and avoid transfinite things and Zorn, we will, unless otherwise stated, assume that $|\text{Const}(\mathcal{L})|, |\text{Rel}(\mathcal{L})|, |\text{Fun}(\mathcal{L})| \leq \aleph_0$, i.e. that our **signatures are countable**. I'm underlying things here to make sure you understand that these are not constants relations or functions in the sense you may be familiar, these are constant relation or function symbols.²

Again, just like when we were talking about the syntax of propositional logic, everything we will discuss in this section is purely syntactic; it's a bunch of strings of letters on a page (or a board) and will have NO MEANING yet.

Ah also, since for any first-order language \mathcal{L} , the logical symbols of \mathcal{L} are always the same, we will be identifying a language with its signature.

Our first goal is to learn how to write formulas in a given first-order language \mathcal{L} . It'll take a minute, but we'll get there. Many things that we will be discussing in this section, have already shown up here and there when we were playing around with propositional logic (cf. the definition of \mathcal{F} -terms).

Definition 1.1.2. Let \mathcal{L} be a first-order language. The *terms* of \mathcal{L} (or \mathcal{L} -terms) are defined inductively, as follows:

- Every variable is an \mathcal{L} -term.

²Slowly as we understand the interplay between syntax and semantics (next section), I may every now and again start dropping the underlines, and probably so will you, but for now I'd like to keep them here.

- (2) Every constant symbol $\underline{c} \in \text{Const}(\mathcal{L})$ is an \mathcal{L} -term.
- (3) If $\underline{f} \in \text{Fun}(\mathcal{L})$ is an n -ary function symbol and t_1, \dots, t_n are \mathcal{L} -terms, then $\underline{f}(t_1, \dots, t_n)$ is an \mathcal{L} -term.
- (4) That's it – all terms are constructed by finitely many applications of (1), (2) and (3).

Remark 1.1.3. Relation symbols don't show up in term building!

Lemma 1.1.4. Let \mathcal{L} be a first-order language. Let \mathcal{T}_0 be the set $\text{Var}(\mathcal{L}) \cup \text{Const}(\mathcal{L})$. Inductively, define \mathcal{T}_{n+1} to be the set:


$$\mathcal{T}_{n+1} := \mathcal{T}_n \cup \{\underline{f}(t_1, \dots, t_k) : k \in \mathbb{N}, \underline{f} \in \text{Fun}(\mathcal{L}), \text{arity}(\underline{f}) = k, t_i \in \mathcal{T}_n\},$$

and let $\mathcal{T} = \bigcup_{n \in \mathbb{N}} \mathcal{T}_n$. Then \mathcal{T} is the set of all \mathcal{L} -terms.

PROOF. It should hopefully be clear that the set of all \mathcal{L} -terms is contained in \mathcal{T} . For the other inclusion, we show, by induction on $n \in \mathbb{N}$ that \mathcal{T}_n is contained in the set of \mathcal{L} -terms. Indeed this is clear for \mathcal{T}_0 and follows by (3) for $n \geq 1$. \square

By the lemma above, the set of all \mathcal{L} -terms is $\mathcal{T} = \bigcup_{n \in \mathbb{N}} \mathcal{T}_n$. Observe that if $t \in \mathcal{T}$, then there is some minimal $n \in \mathbb{N}$ such that $t \in \mathcal{T}_n$. We call this n the **height** of t (some people use the word **length**, I may do so too from time to time). More down to earth, the length of any variable or constant symbol is 0 and the length of any other term is precisely the longest chain of function symbols that appears in it. Here is an example:

Example 1.1.5. Let \mathcal{L} be a first-order language with a single **binary** (i.e. of arity 2) function symbol \underline{g} and a single constant symbol \underline{c} . To compute the length of a term, we can break it up into a tree, as follows:



The term above has height 3.

This isn't all that complicated, really, but to make sure you get it, you can try out the following:

Exercise 1.1.6. Let \mathcal{L} be a language with a single constant symbol \underline{c} , one **unary** (i.e. of arity 1) function symbol \underline{f} and one binary function symbol \underline{g} . Write down all \mathcal{L} -terms of length at most 3 in the variables x and y .

Just like matter is built of atoms and stuff, formulas are built out of inseparable components, the atoms,³ so to speak, of first-order syntax, which we define below:

Definition 1.1.7. The *atomic formulas* of \mathcal{L} (or *atomic \mathcal{L} -formulas*) are defined inductively as follows:

- (1) If t_1, t_2 are \mathcal{L} -terms, then $t_1 \doteq t_2$ is an atomic \mathcal{L} -formula.⁴
- (2) If $\underline{R} \in \text{Rel}(\mathcal{L})$ is an n -ary relation symbol and t_1, \dots, t_n are terms, then $\underline{R}(t_1, \dots, t_n)$ is an atomic \mathcal{L} -formula.

Observe that, whereas in the definition of terms we mixed smaller terms to build bigger ones, in the definition of atomic formulas, we don't mix atomic formulas to build bigger ones. Atomic formulas are atoms indeed! They will form our base case in the definition of formulas proper.

Indeed, in the definition of formulas, atomic formulas will (up to a point of analogy obscuring reality) play the role that propositional variables played in the definition of propositional formulas.

Definition 1.1.8. The *formulas* of \mathcal{L} (or \mathcal{L} -formulas) are defined inductively, as follows:

- (1) Every atomic \mathcal{L} -formula is an \mathcal{L} -formula.
- (2) If ϕ and ψ are \mathcal{L} -formulas, and $x \in \text{Var}$, then:

(a) $(\phi \wedge \psi)$	(c) $(\phi \rightarrow \psi)$	(e) $(\forall x)\phi$
(b) $(\phi \vee \psi)$	(d) $(\neg\psi)$	(f) $(\exists x)\phi$

are all \mathcal{L} -formulas

- (3) That's it – all \mathcal{L} formulas are constructed by finitely many applications of (1) and (2), above.

³From α - (a-, “not”) + $\tau\epsilon\mu\nu\omega$ (témno, “I cut”). Bet you weren’t expecting a lesson in etymology in these notes.

⁴Note here that we use \doteq because these are syntactic objects! Just like with the underlining business, as we go along, I may start dropping the dot above the equality symbol.

In analogy with Lemma 1.1.4 we have the following:

Exercise 1.1.9 (Unique reading of formulas). Let \mathcal{L} be a first-order language. Let \mathcal{F}_0 be the set of atomic \mathcal{L} -formulas. Inductively, define \mathcal{F}_{n+1} to be the set:

$$\mathcal{F}_{n+1} := \mathcal{F}_n \cup \{\phi \wedge \psi, \phi \vee \psi, \phi \rightarrow \psi, \neg\phi, (\forall x)\phi, (\exists x)\phi : \phi, \psi \in \mathcal{F}_n, x \in \text{Var}\}$$

and let $\mathcal{F} = \bigcup_{n \in \mathbb{N}} \mathcal{F}_n$. Prove that \mathcal{F} is the set of all \mathcal{L} -formulas.

This allows us to define the **height** of an \mathcal{L} -formula, as the minimal $n \in \mathbb{N}$ such that $\phi \in \mathcal{F}_n$. Formulas are built in trees, just like terms. Below are some examples that we could certainly care about:

Example 1.1.10.

- (1) The **language of graphs**, $\mathcal{L}_{\text{graph}}$, consists of a single binary relation symbol \underline{E} .
- (2) The **language of groups**, \mathcal{L}_{grp} , consists of a constant symbol $\underline{1}$, a binary function symbol $\underline{\times}$ (usually we write $x\underline{\times}y$ for $\underline{\times}(x, y)$, because we are not maniacs), and a unary function symbol $\underline{\neg 1}$ (again for similar reasons, we usually write $x\underline{\neg 1}$ for $\neg 1(x)$).
- (3) The language of **Peano arithmetic**, $\mathcal{L}_{\text{Peano}}$ (which will be our main focus in the last chapter) consists of a constant symbol $\underline{0}$, a unary function symbol \underline{S} , and two binary function symbols $\underline{+}$ and $\underline{\times}$.

Exercise 1.1.11.

- (1) Write two different $\mathcal{L}_{\text{graph}}$ -formulas of height 4 in this language.
- (2) Write down two \mathcal{L}_{grp} -formulas that use all symbols in the language at least twice.

1.2. Syntax can be annoying, I. In this section, we will handle some of the syntactical quirks of first-order logic. Thankfully, we have gotten some practice with annoying syntactical quirks from our experience with propositional logic.

First of all, just like in the case of propositional logic, formulas have variables. Let's give a preliminary definition (which should remind us of a definition we've seen before). This time it will take a bit longer to state.

Definition 1.2.1. Let \mathcal{L} be a first-order language and t an \mathcal{L} -term. We define the *set of variables of t* , denoted $\text{Var}(t)$, as follows:

- (1) If t is the variable x then $\text{Var}(t) = \{x\}$.
- (2) If t is a constant symbol, then $\text{Var}(t) = \emptyset$.
- (3) If t is of the form $\underline{f}(t_1, \dots, t_n)$ for some n -ary function symbol $\underline{f} \in \text{Fun}(\mathcal{L})$ and \mathcal{L} -terms t_1, \dots, t_n , then $\text{Var}(t) = \bigcup_{i \leq n} \text{Var}(t_i)$.

We say that t is *closed* if $\text{Var}(t) = \emptyset$.

That's it with terms. Now, for formulas, as you may have expected, we have to start with the atomic guys and build from there.

Definition 1.2.2. Let ϕ be an atomic \mathcal{L} -formula. Then, the *set of variables of ϕ* , again denoted $\text{Var}(\phi)$,⁵ is defined as follows:

- (1) If ϕ is of the form $t_1 \doteq t_2$ for \mathcal{L} -terms t_1, t_2 then $\text{Var}(\phi) = \text{Var}(t_1) \cup \text{Var}(t_2)$.
- (2) If ϕ is of the form $\underline{R}(t_1, \dots, t_n)$ for some n -ary relation symbol $\underline{R} \in \text{Rel}(\mathcal{L})$ and \mathcal{L} -terms t_1, \dots, t_n , then $\text{Var}(\phi) = \bigcup_{i \leq n} \text{Var}(t_i)$.

Before we move up a level, to formulas proper, we need to make an important distinction. We want to be able to tell apart if a variables x in a formula appears in some quantifier, say $(\forall x)$ or $(\exists x)$ or not.

The following example could explain a bit what's going on here.

Example 1.2.3. In Calculus, you came across expressions of the form:

$$\int_a^b f(x)dx.$$

Here are two integrals, which SYNTACTICALLY are different objects:

$$\int_a^b 3x^2 + 2x + \log(x)dx, \text{ and } \int_a^b 3y^2 + 2y + \log(y)dy.$$

Of course, any sane person understands that the objects these two integrals represent is the same – we are allowed to switch up the the variable x from the dx as long as we switch up in the same way all occurrences of x in our expression. A more murky example is the following:

$$f(y) = y \int y^2 \left(\int y^3 dy + y \int \log(y) dy \right) dy.$$

⁵I am here overloading the notation Var we introduced earlier, but it should always be clear what I'm referring to. In fact, I overloaded that notation before as well, but I'm running out of meaningful symbols.

This is something a maniac would write. A more reasonable human being would write the above as follows:

$$f(x) = x \int y^2 \left(\int z^3 dz + y \int \log(u) du \right) dy.$$

I'm pretty sure you all can see what's happening here. Of course, this is only an aesthetic choice, but sometimes it's good to have good aesthetics.

In general, $(\forall x)$ and $(\exists x)$ in formulas act like dx in integrals, and we'd like to be able to make the distinctions that in integrals come to us rather naturally.

Definition 1.2.4. Let \mathcal{L} be a first-order formula and ϕ an \mathcal{L} -formula. We define three things at the same time:

- (1) The *variables* of ϕ , denoted $\text{Var}(\phi)$,⁶
- (2) The *free variables* of ϕ , denoted $\text{Free}(\phi)$;
- (3) The *bound variable* of ϕ , denoted $\text{Bound}(\phi)$;

all, by induction on the structure of ϕ , as follows:

- (1) If ϕ is atomic, then $\text{Free}(\phi) = \text{Var}(\phi)$,⁷ $\text{Bound}(\phi) = \emptyset$.
- (2) Suppose we have defined our three sets for the \mathcal{L} -formulas ϕ_1 and ϕ_2 .
 - $\text{Var}(\phi_1 \wedge \phi_2) = \text{Var}(\phi_1) \cup \text{Var}(\phi_2)$,
 $\text{Free}(\phi_1 \wedge \phi_2) = \text{Free}(\phi_1) \cup \text{Free}(\phi_2)$,
 $\text{Bound}(\phi_1 \wedge \phi_2) = \text{Bound}(\phi_1) \cup \text{Bound}(\phi_2)$;
 - $\text{Var}(\phi_1 \vee \phi_2) = \text{Var}(\phi_1) \cup \text{Var}(\phi_2)$,
 $\text{Free}(\phi_1 \vee \phi_2) = \text{Free}(\phi_1) \cup \text{Free}(\phi_2)$,
 $\text{Bound}(\phi_1 \vee \phi_2) = \text{Bound}(\phi_1) \cup \text{Bound}(\phi_2)$;
 - $\text{Var}(\phi_1 \rightarrow \phi_2) = \text{Var}(\phi_1) \cup \text{Var}(\phi_2)$,
 $\text{Free}(\phi_1 \rightarrow \phi_2) = \text{Free}(\phi_1) \cup \text{Free}(\phi_2)$,
 $\text{Bound}(\phi_1 \rightarrow \phi_2) = \text{Bound}(\phi_1) \cup \text{Bound}(\phi_2)$;
 - $\text{Var}(\neg \phi_1) = \text{Var}(\phi_1)$,
 $\text{Free}(\neg \phi_1) = \text{Free}(\phi_1)$,
 $\text{Bound}(\neg \phi_1) = \text{Bound}(\phi_1)$;

⁶You got it, I'm overloading, again.

⁷That is, in an atomic formula, ALL variables are free.

- And now we come to the heart of the cheese, $\text{Var}((\forall x)\phi) = \text{Var}(\phi)$, but

$$\text{Free}((\forall x)\phi) = \text{Free}(\phi) \setminus \{x\}.$$

and

$$\text{Bound}((\forall x)\phi) = \begin{cases} \text{Bound}(\phi) \cup \{x\}, & \text{if } x \in \text{Free}(\phi) \\ \text{Bound}(\phi), & \text{if } x \notin \text{Free}(\phi) \end{cases}$$

- $\text{Var}((\exists x)\phi)$, $\text{Free}((\exists x)\phi)$, and $\text{Bound}((\exists x)\phi)$ are defined analogously.

We call an \mathcal{L} -formula ϕ an \mathcal{L} -sentence if $\text{Free}(\phi) = \emptyset$.

Okay that was a total mess – I’m sorry. There are some somewhat subtle (and mostly silly) points that we have to worry about for a minute! Soon we will realise that we are sensible people and stop worrying about them, but such is life. Here’s the first point:

A variable in a formula can be both bound and free.⁸

The second, less important point is the following:

Not every variable that appears next to a quantifier is bound.

The third and last annoyance we could have is the following:

Quantifiers could bind the same variables.

That last annoyance may not mean much to you right now, but as sensible people, we’d really hope to avoid this. Let’s try to illustrate this with some examples:

Example 1.2.5. I’ll not worry too much about \mathcal{L} here, that’s not the point. Suppose that \mathcal{L} has a single relation symbol \underline{R} of arity 2.

- (1) The formula $\underline{R}(x, y)$ has two variables, x and y and they are both free.
- (2) The formula $\underline{R}(x, y) \wedge (\exists x)\underline{R}(x, y)$ has two variables x and y , two free variables x and y and one bound variable x .
- (3) The formula $\underline{R}(x, y) \wedge (\forall z)\underline{R}(x, y)$ has two variables x and y , two free variables x and y and no bound variables.
- (4) The formula $(\forall x)(\forall y)(x \doteq y)$ is a sentence.
- (5) The formula $\underline{R}(x, x) \vee (\forall x)((\forall x)\underline{R}(x, x) \rightarrow \underline{R}(x, x))$ is a mess.

⁸That is to say it’s not always the case that $\text{Bound}(\phi) \cap \text{Free}(\phi) \neq \emptyset$.

This is all very annoying. We can't really get rid of the second point, so we'll define it away:

Definition 1.2.6. Let \mathcal{L} be a first-order language and ϕ an \mathcal{L} -formula. If ϕ is of the form $(\forall x)\psi$ (respectively $(\exists x)\psi$) and $x \notin \text{Bound}(\phi)$ (i.e. $x \notin \text{Free}(\psi)$), then we say that the quantifier $(\forall x)$ (resp. $(\exists x)$) is *vacuous* in ϕ .

In the next subsubsubsection, we'll do a lot of syntactical yoga to get rid of the first and third annoyances. The issue is that formalising common sense can sometimes be hard. That being said, the contents of this subsubsection are not really required for the remainder of this chapter, but it's there as a warm-up for harder substitution procedures to come. Read at your peril.

1.2.1. *Doing the obvious can be hard.* The first part of the next definition should ring some bells:

Definition 1.2.7. Let \mathcal{L} be a first-order language, t an \mathcal{L} -term and $x, y \in \text{Var}$. We define the term $t[y/x]$ inductively, as follows:

- (1) If t is the variable x then $t[y/x]$ is the variable y .
- (2) If t is a constant symbol, then $t[y/x]$ is just t .
- (3) If t is of the form $\underline{f}(t_1, \dots, t_n)$ for some n -ary function symbol $\underline{f} \in \text{Fun}(\mathcal{L})$ and \mathcal{L} -terms t_1, \dots, t_n , then $t[y/x]$ is the term $\underline{f}(t_1[y/x], \dots, t_n[y/x])$.

That's it for terms, lets go to atomic formulas:

Definition 1.2.8. Let \mathcal{L} be a first-order language, ϕ an atomic \mathcal{L} -formula and $x, y \in \text{Var}$. Then, the atomic formula $\phi[y/x]$ is defined inductively as follows

- (1) If ϕ is of the form $t_1 \doteq t_2$ for \mathcal{L} -terms t_1, t_2 then $\phi[y/x]$ is the formula $t_1[y/x] \doteq t_2[y/x]$.
- (2) If ϕ is of the form $\underline{R}(t_1, \dots, t_n)$ for some n -ary relation symbol $\underline{R} \in \text{Rel}(\mathcal{L})$ and \mathcal{L} -terms t_1, \dots, t_n , then $\phi[y/x]$ is the formula $\underline{R}(t_1[y/x], \dots, t_n[y/x])$.

We now have the tools to define a notion of “*free-variable*” substitution. The definition is somewhat complicated but let's try to bare through it together:

Definition 1.2.9. Let \mathcal{L} be a first-order language, ϕ an \mathcal{L} -formula and $x, y \in \text{Var}$. We define $\phi[y/x]_{\text{free}}$ inductively, as follows:

(1) If ϕ is atomic, then $\phi[y/x]_{free}$ is just the atomic formula $\phi[y/x]$ we defined above.

(2) Suppose we have defined $\phi_1[y/x]_{free}$ and $\phi_2[y/x]_{free}$. Then:

- $(\phi_1 \wedge \phi_2)[y/x]_{free}$ is the formula $(\phi_1[y/x]_{free} \wedge \phi_2[y/x]_{free})$.
- $(\phi_1 \vee \phi_2)[y/x]_{free}$ is the formula $(\phi_1[y/x]_{free} \vee \phi_2[y/x]_{free})$.
- $(\phi_1 \rightarrow \phi_2)[y/x]$ is the formula $(\phi_1[y/x] \rightarrow \phi_2[y/x])$.
- $(\neg\phi_1)[y/x]_{free}$ is the formula $(\neg\phi_1[y/x]_{free})$.
- $(\forall z)\phi[y/x]_{free}$ is defined as follows:

$$(\forall z)\phi[y/x]_{free} := \begin{cases} (\forall x)\phi & \text{if } z = x \\ (\forall z)(\phi[y/x]_{free}) & \text{otherwise.} \end{cases}$$

- $(\exists z)\phi[y/x]_{free}$ is similarly defined as follows:

$$(\exists z)\phi[y/x]_{free} := \begin{cases} (\exists x)\phi & \text{if } z = x \\ (\exists z)(\phi[y/x]_{free}) & \text{otherwise.} \end{cases}$$

Intuitively, to compute $\phi[y/x]_{free}$ from ϕ we go down the construction of ϕ and whenever we are not inside some quantifier of the form $(\forall x)$ or $(\exists x)$ (for the SAME x), we change every occurrence of x with y . This is all annoying, but we have to do it if we're serious about things. Let's look at some examples:

Example 1.2.10. Continuing with our binary language from the previous example, say ϕ is the formula:

$$(\forall z)((\forall x)R(x, z) \vee R(x, x)).$$

Then, $(\forall z)((\forall x)R(x, z) \vee R(x, x))[y/x]_{free}$ is computed as follows:

$$\text{Step 1. } (\forall z)((\forall x)R(x, z) \vee R(x, x))[y/x]_{free} = (\forall z)((\forall x)R(x, z) \vee R(x, x))[y/x]_{free}.$$

$$\text{Step 2. } ((\forall x)R(x, z) \vee R(x, x))[y/x]_{free} = ((\forall x)R(x, z)[y/x]_{free} \vee R(x, x)[y/x]_{free})$$

$$\text{Step 3. } (\forall x)R(x, z)[y/x]_{free} = (\forall x)R(x, z) \text{ and } R(x, x)[y/x]_{free} = R(y, y).$$

Step 4. Putting everything together we have:

$$(\forall z)((\forall x)R(x, z) \vee R(x, x))[y/x]_{free} = (\forall z)((\forall x)R(x, z) \vee R(y, y)),$$

So, all the free occurrences of x were replaced by y . On the other hand,

On the other hand, consider the formula:

$$(\forall z)(\forall x)R(x, z)$$

Carrying out the analogous steps, we see that nothing will happen if we apply the free-occurrence substitution we defined above, since there are no free occurrences of x in this formula.

Just like in propositional logic, we can inductively define

$$\phi[y_1/x_1, \dots, y_n/x_n]_{free}.$$

Exercise 1.2.11. To make sure you understand what happened above, write out the proper inductive definition of $\phi[y_1/x_1, \dots, y_n/x_n]_{free}$.

Exercise 1.2.12. Let \mathcal{L} be a first-order language and ϕ an \mathcal{L} -formula of height h . Prove that for any variables $x_1, \dots, x_n, y_1, \dots, y_n \in \text{Var}$, the formula $\phi[y_1/x_1, \dots, y_n/x_n]_{free}$ is an \mathcal{L} -formula of height n , and has the same structure as ϕ .

All of this for the following definition:

Definition 1.2.13. We say that an \mathcal{L} -formula ϕ is *clean* if:

- (1) $\text{Bound}(\phi) \cap \text{Free}(\phi) = \emptyset$.
- (2) For any $x \in \text{Var}$, if ϕ has a subformula of the form $(\forall x)\psi$ or $(\exists x)\psi$, then $x \notin \text{Bound}(\psi)$.

Thus, a formula ϕ is clean if and only if all of its free and bound variables are distinct, and if $(\forall x)\psi$ is a **subformula**⁹ of ϕ , then x is not a bound variable of a subformula of ψ .

The method to clean up formulas may seem rather convoluted, but I promise that once you get the hang of it, it will become second nature:

Definition 1.2.14. Given an \mathcal{L} -formula ϕ , we define **pre-clean**(ϕ), by induction, as follows:

- (1) If ϕ is atomic, then **pre-clean**(ϕ) = ϕ .

⁹The subformulas of a first-order formula are defined analogously to the subformulas of a propositional formula.

(2) Suppose that ϕ is a formula of height $n+1$ and we have defined $\text{pre-clean}(\psi)$ for all formulas of height n . We define $\text{pre-clean}(\phi)$ as follows:

- If ϕ is of the form $(\phi_1 \wedge \phi_2)$, then $\text{pre-clean}(\phi)$ is $(\text{pre-clean}(\phi_1) \wedge \text{pre-clean}(\phi_2))$.
- When ϕ is of the form $(\phi_1 \vee \phi_2)$, $(\phi_1 \rightarrow \phi_2)$ or $(\neg \phi_1)$, we define $\text{pre-clean}(\phi)$ similarly.
- If ϕ is of the form $(\forall x)\phi_1$, then $\text{pre-clean}(\phi)$ is defined to be $(\forall y)\text{pre-clean}(\phi_1[y/x]_{\text{free}})$ for some $y \notin \text{Var}(\phi_1)$.
- If ϕ is of the form $(\exists x)\phi_1$, then $\text{pre-clean}((\exists x)\phi_1)$ is defined to be $(\exists y)\text{pre-clean}(\phi_1[y/x]_{\text{free}})$ for some $y \notin \text{Var}(\phi_1)$.

Once we have defined $\text{pre-clean}(\phi)$ for any \mathcal{L} -formula ϕ , we define $\text{clean}(\phi)$ to be the formula:

$$(\text{pre-clean}(\phi))[y_1/x_1, \dots, y_m/x_m]_{\text{free}},$$

where $\{x_1, \dots, x_m\} = \text{Free}(\text{pre-clean}(\phi))$, and $\{y_1, \dots, y_m\} \in \text{Var} \setminus \text{Var}(\text{pre-clean}(\phi))$.

Let's make our lives really hard for a minute:

Example 1.2.15. Consider the sentence ϕ in a language with a single relation symbol \underline{R} of arity 1 (these fellas are sometimes called, in increasing level's of fanciness: **predicates** or **unary relations** or **monadic relations**):¹⁰

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right)$$

We compute $\text{pre-clean}(\phi)$:

$$\begin{aligned} \text{Step 1. By definition, } \text{pre-clean} \left((\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right) \right) \text{ is:} \\ (\forall y) \text{pre-clean} \left(\left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right) [y/x]_{\text{free}} \right) \end{aligned}$$

Step 2. Computing the free-variable substitution above we get:

$$(\forall y) \text{pre-clean} \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(y) \right)$$

¹⁰This may have some unnecessary brackets as we go down the steps, but I'm putting these just to make the reading a bit cleaner.

Step 3. We repeat Step 1, for the formula $(\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(y)$ to get:

$$(\forall z) \left((\text{pre-clean}((\forall x) R(x)) \rightarrow R(z)) \rightarrow \text{pre-clean}(R(y)) \right)$$

Step 4. After we start having some faith, we see that once we put everything back together we get the following clean formula:

$$(\forall y) \left((\forall z) \left((\forall u) R(u) \rightarrow R(y) \right) \rightarrow R(z) \right)$$

Exercise 1.2.16. Write out in detail the missing cases from the definition of **pre-clean**.

To get some practice with the definitions, you are welcome to do the following annoying exercise in induction:

Exercise 1.2.17. Let ϕ be an \mathcal{L} -formula. Show that:

- (1) ϕ and $\text{pre-clean}(\phi)$ have the same structure.
- (2) $\text{Free}(\phi) = \text{Free}(\text{pre-clean}(\phi))$.
- (3) $\text{Bound}(\phi) \cap \text{Bound}(\text{pre-clean}(\phi)) = \emptyset$.
- (4) $\text{Bound}(\text{pre-clean}(\phi)) = \text{Bound}(\text{clean}(\phi))$.

Now, for the main course:

Lemma 1.2.18. *For any first-order language \mathcal{L} and any \mathcal{L} -formula ϕ , the \mathcal{L} -formula $\text{clean}(\phi)$ is a clean formula.*

PROOF. We have to show two things:

- (1) For all $x \in \text{Var}(\text{clean})(\phi)$, if $x \in \text{Free}(\text{clean}(\phi))$, then $x \notin \text{Bound}(\text{clean}(\phi))$.
Let $x \in \text{Var}(\text{clean}(\phi))$. By definition, we have that:

$$\text{clean}(\phi) = (\text{pre-clean}(\phi))[y_1/x_1, \dots, y_m/x_m]_{\text{free}},$$

where $\{x_1, \dots, x_m\} = \text{Free}(\text{pre-clean}(\phi))$, and $\{y_1, \dots, y_m\} \in \text{Var} \setminus \text{Var}(\text{pre-clean}(\phi))$. Thus, if $x \in \text{Free}(\text{clean}(\phi))$, we have that $x = y_i$, for some $i \leq m$. But then, $x \notin \text{Var}(\text{pre-clean}(\phi))$ and, in particular, $x \notin \text{Bound}(\text{pre-clean}(\phi)) = \text{Bound}(\text{clean}(\phi))$, by the previous exercise.

(2) Now, to show that if $\text{clean}(\phi)$ has a subformula of the form $(\forall x)\psi$ (or $(\exists x)\psi$), then $x \notin \text{Bound}(\psi)$. It suffices to show this when $\text{clean}(\phi)$ is itself of the form $(\forall x)\psi$ (see next exercise). By definition, we know that if $\text{clean}(\phi)$ is of this form, then $\text{clean}(\phi)$ is the formula:

$$(\forall x)(\text{pre-clean}(\chi)[y_1/x_1, \dots, y_m/x_m]_{\text{free}}),$$

for some formula χ , where $x \notin \{x_1, \dots, x_n\}$. By the previous exercise:

$$\begin{aligned} \text{Bound}(\psi) &= \text{Bound}(\text{clean}(\chi)) \\ &= \text{Bound}(\text{pre-clean}(\chi)[y_1/x_1, \dots, y_m/x_m]_{\text{free}})) \\ &\subseteq \text{Var}(\chi) \end{aligned}$$

so, if $x \in \text{Bound}(\psi)$, then $x \in \text{Var}(\chi)$, which is a contradiction.

□

Exercise 1.2.19. Write down the details of the proof of the lemma above.

From now on, we will assume ALL FORMULAS we are dealing with are clean, by implicitly always replacing a given formula ϕ with $\text{clean}(\phi)$. When we discuss semantics, we will see that this is purely an aesthetic assumption.

End of digression.

1.2.2. *Another approach to cleanliness.* All of this was very annoying, and in fact, just a formal way of doing something very trivial. Given an \mathcal{L} -formula ϕ of the form $(\forall x)\psi$, we say that the **scope** of the quantifier $(\forall x)$ is the subformula ψ . Any *free* occurrence of the variable x in ψ is bounded by the quantifier $(\forall x)$. If somewhere in ψ there is a subformula of the form $(\forall x)\chi$ (or $(\exists x)\chi$), then it is *this* quantifier that binds all the occurrences of x in χ . For example, in the formula that we were playing with before:

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right)$$

the scope of the outermost quantifier is:

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right)$$

so all FREE occurrences of x in the blue part of the formula are bound by this outer quantifier. Similarly, we can see that the scope of the second quantifier is:

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right)$$

so again, the FREE occurrences of x in the red part are bound by this intermediate quantifier. Finally, the scope of the innermost quantifier is:

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right).$$

All in all, the occurrences of x bound by each quantifier are as follows:

$$(\forall x) \left((\forall x) \left((\forall x) R(x) \rightarrow R(x) \right) \rightarrow R(x) \right).$$

(Pre-)Cleaning up the formula simply involved giving different names to the variables in each colour. The final clean up step involves renaming all free variables so that they don't clash with the quantified variables. Now that we no longer have to worry about our formulas being dirty, let's set up some (nice, for a change) notation.

Notation 1.2.20. Let ϕ be an \mathcal{L} -formula. We will write $\phi(x_1, \dots, x_n)$ to indicate the following two things:

- (1) $\text{Free}(\phi) \subseteq \{x_1, \dots, x_n\}$.
- (2) $x_i \neq x_j$ whenever $1 \leq i \neq j \leq n$.