Introductory Notes for Halbach's Logic Manual

Peter Millican, Hertford College, Oxford

The aim of these notes is to help students master some of the basic logical concepts covered in Volker Halbach's *Logic Manual*. Here things are covered slightly differently from in the *Logic Manual* itself, giving a complementary perspective to help solidify students' understanding, with different examples and emphasis. These notes are also designed to be read in advance, without tutorial assistance, and they concentrate on particular areas that might cause confusion later if they are not straightened out from the start. However the *theory* is intended to follow what is laid down in the *Logic Manual (as at September 2009)*, and all definitions reproduce exactly the Halbach text. I am very grateful to Volker Halbach for his willingness to check my text for accuracy, so that users of these notes can be confident that they are following the required system precisely.

My section numbers broadly follow those of the *Logic Manual*, except that I generally divide the sections further into sub-sections. So for example §1.2 below, entitled "Relations", corresponds to Halbach's section §1.2 "Binary Relations". In my presentation, "Binary Relations" are discussed from §1.2.3 onwards, because §1.2.1 and §1.2.2 are devoted to some general points about intensionality and extensionality which lie behind some of the logical treatment of relations. These points can cause confusion, both for novices but also for those who are used to different formal treatments; hence I have given additional space to them.

1 Sets, Relations, and Arguments

1.1 What is a Set?

Put crudely, a set is simply a collection of items, each of which is called a *member* or an *element* of that set. Sets are *extensional*, which means that sets containing the same items (i.e. having the same elements) are one and the same set. So *the set of colours in the French flag*, which contains the three colours red, white and blue, is the same as *the set of colours in the British flag* – these are two different ways of referring to the very same set.

1.1.2 Notation for Sets

A set may be specified *extensionally*, by listing all of its elements, or *intensionally*, by giving a rule for determining its elements. All of the following are alternative ways of specifying the very same set – (a) to (d) are extensional specifications, and (e) to (f) intensional:

(a) $\{1, 2, 3\}$
(b) $\{3, 1, 2\}$
(c) $\{+\sqrt{4}, 1, +\sqrt{9}\}$
(d) $\{2, 3, 2, 3, 1\}$
(e) $\{x : x^3 - 6x^2 + 11x - 6 = 0\}$
(f) $\{x : x$ is an integer **and** $1 \leq x \leq 3\}$

Note that:

- We use curly brackets to indicate a set.
- The order in which the elements are listed makes no difference to the set's identity.
- Repetitions in a specification are irrelevant – any element "counts" only once no matter how many times it is specified.
- A colon is used to indicate an intensional specification of a set. Read "$\{x : ...\}$" as "those items which yield a truth when substituted for x in the following: ...".
- The polynomial $x^3 - 6x^2 + 11x - 6$ can be factorised as $(x-1)(x-2)(x-3)$, so this comes to zero if and only if x is equal to 1, 2, or 3.

We conventionally use capital letters to denote sets and lower case letters to denote their elements (or members). The sign "$\in$" means "is an element of", hence:

$m \in A$	states that	m is an element of set A, e.g. $7 \in \{2,3,5,7,11,13\}$

We won't be needing other set-theoretical terminology, but it's worth acquainting yourself with some other basic notions of set theory, which are used a fair bit in technical areas of philosophy, especially the following. The *cardinality* of a set (i.e. the number of elements it contains):

$\|A\|$	refers to	the cardinality of set A, e.g. $\|\{2,3,5,7,11,13\}\| = 6$;

the *intersection* of two sets (i.e. the set of elements that are common to both of them):

$A \cap B$	refers to	the intersection of sets A and B, e.g. $\{1,3,5\} \cap \{2,3,5,7\} = \{3,5\}$

the *union* of two sets (i.e. the set of items that are in one or both of the sets):

$A \cup B$	refers to	the union of sets A and B, e.g. $\{1,3,5\} \cup \{2,3,5,7\} = \{1,2,3,5,7\}$

Set A is called a *subset* of set B if there is no element of A which is not also an element of B:

$A \subseteq B$	states that	A is a subset of B, e.g. $\{3,5\} \subseteq \{2,3,5,7\}$, and $\{2,3,5,7\} \subseteq \{2,3,5,7\}$

Notice this implies that every set is a subset of itself; A is called a *proper subset* of B if it is a subset of B but not identical to B:

$A \subset B$	states that	A is a *proper* subset of B, e.g. $\{3,5\} \subset \{2,3,5,7\}$, but $\{2,3,5,7\} \not\subset \{2,3,5,7\}$

1.1.3 The Null Set

The set of aardvarks enrolled at Oxford University, and *the set of numbers greater than 3 and less than 1* both have no members whatever. Since their membership is the same, they are therefore *one and the same* set. This unique set, the set which has no elements, is called *the null set*. The symbol for the null set is: $\varnothing$

Note that the null set is a subset of *every* set, because whichever set A you choose, there are no elements in the null set that are not also in A. Note also that any set containing n elements will have 2^n subsets, e.g. $\{1,2,3\}$ has the 8 subsets $\varnothing$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1,2\}$, $\{1,3\}$, $\{2,3\}$, $\{1,2,3\}$.

1.2 Relations

Typically a sentence of English will say something about one or more objects, for example "That house is cold", "The clock made a noise", "John kissed Jane" etc. Here we can distinguish *referring expressions* such as "That house", "The clock", "John", and "Jane", which serve to identify the objects concerned, from *predicate expressions* such as "is cold", "made a noise" and "kissed", which say something about those objects. Predicates can be 1-place, 2-place, or multi-place; in logic we shall be using three dots "…" to indicate the places.[1] Reasoning involving numbers, for example, might make use of the following predicates:

… is prime	1-place or "unary"
… is greater than …	2-place or "binary"
… is a common factor of … and …	3-place or "ternary"

1-place predicates typically correspond to *properties* of individual things (e.g. "… is prime", "… is cold"), whereas predicates with more "places" correspond to *relations* between things. 2-place predicates correspond to *binary relations* (e.g. "… kissed …", "… is greater than …"), which are of particular interest and dealt with at some length in Chapter 1 of Halbach's *Logic Manual*. Note that a property can be thought of as a 1-place relation (because it's represented by a 1-place predicate – don't worry about the oddity in calling it a "relation" when there's only one object involved).

1.2.1 Properties and Extensionality

Given any property, the *extension* of the property will be the set of objects to which it applies. For example the extension of the property of *being an odd number between 3 and 8 inclusive* is the set {3,5,7}. But note that the property of *being a prime number between 3 and 8 inclusive* also has the same extension {3,5,7}. We think of these as different properties, but they have the same extension. Likewise, the property of *being a man weighing over a tonne* has the extension ∅ (the null set), and so does the property of *being a man weighing less than a kilogram*. It seems, then, that having the same extension is not enough to guarantee sameness of property. Properties, as we usually understand them, are therefore *intensional* rather than *extensional*.

However properties can also be interpreted *extensionally* – i.e. as defined by their extension – which though rare in Philosophy is relatively common in Computing, for example. On this understanding of the notion, the property of *being an odd number between 3 and 8 inclusive* is *exactly the same property* as the property of *being a prime number between 3 and 8 inclusive*. Likewise – and perhaps seemingly paradoxically – on this understanding the property of *being a man weighing over a tonne* is exactly the same property as the property of *being a man weighing less than a kilogram*.

To repeat, however, as far as we are concerned here, *properties are intensional*, so that it is possible to have two different properties with the same extension.

1.2.2 Relations are Extensional

Now consider the binary relation expressed by the predicate "… is twice …" restricted to the set {3,4,5,6,7,8}. There are only two cases for which this is true: "6 is twice 3" and "8 is twice 4". Hence the *extension* of this relation consists of two *pairs* of numbers, <6,3> and <8,4>. These are called *ordered pairs* because it matters in which order they are listed: 6 is twice 3, but it's not true that 3 is twice 6. So the extension of a binary relation is a *set of ordered pairs*, in this case:

[1] Technically, these are called *argument* places, with "argument" referring to whatever gets substituted in the gaps. This seems a strange use of the word "argument" (having nothing to do with the more common use), but it is as well to be aware of it. If you see a reference to a predicate "with two argument places", this simply means a binary predicate.

{ <6,3> , <8,4> }

In the same way, the extension of a ternary relation will be a set of ordered triples, and so on. An ordered pair can also be called a "2-tuple"; an ordered triple is a "3-tuple", and the general term for these angle-bracketed sequences of objects is "n-tuples".

As pointed out above, philosophers normally think of properties as intensional (and we are following this convention), but it's far more common to treat *relations* as *extensional*, especially in formal contexts. Suppose, for example, that Chris and Christine love each other, as do Paul and Pauline, but that the two couples are completely unaware of each other's existence – perhaps they even live in different countries. Consider the relations expressed by the predicates "… knows …" and "… loves …", restricted to the set consisting of just those four people:

… knows …	{<Chris, Christine>, <Christine, Chris>, <Paul, Pauline>, <Pauline, Paul>}
… loves …	{<Chris, Christine>, <Christine, Chris>, <Paul, Pauline>, <Pauline, Paul>}

On the extensional understanding of relations, these two relations are *one and the same*: as far as this set is concerned, knowing someone is loving them! This no doubt seems counter-intuitive, but don't let it worry you at all. Just think of it as a formal convenience, which enables us to do our work without worrying about the metaphysical nature of relations: all we're interested in when we handle them in logic is which pairs of objects (or triples etc.) are related in the appropriate way.

From now on, we shall treat binary relations extensionally, and therefore as equivalent to sets of ordered pairs (and, correspondingly, ternary relations as sets of ordered triples, and so on). Consistently with this, we shall treat unary relations extensionally also – so they will *not* be equivalent to properties (which we are interpreting intensionally). Hence the unary relations expressed by the predicates "… is odd" and "… is prime", restricted to the set {3,4,5,6,7,8}, will be understood as follows:[2]

… is odd:	{<3>, <5>, <7>}	*or*	{3, 5, 7}
… is prime:	{<3>, <5>, <7>}	*or*	{3, 5, 7}

When only one object is involved, the angle brackets bring no benefit, so for convenience we shall treat a 1-tuple as equivalent to the object concerned, which makes a unary relation equivalent to a simple set of objects – in this case {3, 5, 7}.

1.2.3 Properties of Binary Relations

The discussion above explains what we mean by a "binary relation":

DEFINITION 1.1 *A set is a binary relation if and only if it contains only ordered pairs*

So a binary relation can contain *nothing that isn't* an ordered pair. Since the null set ∅ contains nothing at all, it satisfies this definition. Hence we can also call the null set *the null relation*.

Here are some examples of non-null binary relations. In each case a description is given to explain how the choice of pairs has been made. But notice that as we are understanding them, a relation is a set of ordered pairs, and the description is quite inessential to it: any other description that captured the same set would do just as well.

---

[2] I have talked of predicates as being "restricted to" a particular set, to make it straightforward to express simple examples involving just a few objects. But this doesn't imply that relations themselves are to be understood as relative to a set. As we shall understand relations, they are simply sets of n-tuples. Hence the very same unary relation as discussed here – namely {3, 5, 7} – can equally well be considered within the domain of all integers, though in that case it could not of course be expressed using only the predicates "… is odd" or "… is prime"; it would be necessary to list the elements, or use a predicate like "… is an odd number between 3 and 8 inclusive".

	Description	Relation (i.e. Set of Ordered Pairs)
(a)	… < … on the set {1,2,3}	{<1,2>, <1,3>, <2,3>}
(b)	… ≤ … on the set {1,2,3,4}	{**<1,1>**, <1,2>, <1,3>, <1,4>, **<2,2>**, <2,3>, <2,4>, **<3,3>**, <3,4>, **<4,4>**}
(c)	0 < |…–…| < 3 on the set {1,2,3,5,7}	{<1,2>, <1,3>, <2,1>, <2,3>, <3,1>, <3,2>, <3,5>, <5,3>, <5,7>, <7,5>}

"|α–β|" *means the absolute value of* α *minus* β*, so for example if* β *is greater than* α*, then* |α-β| *will be equal to* β-α *rather than* α-β*. Hence if* |α-β| *is greater than zero but less than 3, this means that* α *and* β *must differ by just 1 or 2.*

(d)	… mod 3 = … mod 3 on the set {1,4,5,6,9,10}	{**<1,1>**, <1,4>, <1,10>, <4,1>, **<4,4>**, <4,10>, **<5,5>**, **<6,6>**, <6,9>, <9,6>, **<9,9>**, <10,1>, <10,4>, **<10,10>**}

"α *mod 3" means the remainder when* α *is divided by 3, so for example 5 mod 3 is equal to 2. Hence* α *mod 3 =* β *mod 3 if, and only if,* α *and* β *differ by a multiple of 3 (i.e. 0, or 3, or 6, or 9, or …).*

(e)	… + … = 5 on the set {1,2,3,4}	{<1,4>, <2,3>, <3,2>, <4,1>}
(f)	…^{2} + … < 6 on the set {1,2,3}	{**<1,1>**, <1,2>, <1,3>, <2,1>}

Some of the pairs within these relations are of the form <α,α> (sometimes called a "reflexive pair"); these are shown above in bold type. Note accordingly that relation (b) is said to be *reflexive on the set {1,2,3,4}* because for any α in that set, <α,α> is a member of the relation. Likewise relation (d) is reflexive on the set {1,4,5,6,9,10} because for any α in that set, <α,α> is a member of the relation. Notice, however, that (b) is not reflexive on the set {1,2,3,4,5}, because the pair <5,5> is not an element of the relation. So as we are understanding relations here – as simple sets of ordered pairs – a relation will never be said to be *reflexive* absolutely, but only *reflexive on a set*, and a different set might give a different result.[3]

By contrast (a), (c) and (e) are *irreflexive*, since there are no pairs of the form <α,α> within the relation. Notice that (unlike *reflexivity*) this property of the relation need not be defined relative to any other set – it is simply a matter of whether the relation contains any "reflexive pairs" or not. (Note that *irreflexive* is not defined in the *Logic Manual*, and not required for the course.)

Relations (c), (d) and (e) are *symmetric*, since for any α and β, if <α,β> is a member of the relation then <β,α> is too. By contrast, (a) is *asymmetric*, because if <α,β> is a member of the relation then <β,α> is not. Relations (b) and (f) are neither symmetric nor asymmetric (hence they can be called "non-symmetric"). But (b) fails to be asymmetric only in virtue of its containing the "reflexive pairs" <1,1>, <2,2>, <3,3> and <4,4>, so it is therefore *antisymmetric*: this means that although it may be possible for both <α,β> and <β,α> to be members of the relation, this can happen only if α

[3] Note that any relation at all is *reflexive on the null set*. This exemplifies a general point about "any" and "all" as understood in logic. The null set contains no elements at all, hence trivially it follows that all of its elements are even numbers, all of its elements are pink elephants, etc. (It might be helpful to think "All none of them are even.", just as you might say "All 6 of them" if there were 6.) It can be easier to understand this negatively: to say that all elements of a set are pink elephants is to say that *there is no element of the set that is not a pink elephant*. Likewise, the statement that "For any α in set S, <α,α> is a member of relation R" should be understood as meaning that "There is no α in set S for which <α,α> is *not* a member of relation R". Then it is clear that this will automatically be true if S is the null set.

and β are identical. (Think of an antisymmetric relation as being *either asymmetric or very nearly so*: its asymmetry can "break down" – at most – only in the special case when the relation applies between an object and itself.)

Relations (a), (b) and (d) are *transitive*, since for any α, β and γ, if both <α,β> and <β,γ> are members of the relation, then <α,γ> is too. By contrast, Relation (e) is *intransitive*, since for any α, β and γ, if both <α,β> and <β,γ> are members of the relation, then <α,γ> is **not** a member of it. (c) and (f) are neither transitive nor intransitive; hence they can be said to be *non-transitive*. (Note that *intransitive* and *non-transitive* are not defined in the *Logic Manual*, and not required for the course.)

1.2.4 Relativised and Unrelativised Properties

We have now seen that *symmetry*, *asymmetry*, *antisymmetry*, and *transitivity* (unlike *reflexivity*) can all be understood in a form that is not relative to any particular set. Here is what the corresponding unrelativised definition would be like, as given by Halbach (in the text following Definition 1.3 – note in what follows that "iff" is short for "if and only if"):

UNRELATIVISED DEFINITION *A binary relation R is*

symmetric iff for all d, e: if <d,e> ∈R then <e,d> ∈R;

asymmetric iff for no d, e: <d,e> ∈R and <e,d> ∈R;

antisymmetric iff for no two distinct d, e: <d,e> ∈R and <e,d> ∈R;

transitive iff for all d, e, f: if <d,e> ∈R and <e,f> ∈R, then also <d,f> ∈R.

However some logicians prefer to define these concepts (like *reflexivity*) as relativised to a set. So for example relation (f) above – {<1,1>, <1,2>, <1,3>, <2,1>} – although not in itself *symmetric* because it contains the pair <1,3> without containing also <3,1>, would count as *symmetric on the set {1, 2}*, because it contains no pair *from that set* without containing also its symmetric partner.

Halbach's initial definitions of all of these terms specify them as relativised to set *S*:

DEFINITION 1.2 *A binary relation R is*

(i) *reflexive on a set S iff for all elements d of S the pair <d,d> is an element of R;*

(ii) *symmetric on a set S iff for all elements d, e of S: if <d,e> ∈R then <e,d> ∈R;*

(iii) *asymmetric on a set S iff for no elements d, e of S: <d,e> ∈R and <e,d> ∈R;*

(iv) *antisymmetric on a set S iff for no two distinct (that is, different) elements d, e of S: <d,e> ∈R and <e,d> ∈R;*

(v) *transitive on a set S iff for all elements d, e, f of S: if <d,e> ∈R and <e,f> ∈R, then also <d,f> ∈R.*

Then the unrelativised notions are defined in terms of the relativised notions, as follows:

DEFINITION 1.3 *A binary relation R is*

(i) *symmetric iff it is symmetric on all sets;*

(ii) *asymmetric iff it is asymmetric on all sets;*

(iii) *antisymmetric iff it is antisymmetric on all sets;*

(iv) *transitive iff it is transitive on all sets.*

In the *Logic Manual*, this definition is immediately followed by the "Unrelativised Definition" above (though this is not given a definition number, because it counts as merely spelling out the consequences of Definitions 1.2 and 1.3, rather than as a formal definition itself).

1.2.5 Equivalence Relations

If a relation is *reflexive on set S*, *symmetric on S* and *transitive on S*, then it is called an *equivalence relation on set S*. Note that an equivalence relation is reflexive; hence it is not possible to have an unrelativised notion of an equivalence relation.

<u>DEFINITION 1.4</u> *A binary relation R is an equivalence relation on S iff R is reflexive on S, symmetric on S and transitive on S.*

Equivalence relations are given a special name because they have the interesting property of dividing or "partitioning" the set on which they are defined into distinct subsets. Take the case of (d) above, for example:

{<1,1>, <1,4>, <1,10>, <4,1>, <4,4>, <4,10>, <5,5>, <6,6>, <6,9>, <9,6>, <9,9>, <10,1>, <10,4>, <10,10>}

This relation partitions the set {1,4,5,6,9,10} into three subsets: {1,4,10}, {5}, and {6,9}. If you pick any element α from one of these subsets, and element β from the same subset, then <α, β> will be in (d); but if you pick β from a different subset, then <α, β> will not be in (d).

1.2.6 Arrow Diagrams and Properties of Binary Relations

Binary relations can often be represented using an *arrow diagram*, in which:

- a set containing all the relevant *objects* – that is, the objects related by the relation – is represented by a circle or similar shape (as in a Venn diagram);[4]
- the individual objects are represented by labelled points or crosses;
- the elements of the relation (i.e. the ordered pairs of objects) are represented using arrows which join the objects, as follows: *a single arrow* from α to β indicates that the relation holds between α and β (in that order), *a two-way "double arrow"* between α and β indicates that the relation holds both between α and β and between β and α, and *a (doubly-arrowed) "loop"* attached to α indicates that the relation holds between α and itself.

Note that whenever any objects α and β are mutually related, then the arrow between α and β must be a double arrow. Thus, in particular, when any object α is related to itself (think of this as the case when β=α), the loop which indicates this is deemed to be a double arrow (this does make sense, because just like in the case of an ordinary double arrow, the object at each end of the loop is indeed related to the object at the other end).

Using arrow diagrams to represent binary relations can help to make clear the significance of the various properties defined above. Note that with the properties other than reflexivity, the notions can be understood as relativised to set *S* by restricting attention to arrows that both start and end at objects within set *S*.

[4] Note that an arrow diagram does not have to be restricted to the set of those objects that are actually related by the relation – what is often called the *field* of the relation. The diagram can include other objects as well. Clearly any relation that involves an infinite number of pairs cannot be represented by an arrow diagram!

R is reflexive on set S	iff	Within *R*'s arrow diagram, every object in set *S* (or strictly, every point that represents such an object) has a loop attached, i.e. there is no object in set *S* which *lacks* a loop.
R is symmetric	iff	No arrow in *R*'s arrow diagram is single (i.e. one-way).
R is asymmetric	iff	No arrow in *R*'s arrow diagram is double (and hence there must be no loops).
R is antisymmetric	iff	The only double arrows (if any) in *R*'s arrow diagram are loops: there are no double arrows which are *not* loops.
R is transitive	iff	Every "two-step journey" in *R*'s arrow diagram has a direct "short cut". That is, there are no two-step journeys which *lack* a short cut.
R is intransitive	iff	No "two-step journey" in *R*'s arrow diagram has a direct "short cut".[5]

Some of these properties are related in surprising ways. For example a relation can be *both* transitive and intransitive, if it contains no "two-step journies" at all, like the relation which contains all, and only, the <husband, wife> pairs amongst Oxford students. There are no individuals α, β and γ such that α is the husband of β and β is the husband of γ (β cannot be both male and female!). So there are no two-step journeys at all and hence none which either lack, or have, a direct short cut.[6] Note also that the null relation (but *only* the null relation) is both symmetric and asymmetric: it has no single arrows, and no double arrows either. It is indeed worth remembering the strange properties of the null relation, which is *reflexive on the null set* (but not on any other set), *symmetric*, *asymmetric*, *antisymmetric*, *transitive*, and *intransitive*.

1.3 Functions

DEFINITION 1.5 *A binary relation R is a function iff for all d, e, f: if <d,e> ∈R and <d,f> ∈R then e=f.*

This says, in effect, that in the arrow diagram of a function, you can never have two distinct arrows leaving from one point. Functions are of particular interest in mathematics, because they can take us from one number to another, so for example if α is some number we can talk about *the square of* α, or *the cube root of* α, or *the sine of* α, or *the natural logarithm of* α, and so forth. We can do this only because each of these is *well-defined*: a number has one, and only one, cube root or square for example.[7] Notice that in this sense *square root* is not a well-defined function, because a positive number (such as 4) has *two* square roots (in this case, 2 and -2). In practice, mathematicians turn it into a well-defined function by adopting the convention that $\sqrt{\alpha}$ should always be interpreted as referring to the *positive* square root of α, so this has a unique value. You might wonder about negative numbers, which have no square root at all, but this does not prevent $\sqrt{\alpha}$ from being a well-defined function. It just means that there are some numbers for which the function is not defined, or

[5] Again note that the notion of intransitivity, though widely used, is not defined in the *Logic Manual*, and not required to be known for this logic course.

[6] To put this another way, all of the two-step journeys that there are (*all none of them!*) both have, and don't have, a short cut. (Keep in mind some absurd but perhaps convincing example such as the following: "All round squares are round" and "All round squares are square" are both true "by definition" – hence all round squares, *all none of them*, are both round and square.)

[7] For those who know about such things, we're assuming here a restriction to *real* numbers rather than *complex* numbers.

in technical language, that the *domain* of the function – the set of values for which it *is* defined – is the set of *non-negative* numbers rather than the set of *all* numbers.

As another example, take the relation – informally "mother of" – consisting of all ordered pairs $<d,e>$ where *d* is a living human being and *e* is that person's biological mother (whether alive or dead). This is a well-defined function, because nobody has more than one biological mother. And here we can distinguish between the *domain* of the function (i.e. the set of all living people) and the *range* of the function (i.e. the set of all people who are some living person's mother; in other words, the set of women – alive or dead – who have a living child). Notice, also, that "child of" is *not* a well-defined function, because many people have more than one child. However "eldest naturally born child of" would probably be a well-defined function.[8]

DEFINITION 1.6

(i) *The domain of a function R is the set {d: there is an e such that* $<d,e> \in R$*}.*

(ii) *The range of a function R is the set {e: there is a d such that* $<d,e> \in R$*}.*

(iii) *R is a function into the set M if and only if all elements of the range of the function are in M.*

The last clause here means that "mother of" is a function into the set of women, and also a function into the set of people.

Where we do have a well-defined function *R*, it enables us to talk without ambiguity about "*the R* of α". For example we can refer to "*the* mother of α" or "*the* eldest naturally born child of α", whoever α might be, without any risk of ambiguity, but we cannot talk with the same safety about "*the* child of α". This is indeed partly why the notion of a function is so useful in mathematics, as it enables us to talk about *the square of* α, *the cube root of* α etc. The following definition exploits this feature of functions.

DEFINITION 1.7 *If d is in the domain of a function R one writes R(d) for the unique object e such that* $<d,e>$ *is in R.*

1.4 Non-Binary Relations

Binary relations, as we have seen, consist of sets of ordered pairs. In the same way, ternary (3-place) relations consist of ordered triples, and so on. As we saw in §1.2.2 above, the word "tuple" is used to cover these ordered pairs and triples: an ordered pair can also be called a "2-tuple", an ordered triples is a "3-tuple" and so on. The term "n-tuple" is used to refer generally to a tuple of indeterminate length *n*. It's also useful to have a word to refer to the "number" of a relation, and so mathematicians have coined the term "arity" (think of "un*ary*", "bin*ary*", "tern*ary*", and you get "ary-ty"). So a binary relation has arity 2, a ternary relation has arity 3, and so on.

[8] "Probably", because it is possible but very unlikely that some man could have had two children born simultaneously from different mothers – "naturally born" excludes the more likely situation where twins are born simultaneously by Caesarian section, in which case again there would be no unique eldest. An alternative way of dealing with this sort of situation would be to say that someone whose eldest children are born simultaneously has *no* eldest child.

1.5 Arguments, Validity and Contradiction

The logical notion of an *argument* is, of course, quite different from the conversational notion (of a row, or heated dispute, or disagreement). One of the speakers in Monty Python's famous argument sketch was nearer the mark: "An argument is a connected series of statements to establish a definite proposition". So here is an example of an argument in this sense:

(a) John is studying Physics at Oxford University
(b) Every Oxford student either has 'A' levels or is a mature student
(c) So John must either have 'A' levels or be a mature student
(d) But John has a young person's railcard
(e) And nobody with a young person's railcard can be a mature student
(f) So John cannot be a mature student
(g) Hence John must have 'A' levels
(h) Everyone studying Physics at Oxford University who has any 'A' levels at all has at least one Mathematics 'A' level
(i) So if John studies Physics at Oxford University and has 'A' levels, he must have a Mathematics 'A' level
(j) Therefore John must have a Mathematics 'A' level

Within this argument the various "steps" perform different roles:

- (a), (b), (d), (e) and (h) state propositions which are here taken for granted rather than being themselves supported by argument: these are called the "premisses" of the argument.
- (j) is the proposition which the argument is put forward to establish: this is called the argument's "conclusion".
- (c), (f), (g) and (i) are intermediate conclusions, which are (or purport to be) established from the premisses on the way to, and as a means of, establishing the final conclusion.

In the study of logic we usually simplify matters by analysing arguments purely in terms of their premisses and conclusions, ignoring all intermediate steps and all words used to indicate the flow of the argument such as "so", "but then", "must", "hence" and so on. We could therefore represent the argument above in the following simplified form:

1. John is studying Physics at Oxford University
2. Everyone studying at Oxford University either has A-levels or is a mature student
3. John has a young person's railcard
4. Nobody who has a young person's railcard is a mature student
5. Everyone studying Physics at Oxford University who has A-levels has at least one Mathematics A-level

---

∴ John has at least one Mathematics A-level

This suggests a provisional definition of what an argument is, but we shall soon go on to refine it:

<u>PROVISIONAL DEFINITION</u> *An argument consists of a set of premisses, together with a conclusion.*

We need to refine this by getting a bit clearer on what exactly a premiss and conclusion are. Are they *propositions* – and if so, what are *they*? Or are they *sentences*? We shall in fact be opting for a subset of the latter, namely *declarative sentences*, but it might be helpful first to explain why.

1.5.1 Propositions, Declarative Sentences, and Logical Form

In the previous section, the word *proposition* was used for the constituents of an argument. This is a philosopher's term of art, typically defined as being whatever is true or false, asserted, believed, hypothesised, or denied etc. On this understanding, we say that a proposition is *expressed* by a sentence. The following sentences, for example, seem to state that some determinate state of affairs is the case, and as such express propositions – they can be straightforwardly true or false etc.

Grass is not white
1 + 1 = 23
Some pigs can fly

The following sentences, by contrast, do not express propositions:

Are you feeling unwell?
Go away and never come back!
Aardvark window very looking is.

We are interested in only those sentences that are *declarative*: that purport to state some proposition, true or false – so questions, commands, and ungrammatical gibberish are excluded. But even some sentences with a declarative form arguably don't express propositions:

(a) Colourless green ideas sleep furiously
(b) The law reflects the General Will
(c) Abortion is morally wrong
(d) He feels hot just now
(e) The big window faces John's house
(f) Oxford United will win on Saturday

These are problematic for a variety of reasons:

(a) "Colourless green ideas sleep furiously", although grammatically well-formed, is hard to make sense of, and might accordingly be thought meaningless.

(b) "The law reflects the General Will" is problematic because the notion of the "General Will" is dubious and perhaps incoherent.

(c) "Abortion is morally wrong" might be disputed because some would claim that moral words merely express attitudes rather than asserting that some state of affairs obtains.

(d) "He feels hot just now" can be used to express some proposition on a particular occasion, but by itself does not make clear which determinate proposition is intended (which male individual is being referred to, and when is "now"?).

(e) "The big window faces John's house" is another illustration of context-relativity (which indeed probably affects most utterances): here we have a definite description and a name which can be used to speak of different windows and different Johns.

(f) "Oxford United will win on Saturday" is likewise context-relative (which "Saturday" is involved?), and also raises questions about the future: some philosophers claim that a sentence about what *will* happen must express a *judgement* rather than a proposition (involving an act of *prediction* or *conjecture* rather than *assertion*), and so cannot be literally true or false.

In what follows, for the sake of simplicity, we shall be very liberal regarding what counts as a proposition: anything expressed by a grammatical declarative sentence will do (so all of this last group would be allowed). However context-relativity remains potentially a significant problem. Consider, for example, the following argument:

He is 12 years old;
If he is 12 years old, then he is tall for his age;
Therefore he is tall for his age.

This seems to have the form:

O;
If *O* then *T;*
Therefore *T*.

And (to anticipate a notion that will be introduced in §1.5.2 below) this looks like a *valid* argument-form: the conclusion *follows from* the premisses, in the sense that if the premisses are both *true*, then the conclusion also has to be *true* also, doesn't it? The problem is, however, that nothing in the *sentences* – the forms of words used – guarantees that "He" in the first premiss refers to the same "he" who is referred to in the second premiss, or in the conclusion. So here is a strong reason for identifying *propositions* rather than *sentences* as the constituents of arguments, so that we focus on *what each sentence says (about which objects at which time etc.)* rather than just the actual words used. Then if "He" in the first premiss refers to someone different, we can say that the argument does *not* have the valid argument-form above, because we don't have the same *"O"* in the two premisses: they might both contain the same *sentence*, but they don't express the same *proposition*.

Having acknowledged that this is indeed a strong reason for seeing arguments as composed of *propositions* (rather than *declarative sentences*) for the purpose of theoretical analysis, it is not entirely decisive. Propositions have the drawback of being unfamiliar and abstract: we haven't worked out yet what exactly a proposition *is*, and the matter is anyway potentially controversial, so it's a bit unfortunate if we have to sort that out before we start studying basic logic. *Declarative sentences* by contrast seem straightforward and familiar. Admittedly they encounter the context-relativity problem, but this (even if very serious at a theoretical level) is at least manageable in most practical cases, as long as we are allowed to presume that any argument is to be assessed in relation to a single *situation*, so that all referring expressions, including both token-reflexives ("I", "now" etc.) and demonstratives ("he", "that" etc.) are assumed to refer consistently throughout the argument. Given that this is practicable, there is also another reason for taking *declarative sentences* rather than *propositions* to be the constituents of the arguments that interest us: it goes back to the notion of "argument-forms" mentioned above, and is best explained by example.

If we look again at the argument in §1.5 above, we might notice that it has an abstract structure which could be in common to a large number of other arguments (here lower-case letters are used for names of people, subjects and places; upper-case letters for everything else):

1. *j* is *S*-ing *p* at *o*
2. Everyone who is *S*-ing at *o* either has *A* or is *M*
3. *j* has *R*
4. Nobody who has *R* is *M*
5. Everyone *S*-ing *p* at *o* who has *A*, has *L*

∴ *j* has *L*

In the original argument, roughly, *j* is replaced with "John", *S*-ing with "studying", *p* with "Physics", *o* with "Oxford University", *A* with "A levels", *M* with "a mature student", *R* with "a young person's railcard", and *L* with "a Mathematics A level". The argument is actually a strong one. But note that the force of the argument – the strength of the implication between the premisses and the conclusion – does not depend on these details being exactly as they are. There are countless other arguments with the same abstract structure, which would all be equally convincing: such that if their premisses were true, then their conclusion would also have to be true.

There are considerable difficulties in getting clear on what exactly constitutes the abstract structure – or "logical form" – of an argument. Is it, for example, part of the logical form of the argument above that reference is implicitly made to *people* through the words "Everyone", "Nobody", and "who"? Should these perhaps be replaced with "Everything", "Nothing", and "which", on the grounds that the latter are less topic-specific, and will allow us to include a great many more arguments within the same pattern? A more subtle problem concerns the verb *S*: for the argument to work, it has to be the case that "*j* is *S*-ing *p* at *o*" implies that "*j* is *S*-ing at *o*". But this principle sometimes fails, for example "John is growing vegetables at home" does not imply that "John is growing at home", and "John is contemplating study at Oxford" does not imply "John is contemplating at Oxford". The fascinating area of Philosophy often called "Philosophical Logic" (and the closely related "Philosophy of Language") is concerned with sorting out this kind of issue. If you want to take it further, the finals paper on "Philosophy of Logic and Language" is for you!

Another important point here is that a single sentence (and hence a single argument) can instantiate more than one logical form. Consider the following famous example:

> Socrates is human;
> If Socrates is human, then Socrates is mortal;
> therefore Socrates is mortal.

This seems to exemplify the following form (using capital letters here to represent sentences):

1. *H*
2. If *H* then *M*

∴ *M*

But suppose now we add an additional sentence, giving us:

> Socrates is human;
> Every human is mortal;
> If Socrates is human, then Socrates is mortal;
> therefore Socrates is mortal.

Now it looks as though the third line is an intermediate conclusion – inferred from the second line – rather than a premiss. This makes it desirable to represent the argument as having the logical form:

1. *s* is *H*
2. Everything which is *H* is *M*

∴ *s* is *M*

But note that we *could* represent the argument as having the following logical form, which it also exemplifies:

1. *H*
2. *E*
3. If *H* then *M*

∴ *M*

This is not exactly *incorrect*, but it is relatively *unilluminating*, because it makes the second premiss *E* ("Every human is mortal") entirely redundant, and it sheds no light on *why M* is being taken to follow from *H*. So often when we ask what is the "logical form" of an argument, there is no one uniquely correct answer: it is rather a matter of which logical connections between sentences are the most illuminating (and that can in turn depend on what our interests are).

All this gives another reason for preferring to consider arguments as composed of *declarative sentences* rather than *propositions*. Not only are sentences reassuringly straightforward and familiar, unlike the abstract and elusive propositions. But also, they are relatively easy to identify, without having to second-guess what "logical form" might end up being of interest to us. By contrast, if we attempt to identify arguments in terms of their *propositional* structure, we might have to make firm decisions in advance on which sentences express the same (or related) *propositions*, and there might be a significant worry that these decisions will prejudge subtle questions of "logical form" that are best addressed only after the argument has been identified. Hence, at long last, we can refine our provisional definition of an argument, to yield:

DEFINITION 1.8 *An argument consists of a set of declarative sentences (the premisses) and a declarative sentence (the conclusion) marked as the concluded sentence.*

The conclusion can be marked in a number of different ways (e.g. with "therefore" or "so", or using the "∴" symbol as above).

1.5.2 Logical Validity and Related Notions

We have already come across the informal notion of a *valid* argument: an argument in which the truth of the premisses guarantees the truth of the conclusion. (Another – often more useful – way of expressing this is that if the premisses are true, the conclusion cannot be false.) Here is an example of an argument which looks valid in this sense:

Neptune is blue and round

∴ Some round thing is coloured

There is no way that Neptune can be blue and round, without some round thing's being coloured. So if the premiss of this argument is true, the conclusion must also be true. However this argument is not valid in quite the same way as the following:

Neptune is blue and round

∴ Some round thing is blue

Not only does the truth of this argument's premiss guarantee the truth of its conclusion, but also, the argument has a *valid argument-form*:

n is *B* and *R*

∴ Some *R* thing is *B*

Whatever name we put in place of *n*, and *whatever* predicates we put in place of *B* and *R*, it will still be the case that the truth of the premiss guarantees the truth of the conclusion. So *whatever interpretation we give* to *n*, *B*, and *R* – the letters which mark the place of the subject-specific terms ("Neptune", "blue", and "round") in the original argument – the argument that results will be valid: such that it cannot have true premisses and a false conclusion. This idea, of all possible *interpretations* of subject-specific terms (as opposed to logical terms such as "is", "and", and "some", which are part of the form of the argument and are not reinterpreted), is the key to our characterisation of *logical validity*:

CHARACTERISATION 1.9 (LOGICAL VALIDITY) *An argument is logically valid if and only if there is no interpretation under which the premisses are all true and the conclusion is false.*[9]

Logical validity is very closely related to the notion of logical consistency, which we characterise as follows:

CHARACTERISATION 1.10 (CONSISTENCY) *A set of sentences is logically consistent if and only if there is at least one interpretation under which all sentences of the set are true.*[10]

To spell out the connection between logical validity and logical consistency, we need to introduce the notion of a sentence's *negation*:

NEGATION

The *negation* of a sentence is that sentence which denies precisely what the first asserts (and, therefore, asserts precisely what the first denies). So if a sentence is true, its negation must be false, and likewise, if a sentence is false, its negation must be true. The standard way of negating a sentence in English is to precede it with "It is not the case that …". So for example the negation of "All cows eat grass" would be "It is not the case that all cows eat grass". However it is often possible to negate a sentence more elegantly – thus "It is not the case that all cows eat grass" is equivalent to "Some cow does not eat grass".

Now suppose that we have an argument which is logically valid. This is equivalent to saying that there is no interpretation under which the argument's premisses are all *true* and the conclusion *false*. But this is in turn equivalent to saying that there is no interpretation under which the premisses are all *true* and the *negation* of the conclusion is also *true*. Which is equivalent to saying that the premisses, taken together with the negation of the conclusion, form an *inconsistent* set. So we get:

ALTERNATIVE CHARACTERISATION OF VALIDITY *An argument is valid if and only if the set obtained by adding the negation of the conclusion to the premisses is inconsistent.*

Some other important notions are also related to logical validity, as summarised below. For discussion, see Section §1.5 of the *Logic Manual*. Note also that the final section §1.6 of Chapter 1 of the *Logic Manual* gives a brief explanation of the terms *syntax*, *semantics*, and *pragmatics*, which will not be discussed here.

CHARACTERISATION 1.11 (LOGICAL TRUTH) *A sentence is logically true if and only if it is true under any interpretation.*

CHARACTERISATION 1.12 (CONTRADICTION) *A sentence is a contradiction if and only if it is false under any interpretation.*

CHARACTERISATION 1.13 (LOGICAL EQUIVALENCE) *Sentences are logically equivalent if and only if they are true under exactly the same interpretations.*

Peter Millican, September 2009

[9] This is called a "characterisation" rather than a "definition" to reflect the fact that it is imprecise, since the notion of an *interpretation* is not itself precisely defined.

[10] Just as with validity, there is also a less formal notion of *consistency* under which a set of sentences is consistent if it is possible for them all to be true together. Under this interpretation, the set containing the two sentences "All bachelors are tall" and "Some unmarried man is not tall" would be inconsistent, even though they are not *formally* inconsistent. However in what follows, "consistency" should be understood to mean *formal* consistency.