

2.8. Argument by Induction

A well-known principle of argument in elementary arithmetic is the principle of mathematical induction, which is this:

Suppose that the first number, 0, has a certain property; and suppose also that if any number has that property, then so does the next; then it follows that all numbers have the property.

When we speak of 'all numbers' here, we mean, of course, all those numbers that can be reached by starting with 0 and going on from each number to the next some finite number of times. For example, negative numbers, fractional numbers, or infinite numbers are not to count, but only those numbers that are called the 'natural' numbers. I take it that it is sufficiently obvious that the principle of induction is a correct principle for reasoning about such numbers.

The principle as I have just stated it is the *ordinary* principle of (mathematical) induction, but there is also another version of the principle, which is called the principle of complete induction. This is a bit more difficult to grasp. It may be stated thus:

Suppose that, for every number, if all the numbers less than it have a certain property, then so does it; then it follows that every number has the property.

It may help to see how this principle works if I begin by deducing it from the ordinary principle of induction.

Assume as a premiss, then, that for some property P ,

(1) For any number x , if every number less than x has P , then so does x .

We wish to show that it then follows that every number has the property P . We shall show first that it follows that, for every number x , every number less

than x has P . And we shall show this using the ordinary principle of induction. So we observe first:

(a) Every number less than 0 has P .

This is trivially true for every property, since there are no numbers less than 0. (Recall that negative numbers are not to count.) Here is something, then, that holds for 0. Next we let n be any arbitrary number and assume that the same thing holds for n , i.e. that

(b) Every number less than n has P .

Taking x as n in our premiss (1), it then follows that n also has P . But further, the numbers less than $n+1$ are just the numbers less than n together with n itself, so this shows

(c) Every number less than $n + 1$ has P .

This establishes the premisses for an ordinary inductive argument. For in (a) we have established our result for 0, and then in (b) we assumed the result for n and deduced in (c) that it must then hold for $n + 1$ as well. So by ordinary induction it holds for all numbers, i.e.

(2) For every number x , every number less than x has P .

Finally, from (1) and (2) together we at once have the desired conclusion

(3) For every number x , x has P .

Thus we have shown that (1) implies (3), and that is the principle of complete induction.

In logic we shall apply this principle in connection with some special uses of numbers, and the first of these is the use of numbers to measure what we call the *length* of a formula.³ (Other books sometimes speak instead of the *degree* of a formula.) By this we mean simply the *number of occurrences of truth-functors* in that formula. For example, the formula $\neg\neg\neg P$ has length 3, since it contains three occurrences of truth-functors (each an occurrence of the same functor); the formula $P \vee (Q \rightarrow R)$ has length 2, since it contains only two occurrences of truth-functors (each an occurrence of a different functor). (So the formula $\neg\neg\neg P$ counts as *longer than* the formula $P \vee (Q \rightarrow R)$, despite the fact that it occupies less space on paper.) Now suppose that we wish to show that all formulae have a certain property. Since every formula is of some definite length, it will evidently be enough to prove:

³ Later we shall also argue by induction on the length of a proof.

For all numbers x , all formulae of length x have the property.

In order to prove this by the principle of complete induction what we need to do is to set up the *inductive hypothesis*, for an arbitrary number n ,

- (1) For all numbers y less than n , all formulae of length y have the property.

From this hypothesis we then seek to deduce that

- (2) All formulae of length n have the property.

If the deduction succeeds, then—since n was arbitrary—we have established the premiss for the induction, and our result is therefore proved.

As a matter of fact we shall not set out our inductive arguments in quite this form, but in a simpler form which makes no explicit reference to numbers. For in order to deduce our consequence (2), what we need to do is to take an arbitrary formula ϕ of length n , and to show that it has the property. But the inductive hypothesis, from which we hope to deduce this, evidently tells us that all formulae of length less than ϕ have the property, or in other words that all formulae shorter than ϕ have the property. So this is the form in which the inductive hypothesis will actually be stated, namely

- (1') All formulae ψ shorter than ϕ have the property.

And from this hypothesis we then aim to deduce:

- (2') The formula ϕ has the property.

In the deduction, ϕ is to be any arbitrary formula. If the deduction succeeds, then it shows directly that (1') implies (2'), and hence indirectly that (1) implies (2). Since (1) implies (2) it follows, as we have seen, that all formulae of any length have the property, or simply,

All formulae have the property.

So we shall not actually mention particular numbers in the course of the argument, though it helps to think in terms of numbers when reflecting on why this form of argument is justified.

In these arguments by (complete) induction on the length of a formula, it will nearly always be necessary to distinguish, and to consider separately, a number of different cases. The different cases will be the different possibilities for the structure of the formula ϕ , and which these are will depend upon which formulae our argument is intended to cover, i.e. on which *language* is in question. Supposing that the language is that specified by the

formation rules set out on p. 21, there will be four basic kinds of formulae to consider:

- (1) where ϕ is a sentence-letter;
- (2) where ϕ is either \top or \perp ;
- (3) where ϕ is the negation of some formula, i.e. where ϕ is $\neg\psi$ for some formula ψ ;
- (4) where ϕ is the conjunction or disjunction of two other formulae, or some other two-place truth-function of them; i.e. where ϕ is $(\psi*\chi)$ for some formulae ψ and χ , and some two-place truth-functor $*$.

The formation rules specify that every formula must be of one or other of these four kinds, because every formula is constructed by some finite number of applications of the formation rules, and the last rule applied must be one of the four just listed.

In practice, we usually do not need to distinguish between cases (1) and (2), but can roll them together under the case where ϕ is an *atomic* formula, i.e. has no subformulae except itself. In this case, the inductive hypothesis is generally useless. For where ϕ is a sentence-letter there are no formulae shorter than ϕ , and where ϕ is \top or \perp it has no *components* shorter than itself, which has much the same effect. The inductive hypothesis is useful in cases (3) and (4). For if ϕ is $\neg\psi$, then since ψ is shorter than ϕ we can assume that the hypothesis holds of ψ ; and, similarly, if ϕ is $(\psi*\chi)$, then we can assume that it holds both of ψ and of χ . It is quite often the case that we need to consider separately the cases of $(\psi\wedge\chi)$, $(\psi\vee\chi)$, $(\psi\rightarrow\chi)$, and so on. It sometimes happens that we need to consider subcases within these, or (more usually) subcases within the case in which ϕ is $\neg\psi$. I proceed now to illustrate these remarks by several examples of inductive arguments. All the examples in this section will be arguments for results that should be obvious anyway. So the interest is in seeing the technique at work, and is not (yet) in the new results that one can prove by it.

I begin with a couple of theses that follow simply from the formation rules given in Section 2.3.

2.8.A. In any formula, the brackets pair uniquely.

By this I mean that there is one and only one function which pairs each left-hand bracket with just one right-hand bracket, and each right-hand bracket with just one left-hand bracket, and which satisfies the following further conditions: (1) the pair of the pair of any bracket is itself; (2) any left-hand bracket is always to the left of its paired right-hand bracket; (3) for any

paired brackets x and y , a bracket that lies between x and y is paired with another bracket that also lies between x and y , and (hence) a bracket that lies outside x and y is paired with another bracket that also lies outside x and y .

The proof is by induction on the length of an arbitrary formula ϕ . The inductive hypothesis is

In any formula shorter than ϕ , the brackets pair uniquely.

We distinguish cases thus:

Case (1): ϕ is atomic. Then ϕ contains no brackets, and there is nothing to prove.

Case (2): ϕ is $\neg\psi$. Then ψ is a formula shorter than ϕ , so by the inductive hypothesis the brackets in ψ pair uniquely. But the brackets in ϕ just are the brackets in ψ . So it follows that the brackets in ϕ pair uniquely.

Case (3): ϕ is $(\psi*\chi)$ for some two-place truth-functor $*$. Then ψ is a formula shorter than ϕ , so by the inductive hypothesis the brackets in ψ pair uniquely with one another. The same holds for the brackets in χ . The brackets in ϕ are just the brackets in ψ , and the brackets in χ , and the two outer brackets shown. So if we retain the pairing already given for ψ and for χ , and pair the two outer brackets with one another, then clearly we have a pairing that satisfies the conditions. But no other pairing would satisfy them. For by condition (3) the two outer brackets can only be paired with one another, and by condition (2) no bracket in ψ can be paired with any bracket in χ . For by this condition every right-hand bracket in ψ must still be paired with a bracket in ψ , and by the inductive hypothesis there is only one way of doing this, and it does not leave any spare left-hand bracket in ψ to be paired with something in χ .

This completes the induction, and so establishes the desired result for all formulae.

2.8.B. In any formula, the initial symbol is not a symbol of any of its proper subformulae.

A ‘proper’ subformula is a subformula other than, and so shorter than, the whole formula. By ‘the initial symbol’ I mean, more precisely, the first *occurrence* of a symbol. (For example $\neg P$ is a proper subformula of $\neg\neg P$, and both begin with a negation sign. But the first occurrence of the negation sign in $\neg\neg P$ does not fall within any of its proper subformulae. By contrast, the final occurrence of a symbol in $\neg\neg P$, namely the occurrence of P , is an occurrence that falls within both of its proper subformulae.)

The proof is by induction on the length of an arbitrary formula ϕ . The inductive hypothesis is

In any formula shorter than ϕ , the initial symbol is not a symbol of any proper subformula.

We distinguish cases as before:

Case (1): ϕ is atomic. Then ϕ has no proper subformula, so there is nothing to prove.

Case (2): ϕ is $\neg\psi$. Suppose that there were a proper subformula $\neg\chi$ of $\neg\psi$, beginning with the same occurrence of \neg . Then equally χ would be a proper subformula of ψ , beginning with the same occurrence of its initial symbol. But ψ is shorter than ϕ , and so by the inductive hypothesis this cannot happen.

Case (3): ϕ is $(\psi*\chi)$ for some two-place truth-functor $*$. By our first result, the outer brackets in ϕ pair with one another, and cannot be paired in any other way. Hence any subformula of ϕ that includes its first left-hand bracket, i.e. its first occurrence of a symbol, must also include its last right-hand bracket. But that is to say that it cannot be a *proper* subformula of ϕ , but must be the whole of ϕ .

This completes the induction, and so establishes the result.

I postpone to the exercises a further result about the syntax of our formulae, resting upon the results 2.8.A–B already reached. Here I turn to another example of an inductive argument, this time based upon the semantics introduced in Section 2.4.

2.8.C. In any interpretation I , every formula of the language interpreted is assigned one and only one truth-value.

Take any interpretation I of a language for truth-functors, and assume for simplicity that the language is specified by some subset of the formation rules on p. 21. Let ϕ be any formula in that language. We prove the result by induction on the length of ϕ . The inductive hypothesis is

To any formula shorter than ϕ , I assigns a unique truth-value.

We distinguish four cases according to the four kinds of formation rule on p. 21.

Case (1): ϕ is a sentence-letter. Then, by definition of an interpretation, I assigns to ϕ either T or F, but not both.

Case (2): ϕ is \top or \perp . Then it is stipulated that I assigns \top (and only \top) to \top , and assigns \perp (and only \perp) to \perp .

Case (3): ϕ is $\neg\psi$. By the inductive hypothesis I assigns either \top or \perp , but not both, to ψ . Then it is stipulated that I assigns to ϕ either \perp or \top (respectively), but not both.

Case (4): ϕ is $(\psi*\chi)$, for some two-place truth-functor $*$. By the inductive hypothesis I assigns unique truth-values to ψ and to χ . Hence, by whatever is the truth-table for $*$, it is stipulated that I also assigns a unique truth-value to ϕ .

This completes the induction.

Finally I take up an example from Section 2.6, namely the first step of the procedure for reducing any formula to an equivalent formula in DNF or in CNF. This argument should be compared with the informal argument given earlier.

2.8.D. If the language of ϕ contains no truth-functores other than $\top, \perp, \neg, \wedge, \vee$, then there is a formula logically equivalent to ϕ in which no occurrence of \neg has any other occurrence of a truth-functor in its scope.

We assume throughout that our language contains only the five truth-functores listed. To introduce a convenient abbreviation, let us call a formula ‘satisfactory’ if no occurrence of \neg in it has any other occurrence of a truth-functor in its scope. Then we have to show that any formula ϕ is logically equivalent to some satisfactory formula. The proof is by induction on the length of ϕ . The inductive hypothesis is

For any formula ψ shorter than ϕ there is a logically equivalent formula ψ' which is satisfactory.

We distinguish cases as expected, but within case (2) for negation we must further distinguish a variety of subcases. It is in these subcases that the work of the argument is done.

Case (1): ϕ is atomic. Then ϕ contains no occurrence of \neg , and therefore is already satisfactory.

Case (2): ϕ is $\neg\psi$. We distinguish five subcases, being the five possibilities for the formula ψ .

Subcase (a): ψ is a sentence-letter. Then ϕ , i.e. $\neg\psi$, is already satisfactory.

Subcase (b): ψ is \top or \perp . Then ϕ , i.e. $\neg\psi$, is equivalent either to \perp or to \top (respectively), and these are each satisfactory.

Subcase (c): ψ is $\neg\chi$. Then ϕ is logically equivalent to χ , and χ is shorter than ϕ . Hence by inductive hypothesis χ is logically equivalent to some satisfactory formula χ' . So ϕ is also equivalent to χ' .

Subcase (d): ψ is $\chi_1 \wedge \chi_2$. Then by De Morgan's law ϕ is logically equivalent to $\neg\neg\chi_1 \vee \neg\neg\chi_2$. Moreover, $\neg\neg\chi_1$ and $\neg\neg\chi_2$ are each shorter than ϕ . So by inductive hypothesis there are logically equivalent satisfactory formulae $(\neg\neg\chi_1)'$ and $(\neg\neg\chi_2)'$. Hence ϕ is logically equivalent to $(\neg\neg\chi_1)' \vee (\neg\neg\chi_2)'$, and this also is satisfactory.

Subcase (e): ψ is $\chi_1 \vee \neg\chi_2$. The argument is as in subcase (d), with \wedge and \vee interchanged.

Case (3): ϕ is $\psi * \chi$, for $*$ either \wedge or \vee . By inductive hypothesis there are logically equivalent formulae ψ' and χ' which are satisfactory. Hence ϕ is equivalent to $\psi' * \chi'$, which also is satisfactory.

In view of the restricted language that we are here considering, this completes the review of all possible cases, and so establishes our result.

The following exercises offer a few further examples where the reader may practise the technique of arguing by induction. But we are now in a position to turn to something more interesting, and in particular to take up once more the topic of expressive adequacy.

EXERCISES

2.8.1. Return to Exercise 2.3.1, and give an inductive argument for the result there stated.

2.8.2. Show that in every non-atomic formula there is one and only one occurrence of a truth-functor which does not fall within any proper subformula of that formula. (This is the *main* truth-functor of the formula.) [You will need to use the results 2.8.A and 2.8.B established in the text.]