CHAPTER II

Meaning and Form in Mathematics

THIS *Two-Part Invention* was the inspiration for my two characters. Just as Lewis Carroll took liberties with Zeno's Tortoise and Achilles, so have I taken liberties with Lewis Carroll's Tortoise and Achilles. In Carroll's dialogue, the same events take place over and over again, only each time on a higher and higher level; it is a wonderful analogue to Bach's Ever-Rising Canon. The Carrollian Dialogue, with its wit subtracted out, still leaves a deep philosophical problem: *Do words and thoughts follow formal rules, or do they not?* That problem is the problem of this book.

In this Chapter and the next, we will look at several new formal systems. This will give us a much wider perspective on the concept of formal system. By the end of these two Chapters, you should have quite a good idea of the power of formal systems, and why they are of interest to mathematicians and logicians.

The pq-System

The formal system of this Chapter is called the *pq-system.* It is not important to mathematicians or logicians—in fact, it is just a simple invention of mine. Its importance lies only in the fact that it provides an excellent example of many ideas that play a large role in this book. There are three distinct symbols of the pq-system:

p q -

—the letters p, q, and the hyphen.

The pq-system has an infinite number of axioms. Since we can't write them all down, we have to have some other way of describing what they are. Actually, we want more than just a description of the axioms; we want a way to tell whether some given string is an axiom or not. A mere description of axioms might characterize them fully and yet weakly—which was the problem with the way theorems in the MIU-system were characterized. We don't want to have to struggle for an indeterminate—possibly infinite—length of time, just to find out if some string is an axiom or not. Therefore, we will define axioms in such a way that there is an obvious decision procedure for axiomhood of a string composed of p's, q's, and hyphens.

DEFINITION: xp$-$q$x-$ is an axiom, whenever x is composed of hyphens only.

Note that 'x' must stand for the same string of hyphens in both occurrences. For example, --p-q--- is an axiom. The literal expression 'xp$-$q$x-$' is not an axiom, of course (because 'x' does not belong to the pq-system); it is more like a mold in which all axioms are cast—and it is called an *axiom schema.*

The pq-system has only one rule of production:

RULE: Suppose x, y, and z all stand for particular strings containing only hyphens. And suppose that xpyqz is known to be a theorem. Then xp$y-$q$z-$ is a theorem.

For example, take x to be '--', y to be '---', and z to be '-'. The rule tells us:

> If --p---q- turns out to be a theorem, then so will --p----q--.

As is typical of rules of production, the statement establishes a causal connection between the theoremhood of two strings, but without asserting theoremhood for either one on its own.

A most useful exercise for you is to find a decision procedure for the theorems of the pq-system. It is not hard; if you play around for a while, you will probably pick it up. Try it.

The Decision Procedure

I presume you have tried it. First of all, though it may seem too obvious to mention, I would like to point out that every theorem of the pq-system has three separate groups of hyphens, and the separating elements are one p, and one q, in that order. (This can be shown by an argument based on "heredity", just the way one could show that all MIU-system theorems had to begin with M.) This means that we can rule out, from its form alone, a string such as --p--p--p--q--------.

Now, stressing the phrase "from its form alone" may seem silly; what else is there to a string except its form? What else could possibly play a role in determining its properties? Clearly nothing could. But bear this in mind as the discussion of formal systems goes on; the notion of "form" will start to get rather more complicated and abstract, and we will have to think more about the meaning of the word "form". In any case, let us give the name *well-formed string* to any string which begins with a hyphen-group, then has one p, then has a second hyphen-group, then a q, and then a final hyphen-group.

Back to the decision procedure . . . The criterion for theoremhood is that the first two hyphen-groups should add up, in length, to the third

hyphen-group. For instance, --p--q---- is a theorem, since 2 plus 2 equals 4, whereas --p--q- is not, since 2 plus 2 is not 1. To see why this is the proper criterion, look first at the axiom schema. Obviously, it only manufactures axioms which satisfy the addition criterion. Second, look at the rule of production. If the first string satisfies the addition criterion, so must the second one—and conversely, if the first string does not satisfy the addition criterion, then neither does the second string. The rule makes the addition criterion into a hereditary property of theorems: any theorem passes the property on to its offspring. This shows why the addition criterion is correct.

There is, incidentally, a fact about the pq-system which would enable us to say with confidence that it has a decision procedure, even before finding the addition criterion. That fact is that the pq-system is not complicated by the opposing currents of *lengthening* and *shortening* rules; it has only lengthening rules. Any formal system which tells you how to make longer theorems from shorter ones, but never the reverse, has got to have a decision procedure for its theorems. For suppose you are given a string. First check whether it's an axiom or not (I am assuming that there is a decision procedure for axiomhood—otherwise, things are hopeless). If it is an axiom, then it is by definition a theorem, and the test is over. So suppose instead that it's not an axiom. Then, to be a theorem, it must have come from a shorter string, via one of the rules. By going over the various rules one by one, you can pinpoint not only the rules that could conceivably produce that string, but also exactly which shorter strings could be its forebears on the "family tree". In this way, you "reduce" the problem to determining whether any of several new but shorter strings is a theorem. Each of them can in turn be subjected to the same test. The worst that can happen is a proliferation of more and more, but shorter and shorter, strings to test. As you continue inching your way backwards in this fashion, you must be getting closer to the source of all theorems—the axiom schemata. You just can't get shorter and shorter indefinitely; therefore, eventually either you will find that one of your short strings is an axiom, or you'll come to a point where you're stuck, in that none of your short strings is an axiom, and none of them can be further shortened by running some rule or other backwards. This points out that there really is not much deep interest in formal systems with lengthening rules only; it is the interplay of lengthening and shortening rules that gives formal systems a certain fascination.

Bottom-up *vs.* Top-down

The method above might be called a *top-down* decision procedure, to be contrasted with a *bottom-up* decision procedure, which I give now. It is very reminiscent of the genie's systematic theorem-generating method for the MIU-system, but is complicated by the presence of an axiom schema. We are going to form a "bucket" into which we throw theorems as they are generated. Here is how it is done:

(1a) Throw the simplest possible axiom (−p−q−−) into the bucket.
(1b) Apply the rule of inference to the item in the bucket, and put the result into the bucket.

(2a) Throw the second-simplest axiom into the bucket.
(2b) Apply the rule to each item in the bucket, and throw all results into the bucket.

(3a) Throw the third-simplest axiom into the bucket.
(3b) Apply the rule to each item in the bucket, and throw all results into the bucket.

etc., etc.

A moment's reflection will show that you can't fail to produce every theorem of the pq-system this way. Moreover, the bucket is getting filled with longer and longer theorems, as time goes on. It is again a consequence of that lack of shortening rules. So if you have a particular string, such as −−p−−−q−−−−−, which you want to test for theoremhood, just follow the numbered steps, checking all the while for the string in question. If it turns up—theorem! If at some point everything that goes into the bucket is longer than the string in question, forget it—it is not a theorem. This decision procedure is *bottom-up* because it is working its way up from the basics, which is to say the axioms. The previous decision procedure is *top-down* because it does precisely the reverse: it works its way back down towards the basics.

Isomorphisms Induce Meaning

Now we come to a central issue of this Chapter—indeed of the book. Perhaps you have already thought to yourself that the pq-theorems are like additions. The string −−p−−−q−−−−− is a theorem because 2 plus 3 equals 5. It could even occur to you that the theorem −−p−−−q−−−−− is a *statement,* written in an odd notation, whose *meaning* is that 2 plus 3 is 5. Is this a reasonable way to look at things? Well, I deliberately chose 'p' to remind you of 'plus', and 'q' to remind you of 'equals' . . . So, does the string −−p−−−q−−−−− actually *mean* "2 plus 3 equals 5"?

What would make us feel that way? My answer would be that we have perceived an *isomorphism* between pq-theorems and additions. In the Introduction, the word "isomorphism" was defined as an information-preserving transformation. We can now go into that notion a little more deeply, and see it from another perspective. The word "isomorphism" applies when two complex structures can be mapped onto each other, in such a way that to each part of one structure there is a corresponding part in the other structure, where "corresponding" means that the two parts play similar roles in their respective structures. This usage of the word "isomorphism" is derived from a more precise notion in mathematics.

It is cause for joy when a mathematician discovers an isomorphism between two structures which he knows. It is often a "bolt from the blue", and a source of wonderment. The perception of an isomorphism between two known structures is a significant advance in knowledge—and I claim that it is such perceptions of isomorphism which create *meanings* in the minds of people. A final word on the perception of isomorphisms: since they come in many shapes and sizes, figuratively speaking, it is not always totally clear when you really have found an isomorphism. Thus, "isomorphism" is a word with all the usual vagueness of words—which is a defect but an advantage as well.

In this case, we have an excellent prototype for the concept of isomorphism. There is a "lower level" of our isomorphism—that is, a mapping between the parts of the two structures:

$$\begin{aligned} p &\Leftrightarrow \text{plus} \\ q &\Leftrightarrow \text{equals} \\ - &\Leftrightarrow \text{one} \\ -- &\Leftrightarrow \text{two} \\ --- &\Leftrightarrow \text{three} \\ &\text{etc.} \end{aligned}$$

This symbol-word correspondence has a name: *interpretation.*

Secondly, on a higher level, there is the correspondence between true statements and theorems. But—note carefully—this higher-level correspondence could not be perceived without the prior choice of an interpretation for the symbols. Thus it would be more accurate to describe it as a correspondence between true statements and *interpreted* theorems. In any case we have displayed a two-tiered correspondence, which is typical of all isomorphisms.

When you confront a formal system you know nothing of, and if you hope to discover some hidden meaning in it, your problem is how to assign interpretations to its symbols in a meaningful way—that is, in such a way that a higher-level correspondence emerges between true statements and theorems. You may make several tentative stabs in the dark before finding a good set of words to associate with the symbols. It is very similar to attempts to crack a code, or to decipher inscriptions in an unknown language like Linear B of Crete: the only way to proceed is by trial and error, based on educated guesses. When you hit a good choice, a "meaningful" choice, all of a sudden things just feel right, and work speeds up enormously. Pretty soon everything falls into place. The excitement of such an experience is captured in *The Decipherment of Linear B* by John Chadwick.

But it is uncommon, to say the least, for someone to be in the position of "decoding" a formal system turned up in the excavations of a ruined civilization! Mathematicians (and more recently, linguists, philosophers, and some others) are the only users of formal systems, and they invariably have an interpretation in mind for the formal systems which they use and publish. The idea of these people is to set up a formal system whose

theorems reflect some portion of reality isomorphically. In such a case, the choice of symbols is a highly motivated one, as is the choice of typographical rules of production. When I devised the pq-system, I was in this position. You see why I chose the symbols I chose. It is no accident that theorems are isomorphic to additions; it happened because I deliberately sought out a way to reflect additions typographically.

Meaningless and Meaningful Interpretations

You can choose interpretations other than the one I chose. You need not make every theorem come out true. But there would be very little reason to make an interpretation in which, say, all theorems came out false, and certainly even less reason to make an interpretation under which there is no correlation at all, positive or negative, between theoremhood and truth. Let us therefore make a distinction between two types of interpretations for a formal system. First, we can have a *meaningless* interpretation, one under which we fail to see any isomorphic connection between theorems of the system, and reality. Such interpretations abound—any random choice at all will do. For instance, take this one:

p ⇐⇒ horse
q ⇐⇒ happy
\- ⇐⇒ apple

Now -p-q-- acquires a new interpretation: "apple horse apple happy apple apple"—a poetic sentiment, which might appeal to horses, and might even lead them to favor this mode of interpreting pq-strings! However, this interpretation has very little "meaningfulness"; under interpretation, theorems don't sound any truer, or any better, than nontheorems. A horse might enjoy "happy happy happy apple horse" (mapped onto qqq-p) just as much as any interpreted theorem.

The other kind of interpretation will be called *meaningful.* Under such an interpretation, theorems and truths correspond—that is, an isomorphism exists between theorems and some portion of reality. That is why it is good to distinguish between *interpretations* and *meanings.* Any old word can be used as an interpretation for 'p', but 'plus' is the only *meaningful* choice we've come up with. In summary, the meaning of 'p' seems to be 'plus', though it can have a million different interpretations.

Active *vs.* Passive Meanings

Probably the most significant fact of this Chapter, if understood deeply, is this: the pq-system seems to force us into recognizing that *symbols of a formal system, though initially without meaning, cannot avoid taking on "meaning" of sorts, at least if an isomorphism is found.* The difference between meaning in a formal system and in a language is a very important one, however. It is this:

in a language, when we have learned a meaning for a word, we then make new statements based on the meaning of the word. In a sense the meaning becomes *active,* since it brings into being a new rule for creating sentences. This means that our command of language is not like a finished product: the rules for making sentences increase when we learn new meanings. On the other hand, in a formal system, the theorems are predefined, by the rules of production. We can choose "meanings" based on an isomorphism (if we can find one) between theorems and true statements. But this does not give us the license to go out and add new theorems to the established theorems. That is what the Requirement of Formality in Chapter I was warning you of.

In the MIU-system, of course, there was no temptation to go beyond the four rules, because no interpretation was sought or found. But here, in our new system, one might be seduced by the newly found "meaning" of each symbol into thinking that the string

--p--p--p--q--------

is a theorem. At least, one might *wish* that this string were a theorem. But wishing doesn't change the fact that it isn't. And it would be a serious mistake to think that it "must" be a theorem, just because 2 plus 2 plus 2 plus 2 equals 8. It would even be misleading to attribute it any meaning at all, since it is not well-formed, and our meaningful interpretation is entirely derived from looking at well-formed strings.

In a formal system, the meaning must remain *passive;* we can read each string according to the meanings of its constituent symbols, but we do not have the right to create new theorems purely on the basis of the meanings we've assigned the symbols. Interpreted formal systems straddle the line between systems without meaning, and systems with meaning. Their strings can be thought of as "expressing" things, but this must come only as a consequence of the formal properties of the system.

Double-Entendre!

And now, I want to destroy any illusion about having found *the* meanings for the symbols of the pq-system. Consider the following association:

p ⇐⇒ equals
q ⇐⇒ taken from
- ⇐⇒ one
-- ⇐⇒ two
etc.

Now, --p---q----- has a new interpretation: "2 equals 3 taken from 5". Of course it is a true statement. All theorems will come out true under this new interpretation. It is just as meaningful as the old one. Obviously, it is silly to ask, "But which one is *the* meaning of the string?" An interpreta-

tion will be meaningful to the extent that it accurately reflects some isomorphism to the real world. When different aspects of the real world are isomorphic to each other (in this case, additions and subtractions), one single formal system can be isomorphic to both, and therefore can take on two passive meanings. This kind of double-valuedness of symbols and strings is an extremely important phenomenon. Here it seems trivial, curious, annoying. But it will come back in deeper contexts and bring with it a great richness of ideas.

Here is a summary of our observations about the pq-system. Under either of the two meaningful interpretations given, every well-formed string has a grammatical assertion for its counterpart—some are true, some false. The idea of *well-formed strings* in any formal system is that they are those strings which, when interpreted symbol for symbol, yield *grammatical* sentences. (Of course, it depends on the interpretation, but usually, there is one in mind.) Among the well-formed strings occur the theorems. These are defined by an axiom schema, and a rule of production. My goal in inventing the pq-system was to imitate additions: I wanted every theorem to express a true addition under interpretation; conversely, I wanted every true addition of precisely two positive integers to be translatable into a string, which would be a theorem. That goal was achieved. Notice, therefore, that all false additions, such as "2 plus 3 equals 6", are mapped into strings which are well-formed, but which are not theorems.

Formal Systems and Reality

This is our first example of a case where a formal system is based upon a portion of reality, and seems to mimic it perfectly, in that its theorems are isomorphic to truths about that part of reality. However, reality and the formal system are independent. Nobody need be aware that there is an isomorphism between the two. Each side stands by itself—one plus one equals two, whether or not we know that -p-q-- is a theorem; and -p-q-- is still a theorem whether or not we connect it with addition.

You might wonder whether making this formal system, or any formal system, sheds new light on truths in the domain of its interpretation. Have we learned any new additions by producing pq-theorems? Certainly not; but we have learned something about the nature of addition as a process—namely, that it is easily mimicked by a typographical rule governing meaningless symbols. This still should not be a big surprise since addition is such a simple concept. It is a commonplace that addition can be captured in the spinning gears of a device like a cash register.

But it is clear that we have hardly scratched the surface, as far as formal systems go; it is natural to wonder about what portion of reality can be imitated in its behavior by a set of meaningless symbols governed by formal rules. Can all of reality be turned into a formal system? In a very broad sense, the answer might appear to be yes. One could suggest, for instance, that reality is itself nothing but one very complicated formal

system. Its symbols do not move around on paper, but rather in a three-dimensional vacuum (space); they are the elementary particles of which everything is composed. (Tacit assumption: that there is an end to the descending chain of matter, so that the expression "elementary particles" makes sense.) The "typographical rules" are the laws of physics, which tell how, given the positions and velocities of all particles at a given instant, to modify them, resulting in a new set of positions and velocities belonging to the "next" instant. So the theorems of this grand formal system are the possible configurations of particles at different times in the history of the universe. The sole axiom is (or perhaps, *was*) the original configuration of all the particles at the "beginning of time". This is so grandiose a conception, however, that it has only the most theoretical interest; and besides, quantum mechanics (and other parts of physics) casts at least some doubt on even the theoretical worth of this idea. Basically, we are asking if the universe operates deterministically, which is an open question.

Mathematics and Symbol Manipulation

Instead of dealing with such a big picture, let's limit ourselves to *mathematics* as our "real world". Here, a serious question arises: How can we be sure, if we've tried to model a formal system on some part of mathematics, that we've done the job accurately—especially if we're not one hundred per cent familiar with that portion of mathematics already? Suppose the goal of the formal system is to bring us new knowledge in that discipline. How will we know that the interpretation of every theorem is true, unless we've proven that the isomorphism is perfect? And how will we prove that the isomorphism is perfect, if we don't already know all about the truths in the discipline to begin with?

Suppose that in an excavation somewhere, we actually did discover some mysterious formal system. We would try out various interpretations and perhaps eventually hit upon one which seemed to make every theorem come out true, and every nontheorem come out false. But this is something which we could only check directly in a finite number of cases. The set of theorems is most likely infinite. How will we *know* that all theorems express truths under this interpretation, unless we know everything there is to know about both the formal system and the corresponding domain of interpretation?

It is in somewhat this odd position that we will find ourselves when we attempt to match the reality of natural numbers (i.e., the nonnegative integers: 0, 1, 2, . . .) with the typographical symbols of a formal system. We will try to understand the relationship between what we call "truth" in number theory and what we can get at by symbol manipulation.

So let us briefly look at the basis for calling some statements of number theory true, and others false. How much is 12 times 12? Everyone knows it is 144. But how many of the people who give that answer have actually at

any time in their lives drawn a 12 by 12 rectangle, and then counted the little squares in it? Most people would regard the drawing and counting as unnecessary. They would instead offer as proof a few marks on paper, such as are shown below:

$$\begin{array}{r} 12 \\ \times 12 \\ \hline 24 \\ 12 \\ \hline 144 \end{array}$$

And that would be the "proof". Nearly everyone believes that if you counted the squares, you would get 144 of them; few people feel that the outcome is in doubt.

The conflict between the two points of view comes into sharper focus when you consider the problem of determining the value of $987654321 \times 123456789$. First of all, it is virtually impossible to construct the appropriate rectangle; and what is worse, even if it *were* constructed, and huge armies of people spent centuries counting the little squares, only a very gullible person would be willing to believe their final answer. It is just too likely that somewhere, somehow, somebody bobbled just a little bit. So is it ever possible to know what the answer is? If you trust the symbolic process which involves manipulating digits according to certain simple rules, yes. That process is presented to children as a device which gets the right answer; lost in the shuffle, for many children, are the rhyme and reason of that process. The digit-shunting laws for multiplication are based mostly on a few properties of addition and multiplication which are assumed to hold for all numbers.

The Basic Laws of Arithmetic

The kind of assumption I mean is illustrated below. Suppose that you lay down a few sticks:

/ // /// // / /

Now you count them. At the same time, somebody else counts them, but starting from the other end. Is it clear that the two of you will get the same answer? The result of a counting process is independent of the way in which it is done. This is really an assumption about what counting is. It would be senseless to try to prove it, because it is so basic; either you see it or you don't—but in the latter case, a proof won't help you a bit.

From this kind of assumption, one can get to the commutativity and associativity of addition (i.e., first that $b + c = c + b$ always, and second that $b + (c + d) = (b + c) + d$ always). The same assumption can also lead you to the commutativity and associativity of multiplication; just think of

many cubes assembled to form a large rectangular solid. Multiplicative commutativity and associativity are just the assumptions that when you rotate the solid in various ways, the number of cubes will not change. Now these assumptions are not verifiable in all possible cases, because the number of such cases is infinite. We take them for granted; we believe them (if we ever think about them) as deeply as we could believe anything. The amount of money in our pocket will not change as we walk down the street, jostling it up and down; the number of books we have will not change if we pack them up in a box, load them into our car, drive one hundred miles, unload the box, unpack it, and place the books in a new shelf. All of this is part of what we mean by *number*.

There are certain types of people who, as soon as some undeniable fact is written down, find it amusing to show why that "fact" is false after all. I am such a person, and as soon as I had written down the examples above involving sticks, money, and books, I invented situations in which they were wrong. You may have done the same. It goes to show that numbers as abstractions are really quite different from the everyday numbers which we use.

People enjoy inventing slogans which violate basic arithmetic but which illustrate "deeper" truths, such as "1 and 1 make 1" (for lovers), or "1 plus 1 plus 1 equals 1" (the Trinity). You can easily pick holes in those slogans, showing why, for instance, using the plus-sign is inappropriate in both cases. But such cases proliferate. Two raindrops running down a window-pane merge; does one plus one make one? A cloud breaks up into two clouds—more evidence for the same? It is not at all easy to draw a sharp line between cases where what is happening could be called "addition", and where some other word is wanted. If you think about the question, you will probably come up with some criterion involving separation of the objects in space, and making sure each one is clearly distinguishable from all the others. But then how could one count ideas? Or the number of gases comprising the atmosphere? Somewhere, if you try to look it up, you can probably find a statement such as, "There are 17 languages in India, and 462 dialects." There is something strange about precise statements like that, when the concepts "language" and "dialect" are themselves fuzzy.

Ideal Numbers

Numbers as realities misbehave. However, there is an ancient and innate sense in people that numbers ought not to misbehave. There is something clean and pure in the abstract notion of number, removed from counting beads, dialects, or clouds; and there ought to be a way of talking about numbers without always having the silliness of reality come in and intrude. The hard-edged rules that govern "ideal" numbers constitute arithmetic, and their more advanced consequences constitute number theory. There is only one relevant question to be asked, in making the transition from numbers as practical things to numbers as formal things. Once you have

FIGURE 13. Liberation, *by M. C. Escher (lithograph, 1955).*

decided to try to capsulize all of number theory in an ideal system, is it really possible to do the job completely? Are numbers so clean and crystalline and regular that their nature can be completely captured in the rules of a formal system? The picture *Liberation* (Fig. 13), one of Escher's most beautiful, is a marvelous contrast between the formal and the informal, with a fascinating transition region. Are numbers really as free as birds? Do they suffer as much from being crystallized into a rule-obeying system? Is there a magical transition region between numbers in reality and numbers on paper?

When I speak of the properties of natural numbers, I don't just mean properties such as the sum of a particular pair of integers. That can be found out by counting, and anybody who has grown up in this century cannot doubt the mechanizability of such processes as counting, adding, multiplying, and so on. I mean the kinds of properties which mathematicians are interested in exploring, questions for which no counting-process is sufficient to provide the answer—not even theoretically sufficient. Let us take a classic example of such a property of natural numbers. The statement is: "There are infinitely many prime numbers." First of all, there is no counting process which will ever be able to confirm, or refute, this assertion. The best we could do would be to count primes for a while and concede that there are "a lot". But no amount of counting alone would ever resolve the question of whether the number of primes is finite or infinite. There could always be more. The statement—and it is called "Euclid's Theorem" (notice the capital "T")—is quite unobvious. It may seem reasonable, or appealing, but it is not obvious. However, mathematicians since Euclid have always called it true. What is the reason?

Euclid's Proof

The reason is that *reasoning* tells them it is so. Let us follow the reasoning involved. We will look at a variant of Euclid's proof. This proof works by showing that whatever number you pick, there is a prime larger than it. Pick a number—N. Multiply all the positive integers starting with 1 and ending with N; in other words, form the factorial of N, written "$N!$". What you get is divisible by every number up to N. When you add 1 to $N!$, the result

can't be a multiple of 2 (because it leaves 1 over, when you divide by 2);

can't be a multiple of 3 (because it leaves 1 over, when you divide by 3);

can't be a multiple of 4 (because it leaves 1 over, when you divide by 4);

.
.
.

can't be a multiple of N (because it leaves 1 over, when you divide by N);

In other words, $N! + 1$, if it is divisible at all (other than by 1 and itself), only is divisible by numbers greater than N. So either it is itself prime, or its prime divisors are greater than N. But in either case we've shown there must exist a prime above N. The process holds no matter what number N is. Whatever N is, there is a prime greater than N. And thus ends the demonstration of the infinitude of the primes.

This last step, incidentally, is called *generalization,* and we will meet it again later in a more formal context. It is where we phrase an argument in terms of a single number (N), and then point out that N was unspecified and therefore the argument is a general one.

Euclid's proof is typical of what constitutes "real mathematics". It is simple, compelling, and beautiful. It illustrates that by taking several rather short steps one can get a long way from one's starting point. In our case, the starting points are basic ideas about multiplication and division and so forth. The short steps are the steps of reasoning. And though every individual step of the reasoning seems obvious, the end result is not obvious. We can never check directly whether the statement is true or not; yet we believe it, because we believe in reasoning. If you accept reasoning, there seems to be no escape route; once you agree to hear Euclid out, you'll have to agree with his conclusion. That's most fortunate—because it means that mathematicians will always agree on what statements to label "true", and what statements to label "false".

This proof exemplifies an orderly thought process. Each statement is related to previous ones in an irresistible way. This is why it is called a "proof'' rather than just "good evidence". In mathematics the goal is always to give an ironclad proof for some unobvious statement. The very fact of the steps being linked together in an ironclad way suggests that there may be a *patterned structure* binding these statements together. This structure can best be exposed by finding a new vocabulary—a stylized vocabulary, consisting of symbols—suitable only for expressing statements about numbers. Then we can look at the proof as it exists in its translated version. It will be a set of statements which are related, line by line, in some detectable way. But the statements, since they're represented by means of a small and stylized set of symbols, take on the aspect of *patterns.* In other words, though when read aloud, they seem to be statements about numbers and their properties, still when looked at on paper, they seem to be abstract patterns—and the line-by-line structure of the proof may start to look like a slow transformation of patterns according to some few typographical rules.

Getting Around Infinity

Although Euclid's proof is a proof that *all* numbers have a certain property, it avoids treating each of the infinitely many cases separately. It gets around

it by using phrases like "whatever N is", or "no matter what number N is". We could also phrase the proof over again, so that it uses the phrase "all N". By knowing the appropriate context and correct ways of using such phrases, we never have to deal with infinitely many statements. We deal with just two or three concepts, such as the word "all"—which, though themselves finite, embody an infinitude; and by using them, we sidestep the apparent problem that there are an infinite number of facts we want to prove.

We use the word "all" in a few ways which are defined by the thought processes of reasoning. That is, there are *rules* which our usage of "all" obeys. We may be unconscious of them, and tend to claim we operate on the basis of the *meaning* of the word; but that, after all, is only a circumlocution for saying that we are guided by rules which we never make explicit. We have used words all our lives in certain patterns, and instead of calling the patterns "rules", we attribute the courses of our thought processes to the "meanings" of words. That discovery was a crucial recognition in the long path towards the formalization of number theory.

If we were to delve into Euclid's proof more and more carefully, we would see that it is composed of many, many small—almost infinitesimal—steps. If all those steps were written out line after line, the proof would appear incredibly complicated. To our minds it is clearest when several steps are telescoped together, to form one single sentence. If we tried to look at the proof in slow motion, we would begin to discern individual frames. In other words, the dissection can go only so far, and then we hit the "atomic" nature of reasoning processes. A proof can be broken down into a series of tiny but discontinuous jumps which seem to flow smoothly when perceived from a higher vantage point. In Chapter VIII, I will show one way of breaking the proof into atomic units, and you will see how incredibly many steps are involved. Perhaps it should not surprise you, though. The operations in Euclid's brain when he invented the proof must have involved millions of neurons (nerve cells), many of which fired several hundred times in a single second. The mere utterance of a sentence involves hundreds of thousands of neurons. If Euclid's thoughts were that complicated, it makes sense for his proof to contain a huge number of steps! (There may be little direct connection between the neural actions in his brain, and a proof in our formal system, but the complexities of the two are comparable. It is as if nature wants the complexity of the proof of the infinitude of primes to be conserved, even when the systems involved are very different from each other.)

In Chapters to come, we will lay out a formal system that (1) includes a stylized vocabulary in which all statements about natural numbers can be expressed, and (2) has rules corresponding to all the types of reasoning which seem necessary. A very important question will be whether the rules for symbol manipulation which we have then formulated are really of equal power (as far as number theory is concerned) to our usual mental reasoning abilities—or, more generally, whether it is theoretically possible to attain the level of our thinking abilities, by using some formal system.