GÖDEL, ESCHER, BACH:

an Eternal Golden Braid

MU-puzzle solution & discussion

Douglas R. Hofstadter

Basic Books, Inc., Publishers

NEW YORK

© 1979

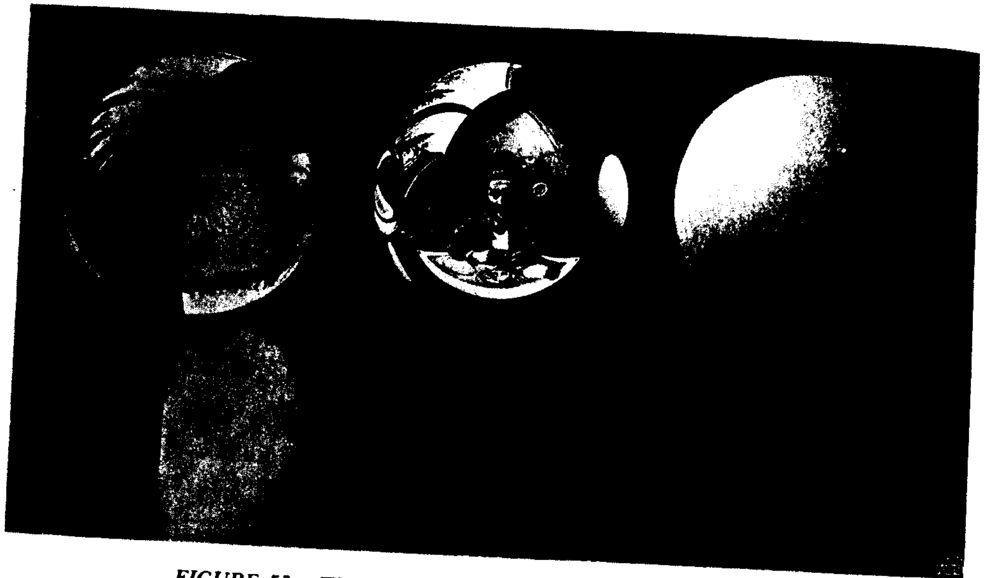

FIGURE 53. Three Spheres II, *by M. C. Escher (lithograph, 1946).*

Indra's Net

Finally, consider *Three Spheres II* (Fig. 53), in which every part of the world seems to contain, and be contained in, every other part: the writing table reflects the spheres on top of it, the spheres reflect each other, as well as the writing table, the drawing of them, and the artist drawing it. The endless connections which all things have to each other is only hinted at here, yet the hint is enough. The Buddhist allegory of "Indra's Net" tells of an endless net of threads throughout the universe, the horizontal threads running through space, the vertical ones through time. At every crossing of threads is an individual, and every individual is a crystal bead. The great light of "Absolute Being" illuminates and penetrates every crystal bead; moreover, every crystal bead reflects not only the light from every other crystal in the net—but also every reflection of every reflection throughout the universe.

To my mind, this brings forth an image of renormalized particles: in every electron, there are virtual photons, positrons, neutrinos, muons . . . ; in every photon, there are virtual electrons, protons, neutrons, pions . . . ; in every pion, there are . . .

But then another image rises: that of people, each one reflected in the minds of many others, who in turn are mirrored in yet others, and so on.

Both of these images could be represented in a concise, elegant way by using Augmented Transition Networks. In the case of particles, there would be one network for each category of particle; in the case of people, one for each person. Each one would contain calls to many others, thus creating a virtual cloud of ATN's around each ATN. Calling one would create calls on others, and this process might cascade arbitrarily far, until it bottomed out.

Mumon on MU

Let us conclude this brief excursion into Zen by returning to Mumon. Here is his comment on Jōshū's MU:[13]

> To realize Zen one has to pass through the barrier of the patriarchs. Enlightenment always comes after the road of thinking is blocked. If you do not pass the barrier of the patriarchs or if your thinking road is not blocked, whatever you think, whatever you do, is like a tangling ghost. You may ask: "What is a barrier of a patriarch?" This one word, 'MU', is it.
>
> This is the barrier of Zen. If you pass through it, you will see Jōshū face to face. Then you can work hand in hand with the whole line of patriarchs. Is this not a pleasant thing to do?
>
> If you want to pass this barrier, you must work through every bone in your body, through every pore of your skin, filled with this question: "What is 'MU'?" and carry it day and night. Do not believe it is the common negative symbol meaning nothing. It is not nothingness, the opposite of existence. If you really want to pass this barrier, you should feel like drinking a hot iron ball that you can neither swallow nor spit out.
>
> Then your previous lesser knowledge disappears. As a fruit ripening in season, your subjectivity and objectivity naturally become one. It is like a dumb man who has had a dream. He knows about it but he cannot tell it.
>
> When he enters this condition his ego-shell is crushed and he can shake the heaven and move the earth. He is like a great warrior with a sharp sword. If a Buddha stands in his way, he will cut him down; if a patriarch offers him any obstacle, he will kill him; and he will be free in his way of birth and death. He can enter any world as if it were his own playground. I will tell you how to do this with this kōan:
>
> Just concentrate your whole energy into this MU, and do not allow any discontinuation. When you enter this MU and there is no discontinuation, your attainment will be as a candle burning and illuminating the whole universe.

From Mumon to the MU-puzzle

From the ethereal heights of Jōshū's MU, we now descend to the prosaic lowlinesses of Hofstadter's **MU** . . . I know that you have already concentrated your whole energy into this **MU** (when you read Chapter I). So now I wish to answer the question which was posed there:

Has **MU** theorem-nature, or not?

The answer to this question is not an evasive MU; rather, it is a resounding NO. In order to show this, we will take advantage of dualistic, logical thinking.

We made two crucial observations in Chapter I:

(1) that the MU-puzzle has depth largely because it involves the interplay of lengthening and shortening rules;

(2) that hope nevertheless exists for cracking the problem by employing a tool which is in some sense of adequate depth to handle matters of that complexity: the theory of numbers.

We did not analyze the MU-puzzle in those terms very carefully in Chapter I; we shall do so now. And we will see how the second observation (when generalized beyond the insignificant MIU-system) is one of the most fruitful realizations of all mathematics, and how it changed mathematicians' view of their own discipline.

For your ease of reference, here is a recapitulation of the MIU-system:

SYMBOLS: M, I, U

AXIOM: MI

RULES:

I. If xI is a theorem, so is xIU.

II. If Mx is a theorem, so is Mxx.

III. In any theorem, III can be replaced by U.

IV. UU can be dropped from any theorem.

Mumon Shows Us How to Solve the MU-puzzle

According to the observations above, then, the MU-puzzle is merely a puzzle about natural numbers in typographical disguise. If we could only find a way to transfer it to the domain of number theory, we might be able to solve it. Let us ponder the words of Mumon, who said, "If any of you has one eye, he will see the failure on the teacher's part." But why should it matter to have one eye?

If you try counting the number of I's contained in theorems, you will soon notice that it seems never to be 0. In other words, it seems that no matter how much lengthening and shortening is involved, we can never work in such a way that all I's are eliminated. Let us call the number of I's in any string the *I-count* of that string. Note that the I-count of the axiom MI is 1. We can do more than show that the I-count can't be 0—we can show that the I-count can never be any multiple of 3.

To begin with, notice that rules I and IV leave the I-count totally undisturbed. Therefore we need only think about rules II and III. As far as rule III is concerned, it diminishes the I-count by exactly 3. After an application of this rule, the I-count of the output might conceivably be a multiple of 3—but only if the I-count of the *input* was also. Rule III, in short, never creates a multiple of 3 from scratch. It can only create one when it began with one. The same holds for rule II, which doubles the

I-count. The reason is that if 3 divides $2n$, then—because 3 does not divide 2—it must divide n (a simple fact from the theory of numbers). Neither rule II nor rule III can create a multiple of 3 from scratch.

But this is the key to the MU-puzzle! Here is what we know:

(1) The I-count begins at 1 (not a multiple of 3);

(2) Two of the rules do not affect the I-count at all;

(3) The two remaining rules which do affect the I-count do so in such a way as never to create a multiple of 3 unless given one initially.

The conclusion—and a typically hereditary one it is, too—is that the I-count can never become any multiple of 3. In particular, 0 is a forbidden value of the I-count. Hence, MU *is not a theorem of the MIU-system.*

Notice that, even as a puzzle about I-counts, this problem was still plagued by the crossfire of lengthening and shortening rules. Zero became the goal; I-counts could increase (rule II), could decrease (rule III). Until we analyzed the situation, we might have thought that, with enough switching back and forth between the rules, we might eventually hit 0. Now, thanks to a simple number-theoretical argument, we know that that is impossible.

Gödel-Numbering the MIU-System

Not all problems of the the type which the MU-puzzle symbolizes are so easy to solve as this one. But we have seen that at least one such puzzle could be embedded within, and solved within, number theory. We are now going to see that there is a way to embed *all* problems about *any* formal system, in number theory. This can happen thanks to the discovery, by Gödel, of a special kind of isomorphism. To illustrate it, I will use the MIU-system.

We begin by considering the notation of the MIU-system. We shall map each symbol onto a new symbol:

M ⇐⇒ 3
I ⇐⇒ 1
U ⇐⇒ 0

The correspondence was chosen arbitrarily; the only rhyme or reason to it is that each symbol looks a little like the one it is mapped onto. Each number is called the *Gödel number* of the corresponding letter. Now I am sure you can guess what the Gödel number of a multiletter string will be:

MU ⇐⇒ 30
MIIU ⇐⇒ 3110
etc.

It is easy. Clearly this mapping between notations is an information-preserving transformation; it is like playing the same melody on two different instruments.

Let us now take a look at a typical derivation in the MIU-system, written simultaneously in both notations:

(1)	MI	axiom	31
(2)	MII	rule 2	311
(3)	MIIII	rule 2	31111
(4)	MUI	rule 3	301
(5)	MUIU	rule 1	3010
(6)	MUIUUIU	rule 2	3010010
(7)	MUIIU	rule 4	30110

The left-hand column is obtained by applying our four familiar typographical rules. The right-hand column, too, could be thought of as having been generated by a similar set of typographical rules. Yet the right-hand column has a dual nature. Let me explain what this means.

Seeing Things Both Typographically and Arithmetically

We could say of the fifth string ('3010') that it was made from the fourth, by appending a '0' on the right; on the other hand we could equally well view the transition as caused by an *arithmetical* operation—multiplication by 10, to be exact. When natural numbers are written in the decimal system, multiplication by 10 and putting a '0' on the right are indistinguishable operations. We can take advantage of this to write an *arithmetical* rule which corresponds to typographical rule I:

ARITHMETICAL RULE Ia: A number whose decimal expansion ends on the right in '1' can be multiplied by 10.

We can eliminate the reference to the symbols in the decimal expansion by arithmetically describing the rightmost digit:

ARITHMETICAL RULE Ib: A number whose remainder when divided by 10 is 1, can be multiplied by 10.

Now we could have stuck with a purely typographical rule, such as the following one:

TYPOGRAPHICAL RULE I: From any theorem whose rightmost symbol is '1' a new theorem can be made, by appending '0' to the right of that '1'.

They would have the same effect. This is why the right-hand column has a "dual nature": it can be viewed either as a series of typographical operations changing one pattern of symbols into another, or as a series of arithmetical operations changing one magnitude into another. But there are powerful reasons for being more interested in the arithmetical version. Stepping out of one purely typographical system into another isomorphic typographical system is not a very exciting thing to do; whereas stepping clear out of the typographical domain into an isomorphic part of number theory has some kind of unexplored potential. It is as if somebody had known musical scores all his life, but purely visually—and then, all of a sudden, someone introduced him to the mapping between sounds and musical scores. What a rich, new world! Then again, it is as if somebody had been familiar with string figures all his life, but purely as string figures, devoid of meaning—and then, all of a sudden, someone introduced him to the mapping between stories and strings. What a revelation! The discovery of Gödel-numbering has been likened to the discovery, by Descartes, of the isomorphism between curves in a plane and equations in two variables: incredibly simple, once you see it—and opening onto a vast new world.

Before we jump to conclusions, though, perhaps you would like to see a more complete rendering of this higher level of the isomorphism. It is a very good exercise. The idea is to give an arithmetical rule whose action is indistinguishable from that of each typographical rule of the MIU-system.

A solution is given below. In the rules, m and k are arbitrary natural numbers, and n is any natural number which is less than 10^m.

RULE 1: If we have made $10m + 1$, then we can make $10 \times (10m + 1)$.

Example: Going from line 4 to line 5. Here, $m = 30$.

RULE 2: If we have made $3 \times 10^m + n$, then we can make $10^m \times (3 \times 10^m + n) + n$.

Example: Going from line 1 to line 2, where both m and n equal 1.

RULE 3: If we have made $k \times 10^{m+3} + 111 \times 10^m + n$, then we can make $k \times 10^{m+1} + n$.

Example: Going from line 3 to line 4. Here, m and n are 1, and k is 3.

RULE 4: If we have made $k \times 10^{m+2} + n$, then we can make $k \times 10^m + n$.

Example: Going from line 6 to line 7. Here, $m = 2$, $n = 10$, and $k = 301$.

Let us not forget our axiom! Without it we can go nowhere. Therefore, let us postulate that:

We can make 31.

Now the right-hand column can be seen as a full-fledged arithmetical process, in a new arithmetical system which we might call the *310-system:*

(1)	31	given
(2)	311	rule 2 ($m=1$, $n=1$)
(3)	31111	rule 2 ($m=2$, $n=11$)
(4)	301	rule 3 ($m=1$, $n=1$, $k=3$)
(5)	3010	rule 1 ($m=30$)
(6)	3010010	rule 2 ($m=3$, $n=10$)
(7)	30110	rule 4 ($m=2$, $n=10$, $k=301$)

Notice once again that the lengthening and shortening rules are ever with us in this "310-system"; they have merely been transposed into the domain of numbers, so that the Gödel numbers go up and down. If you look carefully at what is going on, you will discover that the rules are based on nothing more profound than the idea that shifting digits to left and right in decimal representations of integers is related to multiplications and divisions by powers of 10. This simple observation finds its generalization in the following

CENTRAL PROPOSITION: If there is a typographical rule which tells how certain digits are to be shifted, changed, dropped, or inserted in any number represented decimally, then this rule can be represented equally well by an arithmetical counterpart which involves arithmetical operations with powers of 10 as well as additions, subtractions, and so forth.

More briefly:

Typographical rules for manipulating *numerals* are actually arithmetical rules for operating on *numbers.*

This simple observation is at the heart of Gödel's method, and it will have an absolutely shattering effect. It tells us that once we have a Gödel-numbering for any formal system, we can straightaway form a set of arithmetical rules which complete the Gödel isomorphism. The upshot is that we can transfer the study of any formal system—in fact the study of *all* formal systems—into number theory.

MIU-Producible Numbers

Just as any set of typographical rules generates a set of theorems, a corresponding set of natural numbers will be generated by repeated applications of arithmetical rules. These *producible numbers* play the same role inside number theory as theorems do inside any formal system. Of course, different numbers will be producible, depending on which rules are adopted. "Producible numbers" are only producible *relative to a system* of arithmetical rules. For example, such numbers as 31, 3010010, 3111, and so forth could be called *MIU-producible* numbers—an ungainly name, which might be shortened to *MIU-numbers,* symbolizing the fact that those numbers are the ones that result when you transcribe the MIU-system into number theory, via Gödel-numbering. If we were to Gödel-number the pq-system

and then "arithmetize" its rules, we could call the producible numbers "pq-numbers"—and so on.

Note that the producible numbers (in any given system) are defined by a recursive method: given numbers which are known to be producible, we have rules telling how to make more producible numbers. Thus, the class of numbers known to be producible is constantly extending itself, in much the same way that the list of Fibonacci numbers, or Q-numbers, does. The set of producible numbers of any system is a *recursively enumerable set.* What about its complement—the set of nonproducible numbers? Is that set always recursively enumerable? Do numbers which are nonproducible share some common arithmetical feature?

This is the sort of issue which arises when you transpose the study of formal systems into number theory. For each system which is arithmetized, one can ask, "Can we characterize producible numbers in a simple way?" "Can we characterize *non*producible numbers in a recursively enumerable way?" These are difficult questions of number theory. Depending on the system which has been arithmetized, such questions might prove too hard for us to resolve. But if there is any hope for solving such problems, it would have to reside in the usual kind of step-by-step reasoning as it applies to natural numbers. And that, of course, was put in its quintessential form in the previous Chapter. TNT seemed, to all appearances, to have captured all valid mathematical thinking processes in one single, compact system.

Answering Questions about Producible Numbers by Consulting TNT

Could it be, therefore, that the means with which to answer any question about any formal system lies within just a single formal system—TNT? It seems plausible. Take, for instance, this question:

Is MU a theorem of the MIU-system?

Finding the answer is equivalent to determining whether 30 is a MIU-number or not. Because it is a statement of number theory, we should expect that, with some hard work, we could figure out how to translate the sentence "30 is a MIU-number" into TNT-notation, in somewhat the same way as we figured out how to translate other number-theoretical sentences into TNT-notation. I should immediately caution the reader that such a translation, though it does exist, is immensely complex. If you recall, I pointed out in Chapter VIII that even such a simple arithmetical predicate as "b is a power of 10" is very tricky to code into TNT-notation—and the predicate "b is a MIU-number" is a lot more complicated than that! Still, it can be found; and the numeral SSSSSSSSSSSSSSSSSSSSSSSSSSSSSS0 can be substituted for every b. This will result in a MONstrous string of TNT, a string of TNT which speaks about the MU-puzzle. Let us therefore call that string "MUMON". Through MUMON and strings like it, TNT is now capable of speaking "in code" about the MIU-system.

The Dual Nature of MUMON

In order to gain some benefit from this peculiar transformation of the original question, we would have to seek the answer to a new question:

Is MUMON a theorem of TNT?

All we have done is replace one relatively short string (**MU**) by another (the monstrous MUMON), and a simple formal system (the MIU-system) by a complicated one (TNT). It isn't likely that the answer will be any more forthcoming even though the question has been reshaped. In fact, TNT has a full complement of both lengthening and shortening rules, and the reformulation of the question is likely to be far harder than the original. One might even say that looking at **MU** via MUMON is an intentionally idiotic way of doing things. However, MUMON can be looked at on more than one level.

In fact, this is an intriguing point: MUMON has two different passive meanings. Firstly, it has the one which was given before:

30 is a MIU-number.

But secondly, we know that this statement is tied (via isomorphism) to the statement

MU is a theorem of the MIU-system.

So we can legitimately quote this latter as the second passive meaning of MUMON. It may seem very strange because, after all, MUMON contains nothing but plus signs, parentheses, and so forth—symbols of TNT. How can it possibly express any statement with other than arithmetical content?

The fact is, it can. Just as a single musical line may serve as both harmony and melody in a single piece; just as "BACH" may be interpreted as both a name and a melody; just as a single sentence may be an accurate structural description of a picture by Escher, of a section of DNA, of a piece by Bach, and of the dialogue in which the sentence is embedded, so MUMON can be taken in (at least) two entirely different ways. This state of affairs comes about because of two facts:

Fact 1. Statements such as "**MU** is a theorem" can be coded into number theory via Gödel's isomorphism.

Fact 2. Statements of number theory can be translated into TNT.

It could be said that MUMON is, by Fact 1, a coded message, where the symbols of the code are, by Fact 2, just symbols of TNT.

Codes and Implicit Meaning

Now it could be objected here that a coded message, unlike an uncoded message, does not express anything on its own—it requires knowledge of the code. But in reality there is no such thing as an uncoded message. There are only messages written in more familiar codes, and messages written in less familiar codes. If the meaning of a message is to be revealed, it must be pulled out of the code by some sort of mechanism, or isomorphism. It may be difficult to discover the method by which the decoding should be done; but once that method has been discovered, the message becomes transparent as water. When a code is familiar enough, it ceases appearing like a code; one forgets that there is a decoding mechanism. The message is identified with its meaning.

Here we have a case where the identification of message and meaning is so strong that it is hard for us to conceive of an alternate meaning residing in the same symbols. Namely, we are so prejudiced by the symbols of TNT towards seeing number-theoretical meaning (and *only* number-theoretical meaning) in strings of TNT, that to conceive of certain strings of TNT as statements about the MIU-system is quite difficult. But Gödel's isomorphism compels us to recognize this second level of meaning in certain strings of TNT.

Decoded in the more familiar way, MUMON bears the message:

30 is a MIU-number.

This is a statement of number theory, gotten by interpreting each sign in the conventional way.

But in discovering Gödel-numbering and the whole isomorphism built upon it, we have in a sense broken a code in which messages about the MIU-system are written in strings of TNT. Gödel's isomorphism is a new information-revealer, just as the decipherments of ancient scripts were information-revealers. Decoded by this new and less familiar mechanism, MUMON bears the message

MU is a theorem of the MIU-system.

The moral of the story is one we have heard before: that meaning is an automatic by-product of our recognition of any isomorphism; therefore there are at least two passive meanings of MUMON—maybe more!

The Boomerang: Gödel-Numbering TNT

Of course things do not stop here. We have only begun realizing the potential of Gödel's isomorphism. The natural trick would be to turn TNT's capability of mirroring other formal systems back on itself, as the Tortoise turned the Crab's phonographs against themselves, and as his Goblet G turned against itself, in destroying itself. In order to do this, we