DOUGLAS R. HOFSTADTER

Prelude... Ant Fugue

Prelude . . .

Achilles and the Tortoise have come to the residence of their friend the Crab, to make the acquaintance of one of his friends, the Anteater. The introductions having been made, the four of them settle down to tea.

tortoise: We have brought along a little something for you, Mr. Crab.

crab: That's most kind of you. But you shouldn't have.

tortoise: Just a token of our esteem. Achilles, would you like to give it to Mr. C?

achilles: Surely. Best wishes, Mr. Crab. I hope you enjoy it.

(Achilles hands the Crab an elegantly wrapped present, square and very thin. The Crab begins unwrapping it.)

anteater: I wonder what it could be.

crab: We'll soon find out. ***(Completes the unwrapping, and pulls out the gift.)*** Two records! How exciting! But there's no label. Uh-oh—is this another of your "specials," Mr. T?

tortoise: If you mean a phonograph-breaker, not this time. But it is in fact a custom-recorded item, the only one of its kind in the entire

world. In fact, it's never even been heard before—except, of course, when Bach played it.

crab: When Bach played it? What do you mean, exactly?

achilles: Oh, you are going to be fabulously excited, Mr. Crab, when Mr. T tells you what these records in fact are.

tortoise: Oh, you go ahead and tell him, Achilles.

achilles: May I? Oh, boy! I'd better consult my notes, then. ***(Pulls out a small filing card and clears his voice.)*** Ahem. Would you be interested in hearing about the remarkable new result in mathematics, to which your records owe their existence?

crab: My records derive from some piece of mathematics? How curi ous! Well, now that you've provoked my interest, I must hear about it.

achilles: Very well, then. ***(Pauses for a moment to sip his tea, then resumes.)*** Have you heard of Fermat's infamous "Last Theorem"?

anteater: I'm not sure. ... It sounds strangely familiar, and yet I can't quite place it.

achilles: It's a very simple idea. Pierre de Fermat, a lawyer by vocation but mathematician by avocation, had been reading in his copy of the classic text ***Arithmetica*** by Diophantus and came across a page containing the equation

$$a^2 + b^2 = c^2$$

He immediately realized that this equation has infinitely many solutions ***a, b, c,*** and then wrote in the margin the following notorious comment:

> The equation
>
> $$a^* + b^n = c^n$$
>
> has solutions in positive integers ***a, b, c,*** and n only when ***n*** = 2 (and then there are infinitely many triplets ***a, b, c*** which satisfy the equation); but there are no solutions for ***n*** > 2. I have discovered a truly marvelous proof of this statement, which, unfortunately, this margin is too small to contain.

Ever since that day, some three hundred years ago, mathematicians have been vainly trying to do one of two things: either to prove Fermat's claim and thereby vindicate Fermat's reputation, which, although very high, has been somewhat tarnished by skeptics who

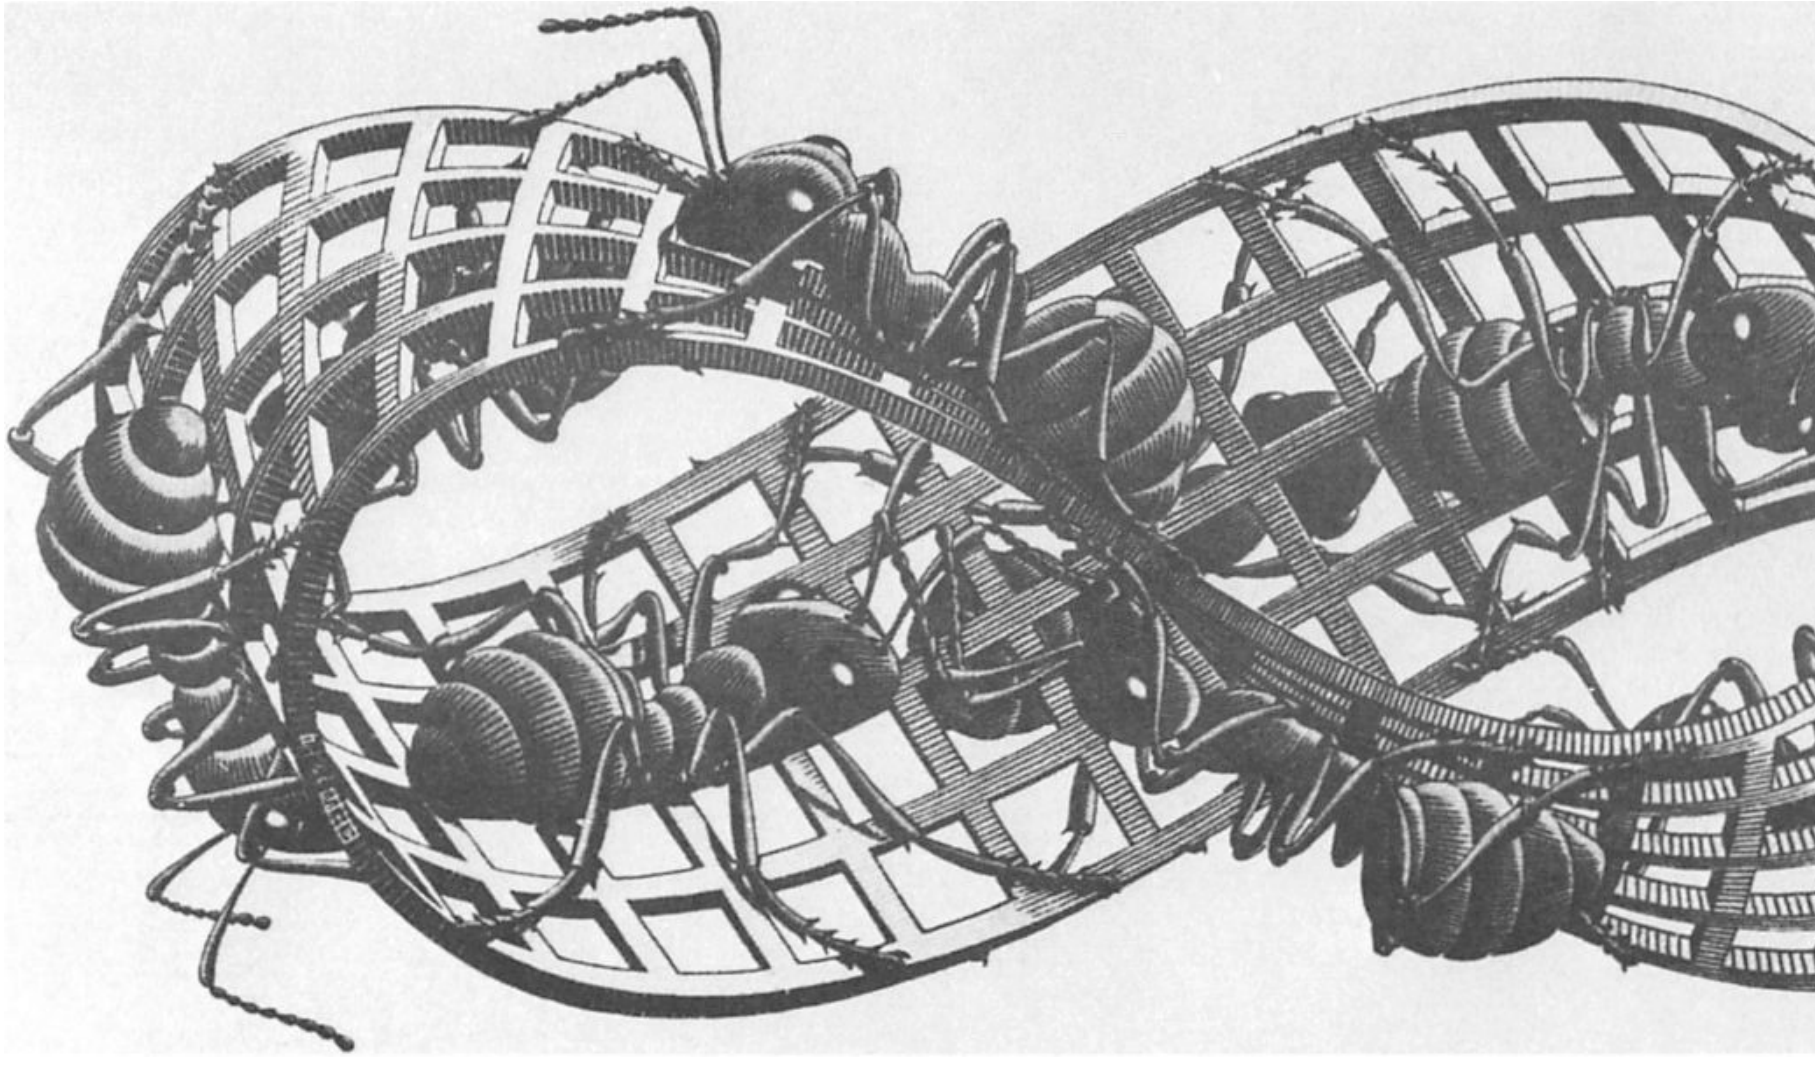

Mobi us Strip // (M. C. Escher, woodcut, 1963).

think he never really found the proof he claimed to have found—or else to refute the claim, by finding a counterexample: a set of four integers ***a, b, c,*** and ***n,*** with ***n*** > 2, which satisfy the equation. Until very recently, every attempt in either direction had met with failure. To be sure, the Theorem has been proven for many specific values of ***n***—in particular, all ***n*** up to 125,000.

anteater: Shouldn't it be called a "Conjecture" rather than a "Theorem," if it's never been given a proper proof?

achilles: Strictly speaking, you're right, but tradition has kept it this way.

crab: Has someone at last managed to resolve this celebrated question?

achilles: Indeed! In fact, Mr. Tortoise has done so, and as usual, by a wizardly stroke. He has not only found a ***proof*** of Fermat's Last Theorem (thus justifying its name as well as vindicating Fermat), but also a ***counterexample,*** thus showing that the skeptics had good intuition!

crab: Oh my gracious! That is a revolutionary discovery.

anteater: But please don't leave us in suspense. What magical integers are they, that satisfy Fermat's equation? I'm especially curious about the value of n.

achilles: Oh, horrors! I'm most embarrassed! Can you believe this? I left the values at home on a truly colossal piece of paper. Unfortunately it was too huge to bring along. I wish I had them here to show to you. If it's of any help to you, I do remember one thing—the value of n is the only positive integer which does not occur anywhere in the continued fraction for ***n.***

crab: Oh, what a shame that you don't have them here. But there's no reason to doubt what you have told us.

anteater: Anyway, who needs to see ***n*** written out decimally? Achilles has just told us how to find it. Well, Mr. T, please accept my hearty felicitations, on the occasion of your epoch-making discovery!

tortoise: Thank you. But what I feel is more important than the result itself is the practical use to which my result immediately led.

crab: I am dying to hear about it, since I always thought number theory was the Queen of Mathematics—the purest branch of mathematics —the one branch of mathematics which has ***no*** applications!

tortoise: You're not the only one with that belief, but in fact it is quite impossible to make a blanket statement about when or how some

Pierre de Fermat.

branch—or even some individual Theorem—of pure mathematics will have important repercussions outside of mathematics. It is quite unpredictable—and this case is a perfect example of that phenomenon.

achilles: Mr. Tortoise's double-barreled result has created a break through in the field of acoustico-retrieval!

anteater: What is acoustico-retrieval?

achilles: The name tells it all: it is the retrieval of acoustic information from extremely complex sources. A typical task of acoustico-retrieval is to reconstruct the sound which a rock made on plummeting into a lake, from the ripples which spread out over the lake's surface.

crab: Why, that sounds next to impossible!

achilles: Not so. It is actually quite similar to what one's brain does, when it reconstructs the sound made in the vocal cords of another person from the vibrations transmitted by the eardrum to the fibers in the cochlea.

crab: I see. But I still don't see where number theory enters the picture, or what this all has to do with my new records.

achilles: Well, in the mathematics of acoustico-retrieval, there arise many questions which have to do with the number of solutions of

certain Diophantine equations. Now Mr. T has been for years trying to find a way of reconstructing the sounds of Bach playing his harpsichord, which took place over two hundred years ago, from calculations involving the motions of all the molecules in the atmosphere at the present time.

anteater: Surely that is impossible! They are irretrievably gone, gone forever!

achilles: Thus think the naive .. . But Mr. T has devoted many years to this problem, and came to the realization that the whole thing hinged on the number of solutions to the equation

$$a^n + b^n = c^n$$

in positive integers, with $n > 2$.

tortoise: I **could explain, of course, just how this equation arises, but** I'm **sure it would bore you.**

achilles: It turned out that acoustico-retrieval theory predicts that the Bach sounds can be retrieved from the motion of all the molecules in the atmosphere, provided that there exists ***either*** at least one solution to the equation—

crab: Amazing!

anteater: Fantastic!

tortoise: Who would have thought!

achilles: I was about to say, "provided that there exists ***either*** such a solution ***or*** a proof that there are ***no*** solutions!" And therefore, Mr. T, in careful fashion, set about working at both ends of the problem simultaneously. As it turns out, the discovery of the counterexample was the key ingredient to finding the proof, so the one led directly to the other.

CRAB: HOW could that be?

tortoise: Well, you see, I had shown that the structural layout of any proof of Fermat's Last Theorem—if one existed—could be described by an elegant formula, which, it so happened, depended on the values of a solution to a certain equation. When I found this second equation, to my surprise it turned out to be the Fermat equation. An amusing accidental relationship between form and content. So when I found the counterexample, all I needed to do was to use those numbers as a blueprint for constructing my proof that there were no solutions to the equation. Remarkably simple, when you think

about it. I can't imagine why no one had ever found the result before.

achilles: As a result of this unanticipatedly rich mathematical success, Mr. T was able to carry out the acoustico-retrieval which he had so long dreamed of. And Mr. Crab's present here represents a palpable realization of all this abstract work.

crab: Don't tell me it's a recording of Bach playing his own works for harpsichord!

achilles: I'm sorry, but I have to, for that is indeed just what it is! This is a set of two records of Johann Sebastian Bach playing all of his ***Well-Tempered Clavier***. Each record contains one of the two volumes of the ***Well-Tempered Clavier;*** that is to say, each record contains twenty-four preludes and fugues—one in each major and minor key.

crab: Well, we must absolutely put one of these priceless records on, immediately! And how can I ever thank the two of you?

TORTOISE: YOU have already thanked us plentifully, with this delicious tea which you have prepared.

(The Crab slides one of the records out of its jacket and puts it on. The sound of an incredibly masterful harpsichordist fills the room, in the highest imaginable fidelity. One even hears—or is it one's imagination?—the soft sounds of Bach singing to himself as he plays. . . .)

crab: Would any of you like to follow along in the score? I happen to have a unique edition of the ***Well-Tempered Clavier,*** specially illuminated by a teacher of mine who happens also to be an unusually fine calligrapher.

tortoise: I would very much enjoy that.

(The Crab goes to his elegant glass-enclosed wooden bookcase, opens the doors, and draws out two large volumes.)

crab: Here you are, Mr. Tortoise. I've never really gotten to know all the beautiful illustrations in this edition. Perhaps your gift will provide the needed impetus for me to do so.

tortoise: I do hope so.

anteater: Have you ever noticed how in these pieces the prelude always sets the mood perfectly for the following fugue?

crab: Yes. Although it may be hard to put it into words, there is always some subtle relation between the two. Even if the prelude and fugue

do not have a common melodic subject, there is nevertheless always some intangible abstract quality which underlies both of them, binding them together very strongly.

tortoise: And there is something very dramatic about the few moments of silent suspense hanging between prelude and fugue—that moment where the theme of the fugue is about to ring out, in single tones, and then to join with itself in ever-increasingly complex levels of weird, exquisite harmony.

achilles: I know just what you mean. There are so many preludes and fugues which I haven't yet gotten to know, and for me that fleeting interlude of silence is very exciting; it's a time when I try to second- guess old Bach. For example, I always wonder what the fugue's tempo will be: allegro or adagio? Will it be in 6/8 or 4/4? Will it have
three voices or five—or four? And then, the first voice starts Such
an exquisite moment.

crab: Ah, yes, well do I remember those long-gone days of my youth, the days when I thrilled to each new prelude and fugue, filled with the excitement of their novelty and beauty and the many unexpected surprises which they conceal.

achilles: And now? Is that thrill all gone?

crab: It's been supplanted by familiarity, as thrills always will be. But in that familiarity there is also a kind of depth, which has its own compensations. For instance, I find that there are always new surprises which I hadn't noticed before.

achilles: Occurrences of the theme which you had overlooked?

crab: Perhaps—especially when it is inverted and hidden among several other voices, or where it seems to come rushing up from the depths, out of nowhere. But there are also amazing modulations which it is marvelous to listen to over and over again, and wonder how old Bach dreamt them up.

achilles: I am very glad to hear that there is something to look forward to, after I have been through the first flush of infatuation with the *Well-Tempered Clavier*—although it also makes me sad that this stage could not last forever and ever.

crab: Oh, you needn't fear that your infatuation will totally die. One of the nice things about that sort of youthful thrill is that it can always be resuscitated, just when you thought it was finally dead. It just takes the right kind of triggering from the outside.

achilles: Oh, really? Such as what?

crab: Such as hearing it through the ears, so to speak, of someone to whom it is a totally new experience—someone such ;as you, Achilles. Somehow the excitement transmits itself, and I can feel thrilled again.

achilles: That is intriguing. The thrill has remained dormant somewhere inside you, but by yourself, you aren't able to fish it up out of your subconscious.

crab: Exactly. The potential of reliving the thrill is "coded," in some unknown way, in the structure of my brain, but I don't have the power to summon it up at will; I have to wait for chance circumstance to trigger it.

achilles: I **have a question about fugues which** I **feel a little embarrassed about asking, but as I am just a novice at fugue-listening, I was wondering if perhaps one of you seasoned fugue-listeners might help me in learning? . . .**

tortoise: I'd certainly like to offer my own meager knowledge, if it might prove of some assistance.

achilles: Oh, thank you. Let me come at the question from an angle. Are you familiar with the print called ***Cube with Magic Ribbons,*** by M. C. Escher?

tortoise: In which there are circular bands having bubblelike distortions which, as soon as you've decided that they are bumps, seem to turn into dents—and vice versa?

achilles: Exactly.

crab: I remember that picture. Those little bubbles always seem to flip back and forth between being concave and convex, depending on the direction that you approach them from. There's no way to see them simultaneously as concave ***and*** convex—somehow one's brain doesn't allow that. There are two mutually exclusive "modes" in which one can perceive the bubbles.

achilles: Just so. Well, I seem to have discovered two somewhat analogous modes in which I can listen to a fugue. The modes are these: either to follow one individual voice at a time, or to listen to the total effect of all of them together, without trying to disentangle one from another. I have tried out both of these modes, and, much to my frustration, each one of them shuts out the other. It's simply not in my power to follow the paths of individual voices and at the same

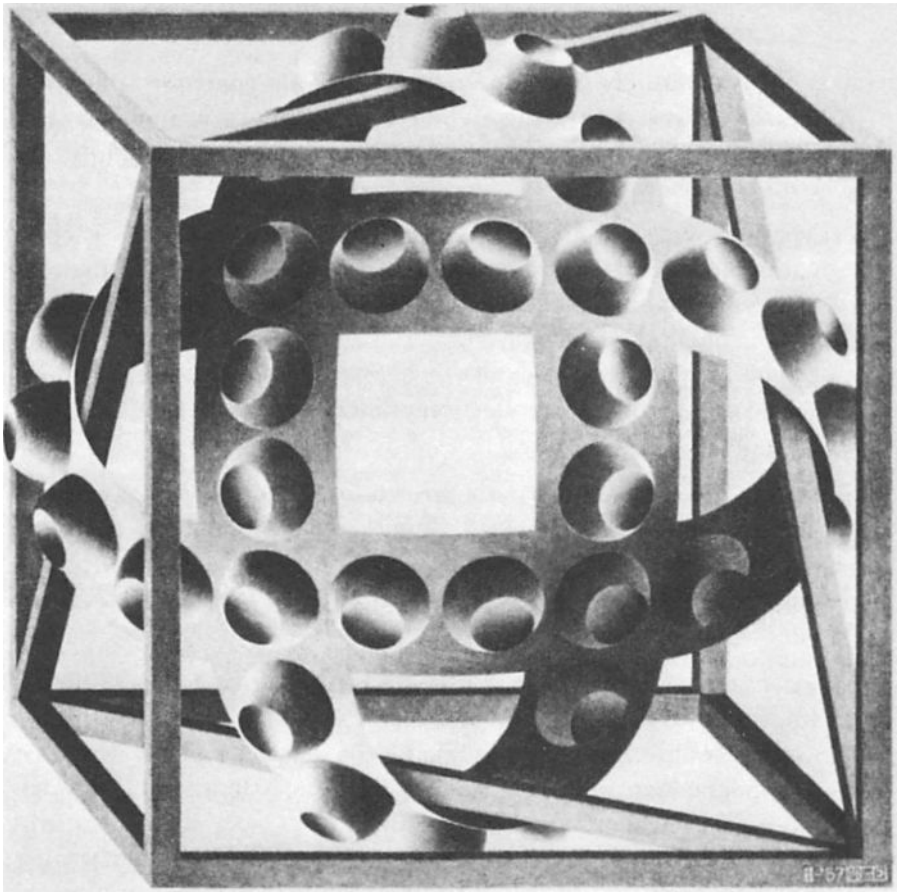

Cube with Magic Ribbons (M. C. Eschcr, lithograph, 1957).

time to hear the whole effect. I find that I flip back and forth between one mode and the other, more or less spontaneously and involuntarily.

anteater: Just as when you look at the magic bands, eh?

achilles: Yes. I was just wondering . . . does my description of these two modes of fugue-listening brand me unmistakably as a naive, inexperienced listener, who couldn't even begin to grasp the deeper modes of perception which exist beyond his ken?

tortoise: No, not at all, Achilles. I can only speak for myself, but I too find myself shifting back and forth from one mode to the other without exerting any conscious control over which mode should be

dominant. I don't know if our other companions here have also experienced anything similar.

crab: Most definitely. It's quite a tantalizing phenomenon, since you feel that the essence of the fugue is Bitting about you, and you can't quite grasp all of it, because you can't quite make yourself function both ways at once.

anteater: Fugues have that interesting property, that each of their voices is a piece of music in itself; and thus a fugue might be thought of as a collection of several distinct pieces of music, all based on one single theme, and all played simultaneously. And it is up to the listener (or his subconscious) to decide whether it should be perceived as a unit, or as a collection of independent parts, all of which harmonize.

achilles: You say that the parts are "independent," yet that can't be literally true. There has to be some coordination between them, otherwise when they were put together one would just have an unsystematic clashing of tones—and that is as far from the truth as could be.

anteater: A better way to state it might be this: if you listened to each voice on its own, you would find that it seemed to make sense all by itself. It could stand alone, and that is the sense in which I meant that it is independent. But you are quite right in pointing out that each of these individually meaningful lines fuses with the others in a highly nonrandom way, to make a graceful totality. The art of writing a beautiful fugue lies precisely in this ability, to manufacture several different lines, each one of which gives the illusion of having been written for its own beauty, and yet which when taken together form a whole, which does not feel forced in any way. Now, this dichotomy between hearing a fugue as a whole and hearing its component voices is a particular example of a very general dichotomy, which applies to many kinds of structures built up from lower levels.

achilles: Oh, really? You mean that my two "modes" may have some more general type of applicability, in situations other than fugue- listening?

anteater: Absolutely.

achilles: I wonder how that could be. I guess it has to do with alternating between perceiving something as a whole and perceiving it as a collection of parts. But the only place I have ever run into that dichotomy is in listening to fugues.

tortoise: Oh, my, look at this! I just turned the page while following the music, and came across this magnificent illustration facing the first page of the fugue.

crab: I have never seen that illustration before. Why don't you pass it 'round?

(The Tortoise passes the book around. Each of the foursome looks at it in a characteristic way—this one from afar, that one from close up, everyone tipping his head this way and that in puzzlement. Finally it has made the rounds and returns to the Tortoise, who peers at it rather intently.)

achilles: Well, I guess the prelude is just about over. I wonder if, as I listen to this fugue, I will gain any more insight into the question "What is the right way to listen to a fugue: as a whole, or as the sum of its parts?"

ttortoise: Listen carefully, and you will!

(The prelude ends. There is a moment of silence; and. . .

[ATTACCA]

. . . Ant Fugue

. . . then, one by one, the four voices of the fugue chime in.)

achilles: I know the rest of you won't believe this, but the answer to the question is staring us all in the face, hidden in the picture. It is simply one word—but what an important one: **"mu"!**

ccrab: I know the rest of you won't believe this, but the answer to the question is staring us all in the face, hidden in the picture. It is simply one word—but what an important one: **"holism"!**

achilles: Now hold on a minute. You must be seeing things. It's plain as day that the message of this picture is **"mu,"** not **"holism"!**

crab: I beg your pardon, but my eyesight is extremely good. Please look again, and then tell me if the picture doesn't say what I said it says!

anteater: I know the rest of you won't believe this, but the answer to the question is staring us all in the face, hidden in the picture. It is simply one word—but what an important one: **"reductionism"!**

Illustration by the Author.

crab: **Now** hold on a minute. You must be seeing things. It's plain as day that the message of this picture is **"holism,"** not **"reduction- ism"!**

achilles: Another deluded one! Not **"holism,"** not **"reductionism,"** but **"mu"** is the message of this picture, and that much is certain.

anteater: I beg your pardon, but my eyesight is extremely clear. Please look again, and then see if the picture doesn't say what I said it says.

achilles: Don't you see that the picture is composed of two pieces, and that each of them is a single letter?

crab: You are right about the two pieces, but you are wrong in your identification of what they are. The piece on the left is entirely composed of three copies of one word: **'"holism";** and the piece on the right is composed of many copies, in smaller letters, of the same word. Why the letters are of different sizes in the two parts, I don't know, but I know what I see, and what I see is **"holism,"** plain as day. How you see anything else is beyond me.

anteater: **You** are right about the two pieces, but you are wrong in your identification of what they are. The piece on the left is entirely composed of many copies of one word: **"reductionism";** and the piece on the right is composed of one single copy, in larger letters, of the same word. Why the letters are of different sizes in the two parts, I don't know, but I know what I see, and what I see is **"reductionism,"** plain as day. How you see anything else is beyond me.

achilles: I know what is going on here. Each of you has seen letters which compose, or are composed of, other letters. In the left-hand piece, there are indeed three **"holism"s,** but each one of them is composed out of smaller copies of the word **"reductionism."** And in complementary fashion, in the right-hand piece, there is indeed one **"reductionism,"** but it is composed out of smaller copies of the word **"holism." Now** this is all fine and good, but in your silly squabble, the two of you have actually missed the forest for the trees. You see, what good is it to argue about whether **"holism"** or **"reductionism"** is right, when the proper way to understand the matter is to transcend the question, by answering **"mu"?**

crab: I now see the picture as you have described it, Achilles, but I have no idea of what you mean by the strange expression "transcending the question."

anteater: I now see the picture as you have described it, Achilles, but I have no idea of what you mean by the strange expression ***"mu."*** "

achilles: I will be glad to indulge both of you, if you will first oblige me, by telling me the meaning of these strange expressions, "***holism***" and "***reductionism***."

crab: ***Holism*** is the most natural thing in the world to grasp. It's simply the belief that "the whole is greater than the sum of its parts." No one in his right mind could reject holism.

anteater: ***Reductionism*** is the most natural thing in the world to grasp. It's simply the belief that "a whole can be understood completely if you understand its parts, and the nature of their 'sum.' " No one in her left brain could reject reductionism.

crab: I reject reductionism. I challenge you to tell me, for instance, how to understand a brain reductionistically. Any reductionistic explanation of a brain will inevitably fall far short of explaining where the consciousness experienced by a brain arises from.

anteater: I reject holism. I challenge you to tell me, for instance, how a holistic description of an ant colony sheds any more light on it than is shed by a description of the ants inside it, and their roles, and their interrelationships. Any holistic explanation of an ant colony will inevitably fall far short of explaining where the consciousness experienced by an ant colony arises from.

achilles: Oh, no! The last thing that I wanted to do was to provoke another argument. Anyway, now that I understand the controversy, I believe that my explanation of "**tom**" will help greatly. You see, "***mu***" is an ancient Zen answer which, when given to a question, ***unasks*** the question. Here, the question seems to be "Should the world be understood via holism or via reductionism?" And the answer of ***"mu"*** here rejects the premises of the question, which are that one or the other must be chosen. By unasking the question, it reveals a wider truth: that there is a larger context into which both holistic and reductionistic explanations fit.

anteater: Absurd! Your **"tom"** is as silly as a cow's moo. I'll have none of this Zen wishy-washiness.

crab: Ridiculous! Your ***"mu"*** is as silly as a kitten's mew. I'll have none of this Zen washy-wishiness.

achilles: Oh, dear! We're getting nowhere fast. Why have you stayed • so strangely silent, Mr. Tortoise? It makes me very uneasy. Surely you must somehow be capable of helping straighten out this mess?

tortoise: I know the rest of you won't believe this, but the answer to the question is staring us all in the face, hidden in the picture. It is simply one word—but what an important one: **"mu"!**

Prelude . . . Ant Fugue

(Just as he says this, the fourth voice in the fugue being played enters, exactly one octave below the first entry.)

achilles: Oh, Mr. T, for once you have let me down. I was sure that you, who always see the most deeply into things, would be able to resolve this dilemma—but apparently, you have seen no further than I myself saw. Oh, well, I guess I should feel pleased to have seen as far as Mr. Tortoise, for once.

tortoise: I beg your pardon, but my eyesight is extremely fine. Please look again, and then tell me if the picture doesn't say what I said it says.

achilles: But of course it does! You have merely repeated my own original observation.

tortoise: Perhaps **"mu"** exists in this picture on a deeper level than you imagine, Achilles—an octave lower (figuratively speaking). But for now I doubt that we can settle the dispute in the abstract. I would like to see both the holistic and reductionistic points of view laid out more explicitly; then there may be more of a basis for a decision. I would very much like to hear a reductionistic description of an ant colony, for instance.

crab: Perhaps Dr. Anteater will tell you something of his experiences in that regard. After all, he is by profession something of an expert on that subject.

tortoise: I am sure that we could learn much from a myrmecologist like you, Dr. Anteater. Could you tell us more about ant colonies, from a reductionistic point of view?

anteater: Gladly. As Mr. Crab mentioned to you, my profession has led me quite a long way into the understanding of ant colonies.

achilles: I can imagine! The profession of Anteater would seem to be synonymous with being an expert on ant colonies!

anteater: I beg your pardon. "Anteater" is not my profession; it is my species. By profession, I am a colony surgeon. I specialize in correcting nervous disorders of the colony by the technique of surgical removal.

achilles: Oh, I see. But what do you mean by "nervous disorders" of an ant colony?

anteater: Most of my clients suffer from some sort of speech impairment. You know, colonies which have to grope for words in everyday situations. It can be quite tragic. I attempt to remedy the situation

by, uhh—removing—the defective part of the colony. These operations are sometimes quite involved, and of course years of study are required before one can perform them.

achilles: But—isn't it true that, before one can suffer from speech impairment, one must have the faculty of speech?

anteater: Right.

achilles: Since ant colonies don't have that faculty, I am a little con
fused.

crab: It's too bad, Achilles, that you weren't here last week, when Dr. Anteater and Aunt Hillary were my house guests. I should have thought of having you over then.

achilles: Is Aunt Hillary your aunt, Mr. Crab?

crab: Oh, no, she's not really anybody's aunt.

anteater: But the poor dear insists that everybody should call her that, even strangers. It's just one of her many endearing quirks.

crab: Yes, Aunt Hillary is quite eccentric, but such a merry old soul. It's a shame I didn't have you over to meet her last week.

anteater: She's certainly one of the best-educated ant colonies I have ever had the good fortune to know. The two of us have spent many a long evening in conversation on the widest range of topics.

achilles: I thought anteaters were devourers of ants, not patrons of ant-intellectualism!

anteater: Well, of course the two are not mutually inconsistent. I am on the best of terms with ant colonies. It's just ***ants*** that I eat, not colonies—and that is good for both parties: me, and the colony.

achilles: How is it possible that—

tortoise: How is it possible that—

ACHILLES: —having its ants eaten can do an ant colony any good?

crab: How is it possible that—

TORTOISE: —having a forest fire can do a forest any good?

anteater: How is it possible that—

crab: —having its branches pruned can do a tree any good?

anteater: —having a haircut can do Achilles any good?

tortoise: Probably the rest of you were too engrossed in the discussion to notice the lovely stretto which just occurred in this Bach fugue.

achilles: What is **a** stretto?

tortoise: Oh, I'm sorry; I thought you knew the term. It is where one theme repeatedly enters in one voice after another, with very little delay between entries.

achilles: If I listen to enough fugues, soon I'll know all of these things and will be able to pick them out myself, without their having to be pointed out.

tortoise: Pardon me, my friends. I am sorry to have interrupted. Dr. Anteater was trying to explain how eating ants is perfectly consistent with being a friend of an ant colony.

achilles: Well, I can vaguely see how it might be possible for a limited and regulated amount of ant consumption to improve the overall health of a colony—but what is far more perplexing is all this talk about having conversations with ant colonies. That's impossible. An ant colony is simply a bunch of individual ants running around at random looking for food and making a nest.

anteater: **You** could put it that way if you want to insist on seeing the trees but missing the forest, Achilles. In fact, ant colonies, seen as wholes, are quite well-defined units, with their own qualities, at times including the mastery of language.

achilles: I find it hard to imagine myself shouting something out loud in the middle of the forest, and hearing an ant colony answer back.

anteater: Silly fellow! That's not the way it happens. Ant colonies don't converse out loud, but in writing. You know how ants form trails leading them hither and thither?

achilles: Oh, yes—usually straight through the kitchen sink and into my peach jam.

anteater: Actually, some trails contain information in coded form. If you know the system, you can read what they're saying just like a book.

achilles: Remarkable. And can you communicate back to them?

anteater: Without any trouble at all. That's how Aunt Hillary and I have conversations for hours. I take a stick and draw trails in the moist ground, and watch the ants follow my trails. Presently, a new trail starts getting formed somewhere. I greatly enjoy watching trails develop. As they are forming, I anticipate how they will continue (and more often I am wrong than right). When the trail is completed, I know what Aunt Hillary is thinking, and I in turn make my reply.

achilles: There must be some amazingly smart ants in that colony, I'll say that.

anteater: I think you are still having some difficulty realizing the difference in levels here. Just as you would never confuse an individual tree with a forest, so here you must not take an ant for the colony. You see, all the ants in Aunt Hillary are as dumb as can be. They couldn't converse to save their little thoraxes!

achilles: Well then, where does the ability to converse come from? It must reside somewhere inside the colony! I don't understand how the ants can all be unintelligent, if Aunt Hillary can entertain you for hours with witty banter.

tortoise: It seems to me that the situation is not unlike the composition of a human brain out of neurons. Certainly no one would insist that individual brain cells have to be intelligent beings on their own, in order to explain the fact that a person can have an intelligent conversation.

achilles: Oh, no, clearly not. With brain cells, I see your point completely. Only ... ants are a horse of another color. I mean, ants just roam about at will, completely randomly, chancing now and then upon a morsel of food. . . . They are free to do what they want to do, and with that freedom, I don't see at all how their behavior, seen as a whole, can amount to anything coherent—especially something so coherent as the brain behavior necessary for conversing.

crab: It seems to me that the ants are free only within certain constraints. For example, they are free to wander, to brush against each other, to pick up small items, to work on trails, and so on. But they never step out of that small world, that, ant-system, which they are in. It would never occur to them, for they don't have the mentality to imagine anything of the kind. Thus the ants are very reliable components, in the sense that you can depend on them to perform certain kinds of tasks in certain ways.

achilles: But even so, within those limits they are still free, and they just act at random, running about incoherently without any regard for the thought mechanisms of a higher-level being which Dr. Anteater asserts they are merely components of.

anteater: Ah, but you fail to recognize one thing, Achilles—the regularity of statistics.

achilles: How is that?

anteater: For example, even though ants as individuals wander about in what seems a random way, there are nevertheless overall trends, involving large numbers of ants, which can emerge from that chaos.

achilles: Oh, I know what you mean. In fact, ant trails are a perfect example of such a phenomenon. There, you have really quite unpredictable motion on the part of any single ant—and yet, the trail itself seems to remain well defined and stable. Certainly that must mean that the individual ants are not just running about totally at random.

anteater: Exactly, Achilles. There is some degree of communication among the ants, just enough to keep them from wandering off completely at random. By this minimal communication they can remind each other that they are not alone but are cooperating with teammates. It takes a large number of ants, all reinforcing each other this way, to sustain any activity—such as trail building—for any length of time. Now my very hazy understanding of the operation of brains leads me to believe that something similar pertains to the firing of neurons. Isn't it true, Mr. Crab, that it takes a group of neurons firing in order to make another neuron fire?

crab: Definitely. Take the neurons in Achilles' brain, for example. Each neuron receives signals from neurons attached to its input lines, and if the sum total of inputs at any moment exceeds a critical threshold, then that neuron will fire and send its own output pulse rushing off to other neurons, which may in turn fire—and on down the line it goes. The neural flash swoops relentlessly in its Achillean path, in shapes stranger then the dash of a gnat-hungry swallow; every twist, every turn foreordained by the neural structure in Achilles' brain, until sensory input messages interfere.

achilles: Normally, I think that ***I'm*** in control of what I think—but the way you put it turns it all inside out, so that it sounds as though "I" am just what comes out of all this neural structure, and natural law. It makes what I consider my ***self*** sound at best like a by-product of an organism governed by natural law and, at worst, an artificial notion produced by my distorted perspective. In other words, you make me feel like I don't know who—or what—I am, if anything.

tortoise: You'll come to understand much better as we go along. But Dr. Anteater—what do you make of this similarity?

anteater: I knew there was something parallel going on in the two very different systems. Now I understand it much better. It seems that group phenomena which have coherence—trail building, for example—will take place only when a certain threshold number of ants get

involved. If an effort is initiated, perhaps at random, by a few ants in some locale, one of two things can happen: either it will fizzle out after a brief sputtering start—

achilles: When there aren't enough ants to keep the thing rolling?

anteater: Exactly. The other thing that can happen is that a critical mass of ants is present, and the thing will snowball, bringing more and more ants into the picture. In the latter case, a whole "team" is brought into being which works on a single project. That project might be trail making, or food gathering, or it might involve nest keeping. Despite the extreme simplicity of this scheme on a small scale, it can give rise to very complex consequences on a larger scale.

achilles: I can grasp the general idea of order emerging from chaos, as you sketch it, but that still is a long way from the ability to converse. After all, order also emerges from chaos when molecules of a gas bounce against each other randomly—yet all that results there is an amorphous mass with but three parameters to characterize it: volume, pressure, and temperature. Now that's a far cry from the ability to understand the world, or to talk about it!

anteater: That highlights a very interesting difference between the explanation of the behavior of an ant colony and the explanation of the behavior of gas inside a container. One can explain the behavior of the gas simply by calculating the statistical properties of the motions of its molecules. There is no need to discuss any higher elements of structure than molecules, except the full gas itself. On the other hand, in an ant colony, you can't even begin to understand the activities of the colony unless you go through several layers of structure.

achilles: I see what you mean. In a gas, one jump takes you from the lowest level—molecules—to the highest level—the full gas. There are no intermediate levels of organization. Now how do intermediate levels of organized activity arise in an ant colony?

anteater: It has to do with the existence of several different varieties of ants inside any colony.

achilles: Oh, yes. I think I have heard about that. They are called "castes," aren't they?

anteater: That's correct. Aside from the queen, there are males, who do practically nothing toward the upkeep of the nest, and then—

achilles: And of course there are soldiers—glorious fighters against communism!

crab: Hmm ... I hardly think that could be right, Achilles. An ant colony is quite communistic internally, so why would its soldiers fight against communism? Or am I right, Dr. Anteater?

anteater: Yes, about colonies you are right, Mr. Crab; they are indeed based on somewhat communistic principles. But about soldiers Achilles is somewhat naive. In fact, the so-called "soldiers" are hardly adept at fighting at all. They are slow, ungainly ants with giant heads, who can snap with their strong jaws, but are hardly to be glorified. As in a true communistic state, it is rather the workers who are to be glorified. It is they who do most of the chores, such as food gathering, hunting, and nursing of the young. It is even they who do most of the fighting.

achilles: Bah. That is an absurd state of affairs. Soldiers who won't fight!

anteater: Well, as I just said, they really aren't soldiers at all. It's the workers who are soldiers; the soldiers are just lazy fatheads.

achilles: Oh, how disgraceful! Why, if I were an ant, I'd put some discipline in their ranks! I'd knock some sense into those fatheads!

tortoise: If you were an ant? How could a myrmedian like you be an ant? There is no way to map your brain onto an ant brain, so it seems to me to be a pretty fruitless question to worry over. More reasonable would be the proposition of mapping your brain onto an ant colony. ... But let us not get sidetracked. Let Dr. Anteater continue with his most illuminating description of castes and their role in the higher levels of organization.

anteater: Very well. There are all sorts of tasks which must be accomplished in a colony, and individual ants develop specializations. Usually an ant's specialization changes as the ant ages. And of course it is also dependent on the ant's caste. At any one moment, in any small area of a colony, there are ants of all types present. Of course, one caste may be be very sparse in some places and very dense in others.

crab: IS the density of a given caste, or specialization, just a random thing? Or is there a reason why ants of one type might be more heavily concentrated in certain areas, and less heavily in others?

anteater: I'm glad you brought that up, since it is of crucial importance in understanding how a colony thinks. In fact, there evolves, over a long period of time, a very delicate distribution of castes inside a colony. And it is this distribution that allows the colony to have the complexity that underlies the ability to converse with me.

achilles: It would seem to me that the constant motion of ants to and fro would completely prevent the possibility of a very delicate distribution. Any delicate distribution would be quickly destroyed by all the random motions of ants, just as any delicate pattern among molecules in a gas would not survive for an instant, due to the random bombardment from all sides.

anteater: In an ant colony, the situation is quite the contrary. In fact, it is just exactly the constant to-ing and fro-ing of ants inside the colony which adapts the caste distribution to varying situations, and thereby preserves the delicate caste distribution. You see, the caste distribution cannot remain as one single rigid pattern; rather, it must constantly be changing so as to reflect, in some manner, the real- world situation with which the colony is dealing, and it is precisely the motion inside the colony which updates the caste distribution, so as to keep it in line with the present circumstances facing the colony.

tortoise: Could you give an example?

anteater: Gladly. When I, an anteater, arrive to pay a visit to Aunt Hillary, all the foolish ants, upon sniffing my odor, go into a panic—which means, of course, that they begin running around completely differently from the way they were before I arrived.

achilles: But that's understandable, since you're a dreaded enemy of the colony.

anteater: Oh, no. I must reiterate that, far from being an enemy of the colony, I am Aunt Hillary's favorite companion. And Aunt Hillary is my favorite aunt. I grant you, I'm quite feared by all the individual ants in the colony—but that's another matter entirely. In any case, you see that the ants' action in response to my arrival completely changes the internal distribution of ants.

achilles: That's clear.

anteater: And that sort of thing is the updating which I spoke of. The new distribution reflects my presence. One can describe the change from old state to new as having added a "piece of knowledge" to the colony.

ACHILLES: HOW can you refer to the distribution of different types of ants inside a colony as a "piece of knowledge"?

anteater: Now there's a vital point. It requires some elaboration. You see, what it comes down to is how you choose to describe the caste distribution. If you continue to think in terms of the lower levels—

individual ants—then you miss the forest for the trees. That's just too microscopic a level, and when you think microscopically, you're bound to miss some large-scale features. You've got to find the proper high-level framework in which to describe the caste distribution—only then will it make sense how the caste distribution can encode many pieces of knowledge.

achilles: Well, how ***do*** you find the proper-sized units in which to describe the present state of the colony, then?

anteater: All right. Let's begin at the bottom. When ants need to get something done, they form little "teams," which stick together to perform a chore. As I mentioned earlier, small groups of ants are constantly forming and unforming. Those which actually exist for a while are the teams, and the reason they don't fall apart is that there really is something for them to do.

achilles: Earlier you said that a group will stick together if its size exceeds a certain threshold. Now you're saying that a group will stick together if there is something for it to do.

anteater: They are equivalent statements. For instance, in food gathering, if there is an inconsequential amount of food somewhere which gets discovered by some wandering ant who then attempts to communicate its enthusiasm to other ants, the number of ants who respond will be proportional to the size of the food sample—and an inconsequential amount will not attract enough ants to surpass the threshold. Which is exactly what I meant by saying there is nothing to do—too little food ought to be ignored.

achilles: I see. I assume that these "teams" are one of the levels of structure falling somewhere in between the single-ant level and the colony level.

anteater: Precisely. There exists a special kind of team, which I call a "signal"—and all the higher levels of structure are based on signals. In fact, all the higher entities are collections of signals acting in concert. There are teams on higher levels whose members are not ants, but teams on lower levels. Eventually you reach the lowest-level teams—which is to say, signals—and below them, ants.

achilles: Why do signals deserve their suggestive name?

anteater: It comes from their function. The effect of signals is to transport ants of various specializations to appropriate parts of the colony. So the typical story of a signal is thus: It comes into existence by exceeding the threshold needed for survival, then it migrates for

some distance through the colony, and at some point it more or less disintegrates into its individual members, leaving them on their own.

achilles: It sounds like a wave, carrying sand dollars and seaweed from afar, and leaving them strewn, high and dry, on the shore.

anteater: In a way that's analogous, since the team does indeed deposit something which it has carried from a distance, but whereas the water in the wave rolls back to the sea, there is no analogous carrier substance in the case of a signal, since the ants themselves compose it.

tortoise: And I suppose that a signal loses its coherency just at some spot in the colony where ants of that type were needed in the first place.

anteater: Naturally.

achilles: Naturally? It's not so obvious to ***me*** that a signal should always go just where it is needed. And even if it goes in the right direction, how does it figure out where to decompose? How does it know it has arrived?

anteater: Those are extremely important matters, since they involve explaining the existence of purposeful behavior—or what seems to be purposeful behavior—on the part of signals. From the description, one would be inclined to characterize the signals' behavior as being oriented toward filling a need, and to call it "purposeful." But you can look at it otherwise.

achilles: Oh, wait. Either the behavior ***is*** purposeful, or it is ***not.*** I don't see how you can have it both ways.

anteater: Let me explain my way of seeing things, and then see if you agree. Once a signal is formed, there is no awareness on its part that it should head off in any particular direction. But here the delicate caste distribution plays a crucial role. It is what determines the motion of signals through the colony, and also how long a signal will remain stable, and where it will "dissolve."

achilles: SO **everything depends on the caste distribution, eh?**

anteater: Right. Let's say a signal is moving along. As it goes, the ants which compose it interact, either by direct contact or by exchange of scents, with ants of the local neighborhoods which it passes through. The contacts and scents provide information about local matters of urgency, such as nest building, or nursing, or whatever. The signal will remain glued together as long as the local needs are different

from what it can supply; but if it ***can*** contribute, it disintegrates, spilling a fresh team of usable ants onto the scene. Do you see now how the caste distribution acts as an overall guide of the teams inside the colony?

achilles: I **do see that.**

anteater: And do you see how this way of looking at things requires attributing no sense of purpose to the signal?

achilles: I think so. Actually, I'm beginning to see things from two different vantage points. From an ant's-eye point of view, a signal has ***no*** purpose. The typical ant in a signal is just meandering around the colony, in search of nothing in particular, until it finds that it feels like stopping. Its teammates usually agree, and at that moment the team unloads itself by crumbling apart, leaving just its members but none of its coherency. No planning is required, no looking ahead; nor is any search required to determine the proper direction. But from the ***colony's*** point of view, the team has just responded to a message which was written in the language of the caste distribution. Now from this perspective, it looks very much like purposeful activity.

crab: What would happen if the caste distribution were entirely random? Would signals still band and disband?

anteater: Certainly. But the colony would not last long, due to the meaninglessness of the caste distribution.

crab: Precisely the point I wanted to make. Colonies survive because their caste distribution has ***meaning***, and that meaning is a holistic aspect, invisible on lower levels. You lose explanatory power unless you take that higher level into account.

anteater: I **see your side; but** I **believe you see things too narrowly.**

CRAB: HOW SO?

anteater: Ant colonies have been subjected to the rigors of evolution for billions of years. A few mechanisms were selected for, and most were selected against. The end result was a set of mechanisms which make ant colonies work as we have been describing. If you could watch the whole process in a movie—running a billion or so times faster than life, of course-^the emergence of various mechanisms would be seen as natural responses to external pressures, just as bubbles in boiling water are natural responses to an external heat source. I don't suppose you see "meaning" and "purpose" in the bubbles in boiling water—or do you?

crab: NO, but—

anteater: NOW that's ***my*** point. No matter how big a bubble is, it owes its existence to processes on the molecular level, and you can forget about any "higher-level laws." The same goes for ant colonies and their teams. By looking at things from the vast perspective of evolution, you can drain the whole colony of meaning and purpose. They become superfluous notions.

achilles: Why, then, Dr. Anteater, did you tell me that you talked with Aunt Hillary? It now seems that you would deny that she can talk or think at all.

anteater: I am not being inconsistent, Achilles. You see, I have as much difficulty as anyone else in seeing things on such a grandiose time scale, so I find it much easier to change points of view. When I do so, forgetting about evolution and seeing things in the here and now, the vocabulary of teleology comes back: the ***meaning*** of the caste distribution and the ***purposefulness*** of signals. This not only happens when I think of ant colonies, but also when I think about my own brain and other brains. However, with some effort I can always remember the other point of view if necessary, and drain all these systems of meaning, too.

crab: Evolution certainly works some miracles. You never know the next trick it will pull out of its sleeve. For instance, it wouldn't surprise me one bit if it were theoretically possible for two or more "signals" to pass through each other, each one unaware that the other one is also a signal; each one treating the other as if it were just part of the background population.

anteater: It is better than theoretically possible; in fact it happens routinely!

achilles: Hmm What a strange image that conjures up in my mind.

I can just imagine ants moving in four different directions, some black, some white, criss-crossing, together forming an orderly pattern, almost like—like—

tortoise: A fugue, perhaps?

achilles: Yes—that's it! An ant fugue!

crab: An interesting image, Achilles. By the way, all that talk of boiling water made me think of tea. Who would like some more?

achilles: I could do with another cup, Mr. C.

crab: Very good.

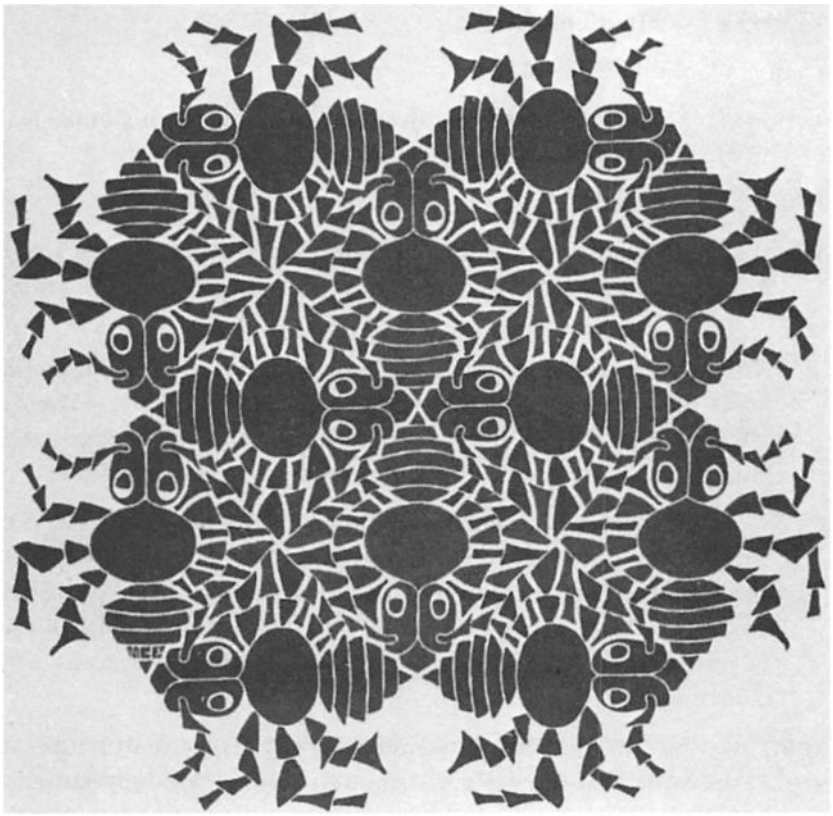

An "Ant Fugue" drawn by M. C. Escher (woodcut, 1953.)

ACHILLES: DO you suppose one could separate out the different visual "voices" of such an "ant fugue"? I know how hard it is for me—

TORTOISE: Not for me, thank you.

ACHILLES: —to track a single voice—

ANTEATER: I'd like some too, Mr. Crab—

ACHILLES: —in a musical fugue—

ANTEATER: —if it isn't too much trouble.

ACHILLES: —when all of them—

CRAB: Not at all. Four cups of tea—

TORTOISE: Three!

achilles: —are going at once.

CRAB: —coming right up!

anteater: That's an interesting thought, Achilles. But it's unlikely that anyone could draw such a picture in a convincing way.

achilles: That's too bad.

tortoise: Perhaps you could answer this, Dr. Anteater. Does a signal, from its creation until its dissolution, always consist of the same set of ants?

ANTEATER: As a matter of fact, the individuals in a signal sometimes break off and get replaced by others of the same caste, if there are a few in the area. Most often, signals arrive at their disintegration points with nary an ant in common with their starting lineup.

crab: I can see that the signals are constantly affecting the caste distribution throughout the colony, and are doing so in response to the internal needs of the colony—which in turn reflect the external situation which the colony is faced with. Therefore the caste distribution, as you said, Dr. Anteater, gets continually updated in a way which ultimately reflects the outer world.

achilles: But what about those intermediate levels of structure? You were saying that the caste distribution should best be pictured not in terms of ants or signals, but in terms of teams whose members were other teams, whose members were other teams, and so on until you come down to the ant level. And you said that that was the key to understanding how it was possible to describe the caste distribution as encoding pieces of information about the world.

anteater: Yes, we are coming to all that. I prefer to give teams of a sufficiently high level the name of "symbols." Mind you, this sense of the word has some significant differences from the usual sense. My "symbols" are ***active subsystems*** of a complex system, and they are composed of lower-level active subsystems. . . . They are therefore quite different from ***passive*** symbols, external to the system, such as letters of the alphabet or musical notes, which sit there immobile waiting for an active system to process them.

achilles: Oh, this is rather complicated, isn't it? I just had no idea that ant colonies had such an abstract structure.

anteater: Yes, it's quite remarkable. But all these layers of structure are necessary for the storage of the kinds of knowledge which enable an

organism to be "intelligent" in any reasonable sense of the word. Any system which has a mastery of language has essentially the same underlying sets of levels.

achilles: Now just a cotton-picking minute. Are you insinuating that my brain consists of, at bottom, just a bunch of ants running around?

anteater: Oh, hardly. You took me a little too literally. The lowest level may be utterly different. Indeed, the brains of anteaters, for instance, are not composed of ants. But when you go up a level or two in a brain, you reach a level whose elements have exact counterparts in other systems of equal intellectual strength—such as ant colonies.

tortoise: That is why it would be reasonable to think of mapping your brain, Achilles, onto an ant colony, but not onto the brain of a mere ant.

achilles: I appreciate the compliment. But how would such a mapping be carried out? For instance, what in my brain corresponds to the low-level teams which you call signals?

anteater: Oh, I but dabble in brains, and therefore couldn't set up the map in its glorious detail. But—and correct me if I'm wrong, Mr. Crab—I would surmise that the brain counterpart to an ant colony's signal is the firing of a neuron; or perhaps it is a larger-scale event, such as a pattern of neural firings.

crab: I would tend to agree. But don't you think that, for the purposes of our discussion, delineating the exact counterpart is not in itself crucial, desirable though it might be? It seems to me that the main idea is that such a correspondence does exist, even if we don't know exactly how to define it right now. I would only question one point, Dr. Anteater, which you raised, and that concerns the level at which one can have faith that the correspondence begins. You seemed to think that a ***signal*** might have a direct counterpart in a brain; whereas I feel that it is only at the level of your ***active symbols*** and above that it is likely that a correspondence must exist.

anteater: Your interpretation may very well be more accurate than mine, Mr. Crab. Thank you for bringing out that subtle point.

achilles: What does a symbol do that a signal couldn't do?

anteater: It is something like the difference between words and letters. Words, which are meaning-carrying entities, are composed of letters, which in themselves carry no meaning. This gives a good idea of the difference between symbols and signals. In fact it is a useful analogy,

as long as you keep in mind the fact that words and letters ar***e passive,*** symbols and signals are ***active.***

achilles: I'll do so, but I'm not sure I understand why it is so vital to stress the difference between active and passive entities.

anteater: The reason is that the meaning which you attribute to any passive symbol, such as a word on a page, actually derives from the meaning which is carried by corresponding active symbols in your brain. So that the meaning of passive symbols can only be properly understood when it is related to the meaning of active symbols.

achilles: All right. But what is it that endows a ***symbol—*** an active one, to be sure—with meaning, when you say that a ***signal,*** which is a perfectly good entity in its own right, has none?

anteater: It all has to do with the way that symbols can cause other symbols to be triggered. When one symbol becomes active, it does not do so in isolation. It is floating about, indeed, in a medium, which is characterized by its caste distribution.

crab: Of course, in a brain there is no such thing as a caste distribution, but the counterpart is the "brain state." There, you describe the states of all the neurons, and all the interconnections, and the threshold for firing of each neuron.

anteater: Very well; let's lump "caste distribution" and "brain state" under a common heading, and call them just the "state." Now the state can be described on a low level or on a high level. A low-level description of the state of an ant colony would involve painfully specifying the location of each ant, its age and caste, and other similar items. A very detailed description, yielding practically no global insight as to ***why*** it is in that state. On the other hand, a description on a high level would involve specifying which symbols could be triggered by which combinations of other symbols, under what conditions, and so forth.

achilles: What about a description on the level of signals, or teams?

anteater: A description on that level would fall somewhere in between the low-level and symbol-level descriptions. It would contain a great deal of information about what is actually going on in specific locations throughout the colony, although certainly less than an ant- by-ant description, since teams consist of clumps of ants. A team-by- team description is like a summary of an ant-by-ant description. However, you have to add extra things which were not present in the ant-by-ant description—such as the relationships between teams,

and the supply of various castes here and there. This extra complication is the price you pay for the right to summarize.

achilles: It is interesting to me to compare the merits of the descriptions at various levels. The highest-level description seems to carry the most explanatory power, in that it gives you the most intuitive picture of the ant colony, although strangely enough, it leaves out seemingly the most important feature—the ants.

anteater: But you see, despite appearances, the ants are not the most important feature. Admittedly, were it not for them, the colony wouldn't exist; but something equivalent—a brain—can exist, ant- free. So, at least from a high-level point of view, the ants are dispensable.

achilles: I'm sure no ant would embrace your theory with eagerness.

anteater: Well, I never met an ant with a high-level point of view.

crab: What a counterintuitive picture you paint, Dr. Anteater. It seems that, if what you say is true, in order to grasp the whole structure, you have to describe it omitting any mention of its fundamental building blocks.

anteater: Perhaps I can make it a little clearer by an analogy. Imagine you have before you a Charles Dickens novel.

achilles: ***The Pickwick Papers*****—will that do?**

anteater: Excellently! And now imagine trying the following game: You must find a way of mapping letters onto ideas, so that the entire ***Pickwick Papers*** makes sense when you read it letter by letter.

achilles: Hmm. . . . You mean that every time I hit a word such as "the," I have to think of three definite concepts, one after another, with no room for variation?

anteater: Exactly. They are the "t"-concept, the "h"-concept, and the "e"-concept—and every time, those concepts are as they were the preceding time.

achilles: Well, it sounds like that would turn the experience of "reading" ***The Pickwick Papers*** into an indescribably boring nightmare. It would be an exercise in meaninglessness, no matter what concept I associated with each letter.

anteater: Exactly. There is no natural mapping from the individual letters into the real world. The natural mapping occurs on a higher level—between words, and parts of the real world. If you wanted to describe the book, therefore, you would make no mention of the letter level.

achilles: Of course not! I'd describe the plot and the characters, and so forth.

anteater: SO there you are. You would omit all mention of the building blocks, even though the book exists thanks to them. They are the medium, but not the message.

achilles: All right—but what about ant colonies?

anteater: Here, there are active signals instead of passive letters, and active symbols instead of passive words—but the idea carries over.

achilles: DO you mean I couldn't establish a mapping between signals and things in the real world?

anteater: **You** would find that you could not do it in such a way that the triggering of new signals would make any sense. Nor could you succeed on any lower level—for example, the ant level. Only on the symbol level do the triggering patterns make sense. Imagine, for instance, that one day you were watching Aunt Hillary when I arrived to pay a call. You could watch as carefully as you wanted, and yet you would probably perceive nothing more than a rearrangement of ants.

achilles: I'm **sure that's accurate.**

anteater: And yet, as I watched, reading the higher level instead of the lower level, I would see several dormant symbols being awakened, those which translate into the thought "Oh, here's that charming Dr. Anteater again—how pleasant!"—or words to that effect.

achilles: That sounds like what happened when the four of us all found different levels to read in the MU-picture—or at least ***three*** of us did. . . .

tortoise: What an astonishing coincidence that there should be such a resemblance between that strange picture which I chanced upon in the ***Well-Tempered Clavier*** and the trend of our conversation.

achilles: DO **you think it's just coincidence?**

tortoise: Of **course.**

anteater: Well, I hope you can grasp now how the thoughts in Aunt Hillary emerge from the manipulation of symbols composed of signals composed of teams composed of lower-level teams, all the way down to ants.

achilles: Why do you call it "symbol manipulation"? Who does the manipulating, if the symbols are themselves active? Who is the agent?

anteater: This gets back to the question that you earlier raised about purpose. You're right that symbols themselves are active, but the activities which they follow are nevertheless not absolutely free. The activities of all symbols are strictly determined by the state of the full system in which they reside. Therefore, the full system is responsible for how its symbols trigger each other, and so it is quite reasonable to speak of the full system as the "agent." As the symbols operate, the state of the system gets slowly transformed, or updated. But there are many features that remain over time. It is this partially constant, partially varying system that is the agent. One can give a name to the full system. For example, Aunt Hillary is the "who" who can be said to manipulate her symbols; and you are similar, Achilles.

achilles: That's quite a strange characterization of the notion of who I am. I'm not sure I can fully understand it, but I will give it some thought.

tortoise: It would be quite interesting to follow the symbols in your brain as you do that thinking about the symbols in your brain.

achilles: That's too complicated for me. I have trouble enough just trying to picture how it is possible to look at an ant colony and read it on the symbol level. I can certainly imagine perceiving it at the ant level; and with a little trouble, I can imagine what it must be like to perceive it at the signal level; but what in the world can it be like to perceive an ant colony at the symbol level?

anteater: One learns only through long practice. But when one is at my stage, one reads the top level of an ant colony as easily as you yourself read the **"mu"** in the MU-picture.

achilles: Really? That must be an amazing experience.

anteater: In a way—but it is also one which is quite familiar to you, Achilles.

achilles: Familiar to me? What do you mean? I have never looked at an ant colony on anything but the ant level.

anteater: Maybe not; but ant colonies are no different from brains in many respects.

achilles: I **have never seen nor read any brain either, however.**

anteater: What about your ***own*** brain? Aren't you aware of your own thoughts? Isn't that the essence of consciousness? What else are you doing but reading your own brain directly at the symbol level?

achilles: I never thought of it that way. You mean that I bypass all the lower levels, and see only the topmost level?

anteater: That's the way it is, with conscious systems. They perceive themselves on the symbol level only, and have no awareness of the lower levels, such as the signal levels.

achilles: Does it follow that in a brain, there are active symbols that are constantly updating themselves so that they reflect the overall state of the brain itself, always on the symbol level?

anteater: Certainly. In any conscious system there are symbols that represent the brain state, and they are themselves part of the very brain state which they symbolize. For consciousness requires a large degree of self-consciousness.

achilles: That is a weird notion. It means that although there is frantic activity occurring in my brain at all times, I am capable of registering that activity in only one way—on the symbol level; and I am completely insensitive to the lower levels. It is like being able to read a Dickens novel by direct visual perception, without ever having learned the letters of the alphabet. I can't imagine anything as weird as that really happening.

crab: But precisely that sort of thing ***did*** happen when you read **"mu,"** without perceiving the lower levels **"holism"** and **"reductionism."**

achilles: You're right—I bypassed the lower levels, and saw only the top. I wonder if I'm missing all sorts of meaning on lower levels of my brain as well, by reading only the symbol level. It's too bad that the top level doesn't contain all the information about the bottom level, so that by reading the top, one also learns what the bottom level says. But I guess it would be naive to hope that the top level encodes anything from the bottom level—it probably doesn't percolate up. The MU-picture is the most striking possible example of that: There, the topmost level says only **"mu,"** which bears no relation whatever to the lower levels!

crab: That's absolutely true. ***(Picks up the MU-picture, to inspect it more closely.)*** Hmm. . . . There's something strange about the smallest letters in this picture; they're very wiggly . . .

anteater: Let me take a look. ***(Peers closely at the MU-picture.)*** I think there's yet another level, which all of us missed!

tortoise: Speak for yourself, Dr. Anteater.

achilles: Oh, no—that can't be! Let me see. ***(Looks very carefully.)*** I know the rest of you won't believe this, but the message of this picture is staring us all in the face, hidden in its depths. It is simply one word, repeated over and over again, like a mantra—but what an important

one: **"mu"!** What do you know! It is the same as the top level! And none of us suspected it in the least.

crab: We would never have noticed it if it hadn't been for you, Achilles.

anteater: I wonder if the coincidence of the highest and lowest levels happened by chance? Or was it a purposeful act carried out by some creator?

CRAB: HOW could one ever decide that?

tortoise: I don't see any way to do so, since we have no idea why that particular picture is in the Crab's edition of the ***Well-Tempered Clavier.***

anteater: Although we have been having a lively discussion, I have still managed to listen with a good fraction of an ear to this very long and complex four-voice fugue. It is extraordinarily beautiful.

tortoise: It certainly is. And now, in just a moment, comes an organ point.

achilles: Isn't an organ point what happens when a piece of music slows down slightly, settles for a moment or two on a single note or chord, and then resumes at normal speed after a short silence?

tortoise: NO, you're thinking of a "fermata"—a sort of musical semicolon. Did you notice there was one of those in the prelude?

achilles: I guess I must have missed it.

tortoise: Well, you have another chance coming up to hear a fermata —in fact, there are a couple of them coming up, toward the end of this fugue.

achilles: Oh, good. You'll point them out in advance, won't you?

tortoise: If **you like.**

achilles: But do tell me, what is an organ point?

tortoise: An organ point is the sustaining of a single note by one of the voices in a polyphonic piece (often the lowest voice), while the other voices continue their own independent lines. This organ point is on the note of G. Listen carefully, and you'll hear it.

anteater: There occurred an incident one day when I visited with Aunt Hillary which reminds me of your suggestion of observing the symbols in Achilles' brain as they create thoughts which are about themselves.

crab: DO tell us about it.

anteater: Aunt Hillary had been feeling very lonely, and was very happy to have someone to talk to that day. So she gratefully told me

to help myself to the juiciest ants I could find. (She's always been most generous with her ants.)

achilles: Gee!

anteater: It just happened that I had been watching the symbols which were carrying out her thoughts, because in them were some particularly juicy-looking ants.

achilles: Gee!

anteater: SO I helped myself to a few of the fattest ants which had been parts of the higher-level symbols which I had been reading. Specifically, the symbols which they were part of were the ones which had expressed the thought "Help yourself to any of the ants which look appetizing."

achilles: Gee!

anteater: Unfortunately for them, but fortunately for me, the little bugs didn't have the slightest inkling of what they were collectively telling me, on the symbol level.

achilles: Gee! That is an amazing wraparound. They were completely unconscious of what they were participating in. Their acts could be seen as part of a pattern on a higher level, but of course they were completely unaware of that. Ah, what a pity—a supreme irony, in fact —that they missed it.

crab: **You** are right, Mr. **T**—that was a lovely organ point.

anteater: I had never heard one before, but that one was so conspicuous that no one could miss it. Very effective.

achilles: What? Has the organ point already occurred? How can I not have noticed it, if it was so blatant?

tortoise: Perhaps you were so wrapped up in what you were saying that you were completely unaware of it. Ah, what a pity—a supreme irony, in fact—that you missed it.

crab: Tell me, does Aunt Hillary live in an anthill?

anteater: Well, she owns a rather large piece of property. It used to belong to someone else, but that is rather a sad story. In any case, her estate is quite expansive. She lives rather sumptuously, compared to many other colonies.

achilles: **How** does that jibe with the communistic nature of ant colonies which you earlier described to us? It sounds quite inconsistent, to me, to preach communism and to live in a fancy estate!

Prelude . . . Ant Fugue

anteater: The communism is on the ant level. In an ant colony all ants work for the common good, even to their own individual detriment at times. Now this is simply a built-in aspect of Aunt Hillary's structure, but for all I know, she may not even be aware of this internal communism. Most human beings are not aware of anything about their neurons; in fact they probably are quite content not to know anything about their brains, being somewhat squeamish creatures. Aunt Hillary is also somewhat squeamish; she gets rather antsy whenever she starts to think about ants at all. So she avoids thinking about them whenever possible. I truly doubt that she knows anything about the communistic society which is built into her very structure. She herself is a staunch believer in libertarianism—you know, laissez- faire and all that. So it makes perfect sense, to me at least, that she should live in a rather sumptuous manor.

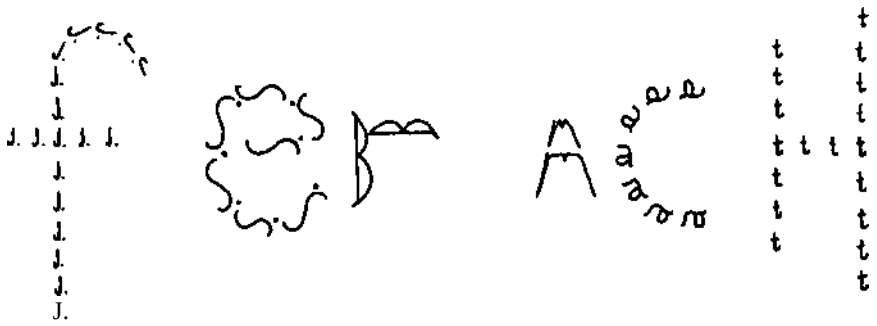

ILLUSTRATION BY THE AUTHOR.

tortoise: As I turned the page just now, while following along in this lovely edition of the ***Well-Tempered Clavier,*** I noticed that the first of the two fermatas is coming up soon—so you might listen for it, Achilles.

achilles: I will, I will.

tortoise: Also, there's a most curious picture facing this page.

crab: Another one? What next?

tortoise: See for yourself. ***(Passes the score over to the Crab.)***

crab: Aha! It's just a few bunches of letters. Let's see—there are various numbers of the letters "J," "S," "B," "m," "a," and "t." It's strange, how the first three letters grow, and then the last three letters shrink again.

anteater: May I see it?

crab: Why, certainly.

anteater: Oh, by concentrating on details, you have utterly missed the big picture. In reality, this group of letters is "f," "e," "r," "A," "C," "H," without any repetitions. First they get smaller, then they get bigger. Here, Achilles—what do you make of it?

achilles: Let me see. Hmm. Well, I see it as a set of uppercase letters which grow as you move to the right.

tortoise: **Do** they spell anything?

achilles: Ah . . . "J. S. BACH." Oh! I understand now. It's Bach's name!

tortoise: Strange that you should see it that way. **I** see it as a set of lower-case letters, shrinking as they move to the right-, and . . . spelling o u t . . . the name o f . . . ***(Slow down slightly, especially drawing out the last few words. Then there is a brief silence. Suddenly he resumes as if nothing unusual had happened.)***—"fermat."

achilles: Oh, you've got Fermat on the brain, I do believe. You see Fermat's Last Theorem everywhere.

anteater: **You** were right, Mr. Tortoise—**I** just heard a charming little fermata in the fugue.

CRAB: **SO did I.**

achilles: DO you mean everybody heard it but me? I'm beginning to feel stupid.

tortoise: There, there, Achilles—don't feel bad. I'm sure you won't miss Fugue's Last Fermata (which is coming up quite soon). But, to return to our previous topic, Dr. Anteater, what is the very sad story which you alluded to, concerning the former owner of Aunt Hillary's property?

anteater: The former owner was an extraordinary individual, one of the most creative ant colonies who ever lived. His name was Johant Sebastiant Fermant, and he was a mathematiciant by vocation, but a musiciant by avocation.

ACHILLES: **HOW very versantile of him!**

anteater: At the height of his creative powers, he met with a most untimely demise. One day, a very hot summer day, he was out soaking up the warmth, when a freak thundershower—the kind that hits only once every hundred years or so—appeared from out of the blue and thoroughly drenched J. S. F. Since the storm came utterly without warning, the ants got completely disoriented and confused. The

intricate organization that had been so finely built up over decades all went down the drain in a matter of minutes. It was tragic.

achilles: DO you mean that all the ants drowned, which obviously would spell the end of poor J. S. F.?

anteater: Actually, no. The ants managed to survive, every last one of them, by crawling onto various sticks and logs that floated above the raging torrents. But when the waters receded and left the ants back on their home grounds, there was no organization left. The caste distribution was utterly destroyed, and the ants themselves had no ability to reconstruct what had once before been such a finely tuned organization. They were as helpless as the pieces of Humpty Dumpty in putting themselves back together again. I myself tried, like all the king's horses and all the king's men, to put poor Fermant together again. I faithfully put out sugar and cheese, hoping against hope that somehow Fermant would reappear . . . ***(Pulls out a handkerchief and wipes his eyes.)***

achilles: ' **How** valiant of you! **I** never knew Anteaters had such big hearts.

anteater: But it was all to no avail. He was gone, beyond reconstitution. However, something very strange then began to take place: over the next few months, the ants that had been components of J. S. F. slowly regrouped, and built up a new organization. And thus was Aunt Hillary born.

crab: Remarkable! Aunt Hillary is composed of the very same ants as Fermant was?

anteater: Well, originally she was, yes. By now, some of the older ants have died, and been replaced. But there are still many holdovers from thej. S. F.-days.

crab: And can't you recognize some ofj. S. F.'s old traits coming to the fore, from time to time, in Aunt Hillary?

anteater: Not a one. They have nothing in common. And there is no reason they should, as I see it. There are, after all, often several distinct ways to rearrange a group of parts to form a "sum." And Aunt Hillary was just a new "sum" of the old parts. Not ***more*** than the sum, mind you—just that particular ***kind*** of sum.

tortoise: Speaking of sums, I am reminded of number theory, where occasionally one will be able to take apart a theorem into its component symbols, rearrange them in a new order, and come up with a new theorem.

anteater: I've never heard of such a phenomenon, although I confess to being a total ignoramus in the field.

achilles: Nor have I heard of it—and I am rather well versed in the field, if I don't say so myself. I suspect Mr. T is just setting up one of his elaborate spoofs. I know him pretty well by now.

anteater: Speaking of number theory, I am reminded of J. S. F. again, . for number theory is one of the domains in which he excelled. In fact, he made some rather remarkable contributions to number theory. Aunt Hillary, on the other hand, is remarkably dull-witted in anything that has even the remotest connection with mathematics. Also, she has only a rather banal taste in music, whereas Sebastiant was extremely gifted in music.

achilles: I am very fond of number theory. Could you possibly relate to us something of the nature of Sebastiant's contributions?

anteater: Very well, then. ***(Pauses for a moment to sip his tea, then resumes.)*** Have you heard of Fourmi's infamous "Well-Tested Conjecture"?

achilles: I'm not s u r e . . . . It sounds strangely familiar, and yet I can't quite place it.

anteater: It's a very simple idea. Lierre de Fourmi, a mathemadciant by vocation but lawyer by avocation, had been reading in his copy of the classic text ***Arithmetica*** by Di of Antus, and came across a page containing the equation

$$2^a + 2^b = 2^c$$

He immediately realized that this equation has infinitely many solutions ***a, b, c,*** and then wrote in the margin the following notorious comment:

> The equation
>
> $$n^a + n^b = n^c$$
>
> has solutions in positive integers *a*, *b*, *c*, and n only when $n = 2$ (and then there are infinitely many triplets *a*, *b*, *c* which satisfy the equation); but there are no solutions for n > 2.1 have discovered a truly marvelous proof of this statement, which, unfortunately, is so small that it would be well-nigh invisible if written in the margin.

Ever since that year, some three hundred days ago, mathematiciants have been vainly trying to do one of two things: either to prove Fourmi's claim, and thereby vindicate Fourmi's reputation, which,

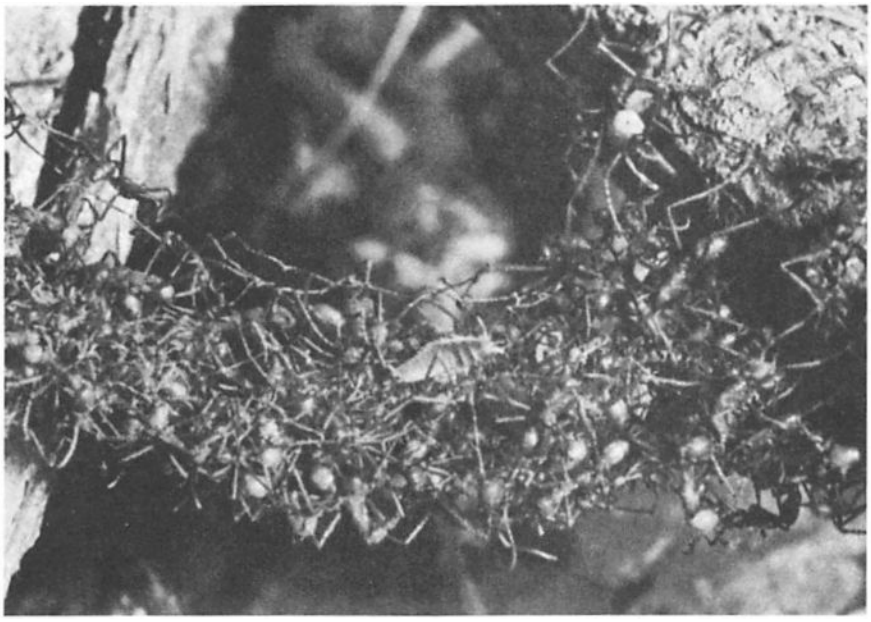

During emigrations army anls sometimes crcate bridges of their own bodies. In this photograph of such a bridge (de Fourmi Lierre), the workers *Eaton burchelli* colony can be seen linking their legs and, along the top of the bridge, hooking their tarsal claws together to form irregular systems of chains. A symbiotic silverfish. *Trichalelura manni*, is seen crossing the bridge in the ccnter (From E. O. Wilson, *The Insect Societies.* Photograph courtesy of C. W. Rettenmeyer.).

> although very high, has been somewhat tarnished by skeptics who think he never really found the proof he claimed to have found—or else to refute the claim, by finding a counterexample: a set of four integers *a, b, c,* and *n,* with $n > 2$, which satisfy the equation. Until very recently, every attempt in either direction had met with failure. To be sure, the Conjecture has been verified for many specific values of *n*—in particular, all *n* up to 125,000. But no one had succeeded in proving it for *all n*—no one, that is, until Johant Sebastiant Fermant came upon the scene. It was he who found the proof that cleared Fourmi's name. It now goes under the name "Johant Sebas- tiant's Well-Tested Conjecture."

achilles: Shouldn't it be called a "Theorem" rather than a "Conjecture," if it's finally been given a proper proof?

anteater: Strictly speaking, you're right, but tradition has kept it this way.

tortoise: What sort of music did Sebastiant do?

anteater: He had great gifts for composition. Unfortunately, his greatest work is shrouded in mystery, for he never reached the point of publishing it. Some believe that he had it all in his mind; others are more unkind, saying that he probably never worked it out at all, but merely blustered about it.

achilles: What was the nature of this magnum opus?

anteater: It was to be a giant prelude and fugue; the fugue was to have twenty-four voices, and to involve twenty-four distinct subjects, one in each of the major and minor keys.

achilles: It would certainly be hard to listen to a twenty-four-voice fugue as a whole!

crab: Not to mention composing one!

anteater: But all that we know of it is Sebastiant's description of it, which he wrote in the margin of his copy of Buxtehude's Preludes and Fugues for Organ. The last words which he wrote before his tragic demise were:

> I have composed a truly marvelous fugue. In it, I have added together the power of 24 keys, and the power of 24 themes; I came up with a fugue with the power of 24 voices. Unfortunately, this margin is too narrow to contain it.

And the unrealized masterpiece simply goes by the name "Fermant's Last Fugue."

achilles: Oh, that **is** unbearably tragic.

tortoise: Speaking of fugues, this fugue that we have been listening to is nearly over. Toward the end, there occurs a strange new twist on its theme. ***(Flips the page in the*** Well-Tempered Clavier J Well, what have we here? A new illustration—how appealing! ***(Shows it to the Crab.)***

crab: Well, what have we here? Oh, I see: it's **"holismionism,"** written in large letters that first shrink and then grow back to their original size. But that doesn't make any sense, because it's not a word. Oh me, oh my! ***(Passes it to the Anteater.)***

anteater: Well, what have we here? Oh, I see: it's **"reductholism,"** written in small letters that first grow and then shrink back to their original size. But that doesn't make any sense, because it's not a word. Oh my, oh me! ***(Passes it to Achilles.)***

achilles: I know the rest of you won't believe this, but in fact this picture consists of the word **"holism"** written twice, with the letters

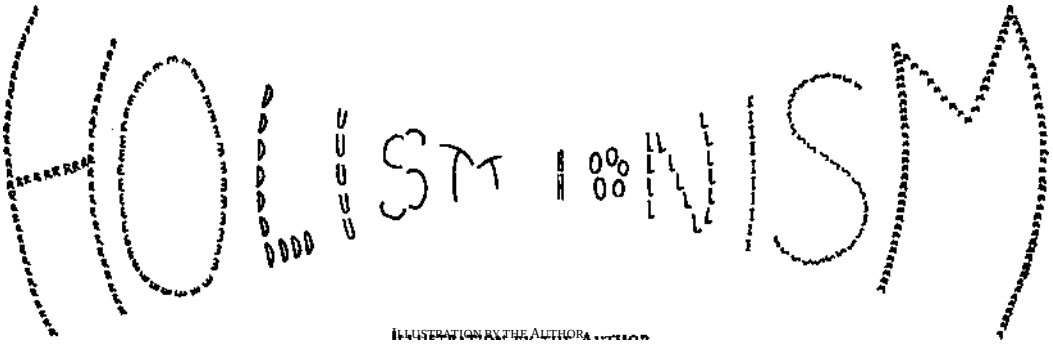

ILLUSTRATION BY THE AUTHOR

continually shrinking as they proceed from left to right. ***(Returns it to the Tortoise.)***

tortoise: I know the rest of you won't believe this, but in fact this picture consists of the word **"reductionism"** written once, with the letters continually growing as they proceed from left to right.

achilles: At last—I heard the new twist on the theme this time! I am so glad that you pointed it out to me, Mr. Tortoise. Finally, I think I am beginning to grasp the art of listening to fugues.

Reflections

Is a soul greater than the hum of its parts? The participants in the preceding dialogue seem to have divergent views on this question. What is certain and agreed upon, however, is that the collective behavior of a system of individuals can have many surprising properties.

Many people, on reading this dialogue, are reminded of the seemingly purposive, selfish, survival-oriented behavior of countries that emerges somehow from the habits and institutions of their citizens: their educational system, legal structure, religions, resources, style of consumption and level of expectations, and so on. When a tight organization forms from distinct individuals—particularly when contributions to the organization are not traceable to specific individuals in the lower level— we tend to see it as a higher-level individual and often speak of it in

anthropomorphic terms. A newspaper article about a terrorist group described it as "playing its cards extremely close to its chest." It is often said of Russia that it "desires" world recognition of its might because it "suffers" from a "long-standing inferiority complex" with respect to Western Europe. While admittedly metaphors, these examples serve to demonstrate how strong the urge is to personify organizations.

The component individuals of organizations—secretaries, workers, bus drivers, executives, and so on—have their own goals in life, which, one might expect, would come into conflict with any higher-level entity of which they formed a part, but there is an effect (which many students of political science would regard as insidious and sinister) whereby the organization co-opts and exploits these very goals, taking advantage of the individuals' pride, need for self-esteem, and so on, and turning them back to its own profits. There emerges from all the many low-level goals a kind of higher-level momentum that subsumes all of them, that sweeps them along and thereby perpetuates itself.

Therefore it is perhaps not so silly for the Tortoise to object to Achilles' comparison of himself to an ant and to prefer an attempt by Achilles to "map himself," at a suitable level, onto an ant colony. Similarly, we may sometimes wonder to ourselves "What is it like to be China? How different from that would it feel to be the United States?" Do such questions makes any kind of sense at all? We shall postpone detailed discussion of them until after Nagel's piece on bats (selection 24). Nonetheless, let us think a bit right now about whether it makes sense to think of "being" a country. Does a country have thoughts or beliefs? It all comes down to whether a country has a ***symbol*** level, in the sense that Aunt Hillary does. Instead of saying that a system "has a symbol level," we might instead say, "It is a representational system."

This concept of "representational system" is a crucial one in this book, and needs a somewhat precise definition. By "representational system" we will mean an active, self-updating collection of structures organized to "mirror" the world as it evolves. A painting, no matter how representational, would thus be excluded, since it is static. Curiously, we mean also to exclude mirrors themselves, although the argument could be made that the set of images in a mirror keeps quite up to date with the world! The lack in this case is twofold. First, the mirror itself does not make any distinction between images of different objects—it mirrors the universe, but sees no ***categories.*** In fact, a mirror makes only ***one*** image— it is in the eye of the beholder that the mirror's single image breaks up into "separate" images of many distinct objects. A mirror cannot be said to perceive—only to reflect. Second, the image in a mirror is not an autonomous structure with its own "life"; it depends directly on the

external world. If the lights are turned off, it goes away. A representational system should be able to keep on going even if cut off from contact with the reality it is "reflecting"—although you now see that "reflection" is not quite a rich enough metaphor. The isolated representational structures should now continue to evolve in a way that reflects, if not the true way the world will evolve, at least a probable way. Actually, a good representational system will sprout parallel branches for various possibilities that can be reasonably anticipated. Its internal models will, in the metaphorical sense defined in the Reflections on "Rediscovering the Mind," go into superpositions of states, each with an associated subjective estimate of likelihood.

In brief, then, a representational system is built on categories; it sifts incoming data into those categories, when necessary refining or enlarging its network of internal categories; its representations or "symbols" interact among themselves according to their own internal logic; this logic, although it runs without ever consulting the external world, nevertheless creates a faithful enough model of the way the world works that it manages to keep the symbols pretty much "in phase" with the world they are supposed to be mirroring. A television is thus not a representational system, as it indiscriminately throws dots onto its screen without regard to what kinds of things they represent, and the patterns on the screen do not have autonomy—they are just passive copies of things "out there." By contrast, a computer program that can "look" at a scene and tell you what is in that scene comes closer to being a representational system. The most advanced artificial intelligence work on computer vision hasn't yet cracked that nut. A program that could look at a scene and tell you not only what kinds of things are in the scene, but also what probably caused that scene and what will probably ensue in it—that is what we mean by a representational system. In this sense, is a country a representational system? Does a country have a symbol level? We'll leave this one for you to ponder on.

One of the crucial notions of the Ant Fugue is the "caste distribution" or "state," for it is claimed that that is a causal agent in determining the future of the organism. Yet this seems to contradict the idea that all of a system's behavior comes from underlying laws—those of ants or neurons, in the case of colonies or brains—but ultimately, in either case, those of particles. Is there such a thing as "downward causality"—. put starkly, the notion that "a thought can influence the path of an electron"?

In ***Inside the Brain*** by William Calvin and George Ojemann, there is a provocative series of questions asked about a neural firing. "What starts it?" they ask. What causes the sodium channels to open up? (The function

of the sodium channels is to let sodium ions into the neuron, and when their concentration is high enough, that then triggers the release of the neurotransmitters, whose flow from one neuron to another constitutes the essence of neural firing.) The answer is, the sodium channels are voltage-sensitive, and they have just been hit by a strong enough voltage pulse to flip their state from closed to open.

"But what causes the voltage to rise originally, so that it crosses this threshold . . . and sets off this sequence of events called the impulse?" they go on. The answer is, various "nodes" along the neuron's axon have simply relayed this high voltage from one station to the next. So then the question is again transformed. This time they ask, "But what causes the very first impulse to occur at the very first node? Where does ***that*** voltage shift come from? What precedes the impulse?"

Well, for most neurons inside the brain—"intemeurons," meaning neurons that are fed into not by sensory input but only by other neurons —the answer is, their first node's voltage shift is provoked by the total effect of the pulses of neurotransmitters coming in from other neurons. (We could call those neurons "upstream" neurons, but that would imply, quite falsely, that the flow of neural activity in the brain follows a line in only one direction, in the manner of a river. In fact, as a rule, neural flow patterns are far from linear and make loops all over the place, quite unlike rivers.)

Thus we seem to get into a vicious circle—a chicken-and-egg type of riddle. Question: "What triggers a neural firing?" Answer: "Other neural firings!" But the real question remains unanswered: "Why those neurons, and not others? Why this vicious circle and not another neural loop in another part of the brain?" To answer this, we have to shift levels and talk about the relationship of the brain to the ideas it encodes, which then would require us to talk about how the brain encodes, or represents, its concepts about the world. Since we do not wish to theorize in this book on the details of such matters, we will talk about a related but simpler concept.

Imagine an intricately bifurcating and rejoining domino-chain network. Suppose that each domino has a little time-delayed spring underneath it that stands it up again five seconds after it has fallen. By setting up the network in various configurations, one could actually program the system of dominoes to perform calculations with numbers, exactly as one could a full-scale computer. Various pathways would carry out various parts of the calculation, and elaborate branching loops could be set up. (Note how this image is not too different, then, from that of networks of neurons in a brain.)

One could imagine a "program" trying to break the integer 641 into

the product of its prime factors. "Why isn't this particular domino ever falling down?" you might ask, pointing at one that you've been watching for a long time. An answer on one level would be "Because its predecessor never falls." But that low-level "explanation" only begs the question. What one really wants—the only satisfying answer, in fact—is an answer on the level of the concepts of the program: "It never falls because it is in a stretch of dominoes that gets activated only when a divisor is found. But 641 has no divisors—it is prime. So the reason that domino never falls has nothing to do with physics or domino chains—it is simply the fact that 641 is prime."

But have we then admitted that higher-level laws actually are responsible, and govern the system above and beyond lower-level laws? No. It is simply that an explanation that makes any ***sense*** demands higher-level concepts. The dominos certainly don't know they are part of a program, nor do they need to—any more than the keys of a piano know, or need to know, which piece you are playing. Think how strange it would be if they did! Nor do your neurons know that they are involved in thinking these thoughts right now, nor ants that they are part of the grand scheme of their colony.

There is a further-back question that might arise in your mind: "What laws, at what level, are responsible for the existence of the program and the domino chains—indeed, for the manufacturing of the dominoes at all?" To answer this and the many questions it inevitably triggers, we are sent sailing backward in time over larger and larger spans, back into all the reasons our society exists, back to the origin of life, and so on. It is more convenient to sweep these many questions under the rug and simply to leave our reason as: the primeness of 641. We prefer this kind of compact higher-level explanation, one that eliminates long views into the past and that concentrates on the present or the timeless. But if we want to trace events to their ultimate causes, we are forced into reductionistic views as described by Dawkins or the Tortoise. Indeed, ultimately we are sent back to the physicists, who will refer us to the "Big Bang" as the primordial cause of everything. This is not satisfying, however, because we want an answer at a level that appeals to concepts familiar to people—and, fortunately, nature is stratified enough that this is often possible.

We asked whether a thought can influence the course of an electron in flight. The reader could easily conjure up an image we do not have in mind—namely, of a deeply concentrating "psychic" with furrowed brow beaming his "waves of Plutonian energy" (or whatever he calls them) outwards toward an object—say a tumbling die—and influencing the way it will land. We do not believe in anything of the sort. We do not believe

that there is some as-yet undiscovered "mental magnetism" through which concepts could "reach down" and, through some sort of "semantic potential," alter the paths of particles, making them deviate from what present-day physics would predict. We are talking about something else. It is more a question of where explanatory power comes from—perhaps a question of the proper ways of using words, a question of how to reconcile everyday usage of terms like "cause" with the scientific usage of those terms. Thus, is it reasonable to explain the trajectories of particles by making references to higher-level notions such as "beliefs," "desires," and so forth? The reader may detect that we see much utility in adopting this way of speaking. Just as evolutionary biologists feel free to use "teleological shorthand" to condense their concepts down to an intuitively reasonable size, so we feel that people who study the mechanisms of thought must necessarily become conversant with ways of translating back and forth between purely reductionistic language and a sort of "holistic" language in which wholes do indeed exert a visible effect on their parts, do indeed possess "downward causality."

In physics, when a shift of point of view is made, sometimes the laws may appear to be different. Think of the amusement park ride in which people line the inner walls of a large cylinder. The cylinder starts spinning and as it does so, its floor falls away, as if a giant can opener had just opened this can from below. The people are left hanging, with their backs powerfully pressed against the wall by the so-called centrifugal force. If you were on this ride and attempted to throw a tennis ball to a friend directly across the cylinder, you would see the ball flying crazily off course, perhaps even boomeranglike returning to you! Of course, this is simply because you would move around in the same amount of time as the ball sailed (in a straight line) across the cylinder. But if you were unaware that you were in a rotating frame, you might invent a name for the strange deflecting force that makes your ball veer away from its intended destination. You might think it was some bizarre variation of gravity. This would be strongly supported by the observation that this force acted identically on any two objects with the same mass, as gravity does. Amazingly enough, this simple observation—that "fictitious forces" and gravity are easily confused—is at the heart of Einstein's great theory of general relativity. The point of this example is that a shift of frame of reference can induce a shift of perceptions and concepts—a shift in ways of perceiving causes and effects. If it is good enough for Einstein, it ought to be good enough for us!

We will not belabor the reader further with descriptions of the tricky shifts of point of view as one swings back and forth between the level of wholes and the level of their parts. We will simply introduce some catchy

terminology which may titillate the reader into thinking further about these issues. We have contrasted "reductionism" and "holism." Now you can see that "reductionism" is synonymous with "upward causality" and "holism" with "downward causality." These are concepts having to do with how events on different size-scales in ***space*** determine each other. They have counterpart notions in the ***time*** dimension: to reductionism corresponds the idea of predicting the future from the past without regard to "goals" of organisms; to holism corresponds the idea that only inanimate objects can be so predicted, but that in the case of animate objects, purposes and goals and desires and so on are essential to explain their actions. This view, often called "goal-oriented" or "teleological," could equally well be termed "goalism"—and its opposite could be termed "predictionism." Thus predictionism emerges as the temporal counterpart to reductionism, with goalism being the temporal counterpart to holism. Predictionism is the doctrine that only "upstream" events —and nothing "downstream"—need be taken into account in determining the way the present flows into the future. Goalism, its opposite, sees animate objects as moving toward goals in the future—thus it sees future events in some sense projecting causal power backward in time, or retroactively. We can call this "retroactive causality"; it is the temporal counterpart to holism's "introactive causality," where causes are seen to flow "inward" (from wholes to their parts). Put goalism and holism together, and you have—you guessed it—soulism! Put predictionism and reductionism together, and you get—mechanism.

To summarize, we can draw a little chart:

Hard scientists	Soft scientists
Reductionism	Holism
(upward causality)	(downward causality)
+	*f
Predictionism	Goalism
(upstream causality)	(downstream causality)
= Mechanism	= Soulism

Well, now that we have indulged our fancy for wordplay, let us move on. A fresh perspective is offered us by another metaphor for brain activity: that of the "thinking wind chime." Think of a complex wind

chime structured like a mobile, with glass "tinklers" dangling like leaves off branches, branches dangling from larger branches, and so on. When wind strikes the chime, many tinklers flutter and slowly the whole structure changes on all levels. It is obvious that not just the wind, but also the chime state, determines how the little glass tinklers move. Even if only one single glass tinkler were dangling, the twistedness of its string would have as much to do with how the chime would move as the wind would.

Just as people do things "of their own volition," so the chime seems to have a "will of its own." What is volition? A complicated internal configuration, established through a long history, that encodes tendencies toward certain future internal configurations and away from others. This is present in the lowliest wind chime.

But is this fair? Does a wind chime have desires? Can a wind chime think? Let's fantasize a bit, adding many features to our chime. Suppose there is a fan on a track near the chime, whose position is electronically controlled by the angle of one particular branch in the chime, and whose blades' rotational speed is controlled by the angle of another branch. Now the chime has some control over its environment, like having big hands that are guided by groups of tiny, insignificant-seeming neurons: the chime plays a larger role in determining its own future.

Let's go further and suppose that many of the branches control blowers, one blower per branch. Now when wind—natural or blower- caused—blows, a group of tinklers will shimmer, and subtly and delicately they will transmit a soft shimmer to various other portions of the chime. That in turn propagates around, gradually twisting branches, thus creating a new chime state that determines where the blowers point and how hard they blow, and this will set up more responses in the chime. Now the external wind and the internal chime state are intertwined in a very complicated way—so complicated, in fact, that it would be very hard to disentangle them conceptually from each other.

Imagine two chimes in the same room, each affecting the other by blowing small gusts of wind in the direction of the other. Who can say that it makes sense to decompose the system into two natural parts? It might be that the best way to look at the system is in terms of top-level branches, in which case there might be five or ten natural parts in each of the two chimes—or perhaps the branches a level below that are the best units to look at, in which case we might see twenty or more per chime It is all a matter of convenience. All parts interact in some sense with all others, but there might be two parts that are somewhat discernible as separate in space or in coherence of organization—certain types of shimmering might stay localized in one region, for instance—and we

could then speak of distinct "organisms." But note how the whole thing is still explicable in terms of physics.

We could now posit a mechanical hand whose motions are controlled by the angles of, say, two dozen high-level branches. These branches are of course intimately tied in with the entire chime state. We could imagine the chime state determining the hand's motions in a curious way—namely telling the hand which chess piece to pick up and move on a board. Wouldn't it be a marvelous coincidence if it always picked up a sensible piece and made a legal move? And an even more marvelous coincidence if its moves were always ***good*** moves? Hardly. If this were to happen, it would be precisely because it was ***not*** a coincidence. It would be because the chime's internal state had ***representational power.***

Once again we'll back away from trying to describe precisely how ideas could be stored in this strange shimmering structure, reminiscent of a quaking aspen. The point has been to suggest to the reader the potential delicacy, intricacy, and self-involvedness of a system that responds to external stimuli and to features at various levels of its own internal configuration.

It is well-nigh impossible to disentangle such a system's response to the outside world from its own self-involved response, for the tiniest external perturbation will trigger a myriad tiny interconnected events, and a cascade will ensue. If you think of this as the system's "perception" of input, then clearly its own state is also "perceived" in a similar way. Self-perception cannot be disentangled from perception.

The existence of a higher-level way of looking at such a system is not a foregone conclusion; that is, there is no guarantee that we could decode the chime state into a consistent set of English sentences expressing the beliefs of the system, including, for instance, the set of rules of chess (as well as how to play a good game of chess!). However, when systems like that have ***evolved*** by means of natural selection, there ***will*** be a reason that some have survived and most others failed to: meaningful internal organization allowing the system to take advantage of its environment and to control it, at least partially.

In the wind chime, the hypothetical conscious ant colony, and the brain, that organization is stratified. The levels in the wind chime corresponded to the different levels of branches dangling from other branches, with the spatial disposition of the highest branches representing the most compact and abstract summary of the global qualities of the chime state, and the disposition of the many thousands (or millions?) of quivering individual tinklers giving a totally unsummarized, unintuitive, but concrete and local description of the chime state. In the ant colony, there were ants, teams, signals at various levels, and finally the caste distribu

tion or "colony state"—again the most incisive yet abstract view of the colony. As Achilles marveled, it is so abstract that the ants themselves are never mentioned! In the brain, we just do not know how to find the high-level structures that would provide a readout in English of the beliefs stored in the brain. Or rather, we do—we just ask the brain's owner to tell us what he or she believes! But we have no way of physically determining where or how beliefs are coded.*

In our three systems, various semiautonomous subsystems exist, each of which represents a concept, and various input stimuli can awaken certain concepts, or symbols. Note that in this view there is no "inner eye" that watches all the activity and "feels" the system; instead the system's state itself represents the feelings. The legendary "little person" who would play that role would have to have yet a smaller "inner eye," after all, and that would lead to further little people and ever-tinier "inner eyes"—in short, to infinite regress of the worst and silliest kind. In this kind of system, contrariwise, the self-awareness comes from the system's intricately intertwined responses to both external and internal stimuli. This kind of pattern illustrates a general thesis: "Mind is a pattern perceived by a mind." This is perhaps circular, but it is neither vicious nor paradoxical.

The closest one could come to having a "little person" or an "inner eye" that perceives the brain's activity would be in the ***self-symbol***—a complex subsystem that is a model of the full system. But the self-symbol does not perceive by having its own repertoire of smaller symbols (including its ***own*** self-symbol—an obvious invitation to infinite regress). Rather, the self-symbol's ***joint activation*** with ordinary (nonreflexive) symbols constitutes the system's perception. Perception resides at the level of the full system, not at the level of the self-symbol. If you want to say that the self-symbol perceives something, it is only in the sense that a male moth perceives a female moth, or in the sense that your brain perceives your heart rate—at a level of microscopic intercellular chemical messages.

The last point to be made here is that the brain needs this multileveled structure because its mechanisms must be extraordinarily flexible in order to cope with an unpredictable, dynamic world. Rigid programs will go extinct rapidly. A strategy exclusively for hunting dinosaurs will be no good when it comes to hunting woolly mammoths, and much less good when it comes to tending domestic animals or commuting to work on the subway. An intelligent system must be able to reconfigure itself—to sit back, assess the situation, and regroup—in rather deep ways; such flexi

*See selection 25, "An Epistemological Nightmare," for a story featuring a machine that can outdo a person at "brain reading."

bility requires only the most abstract kinds of mechanisms to remain unchanged. A many-layered system can have programs tailored to very specific needs (e.g., programs for chess playing, woolly-mammoth hunting, and so on) at its most superficial level, and progressively more abstract programs at deeper layers, thus getting the best of both worlds. Examples of the deeper type of program would be ones for recognizing patterns; for evaluating conflicting pieces of evidence; for deciding which, among rival subsystems clamoring for attention, should get higher priority; for deciding how to label the currently perceived situation for possible retrieval on future occasions that may be similar; for deciding whether two concepts really are or are not analogous; and so on.

Further description of this kind of system would carry us deep into the philosophical and technical territory of cognitive science, and we will not go that far. Instead, we refer readers to the "Further Readings" section for discussions of the strategies of knowledge representation in humans and in programs. In particular, Aaron Sloman's book ***The Computer Revolution in Philosophy*** goes into great detail on these issues.

D.R.H.