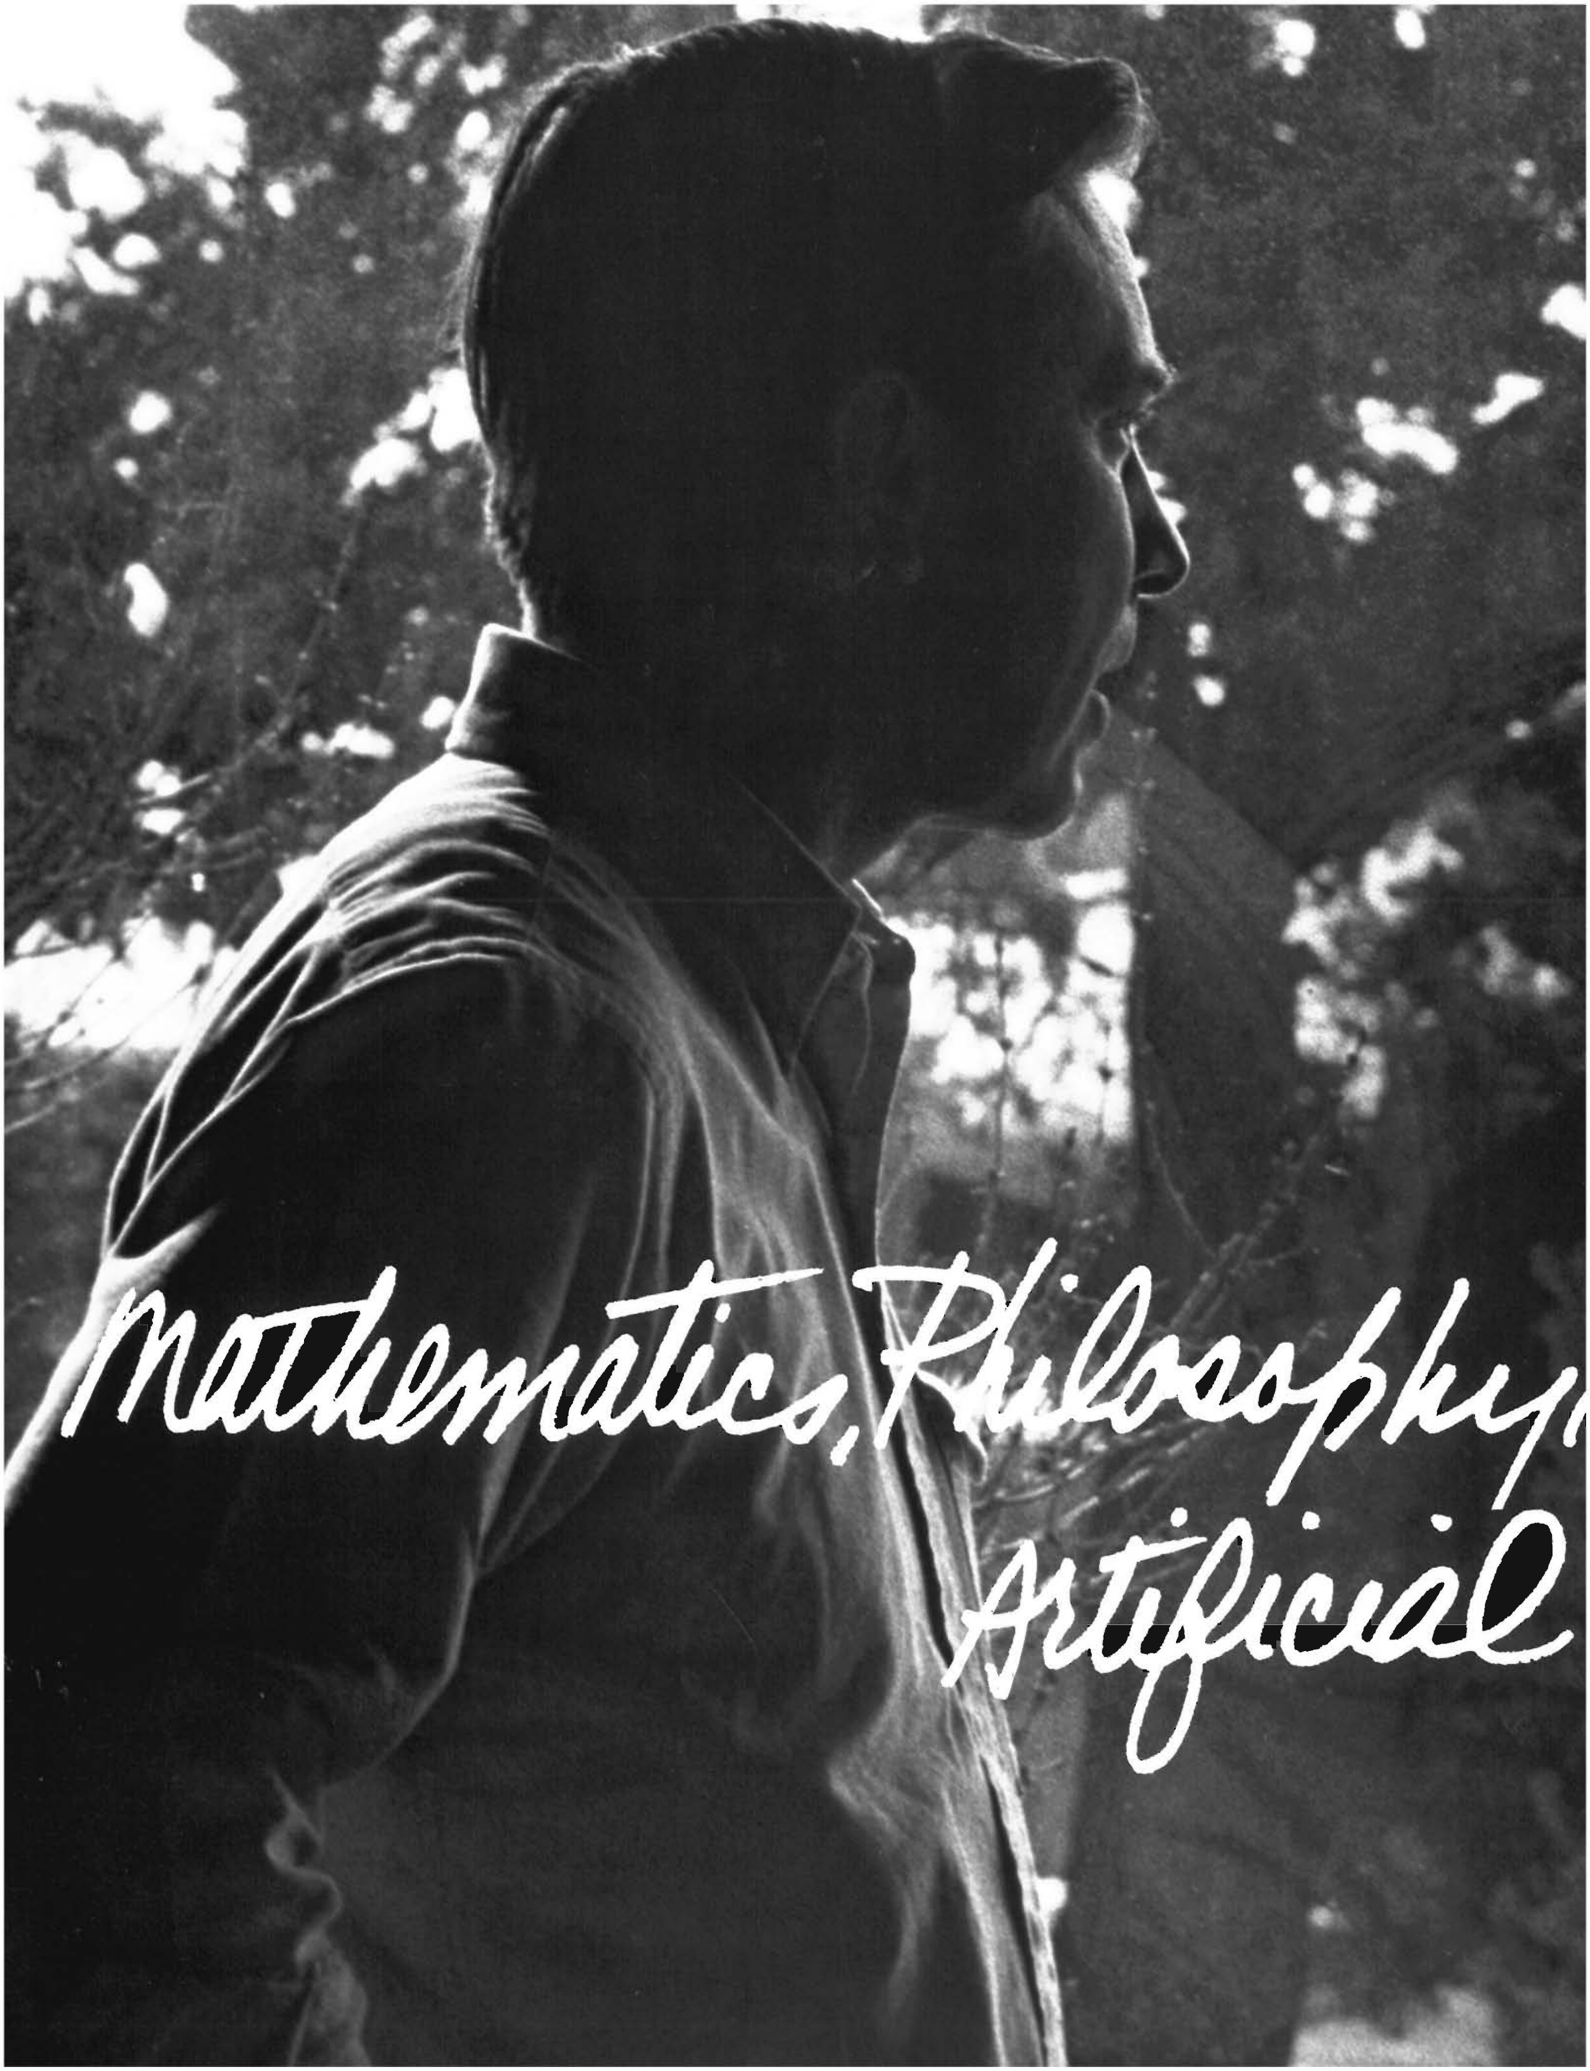
Mathematics, Philosophy,
Artificial

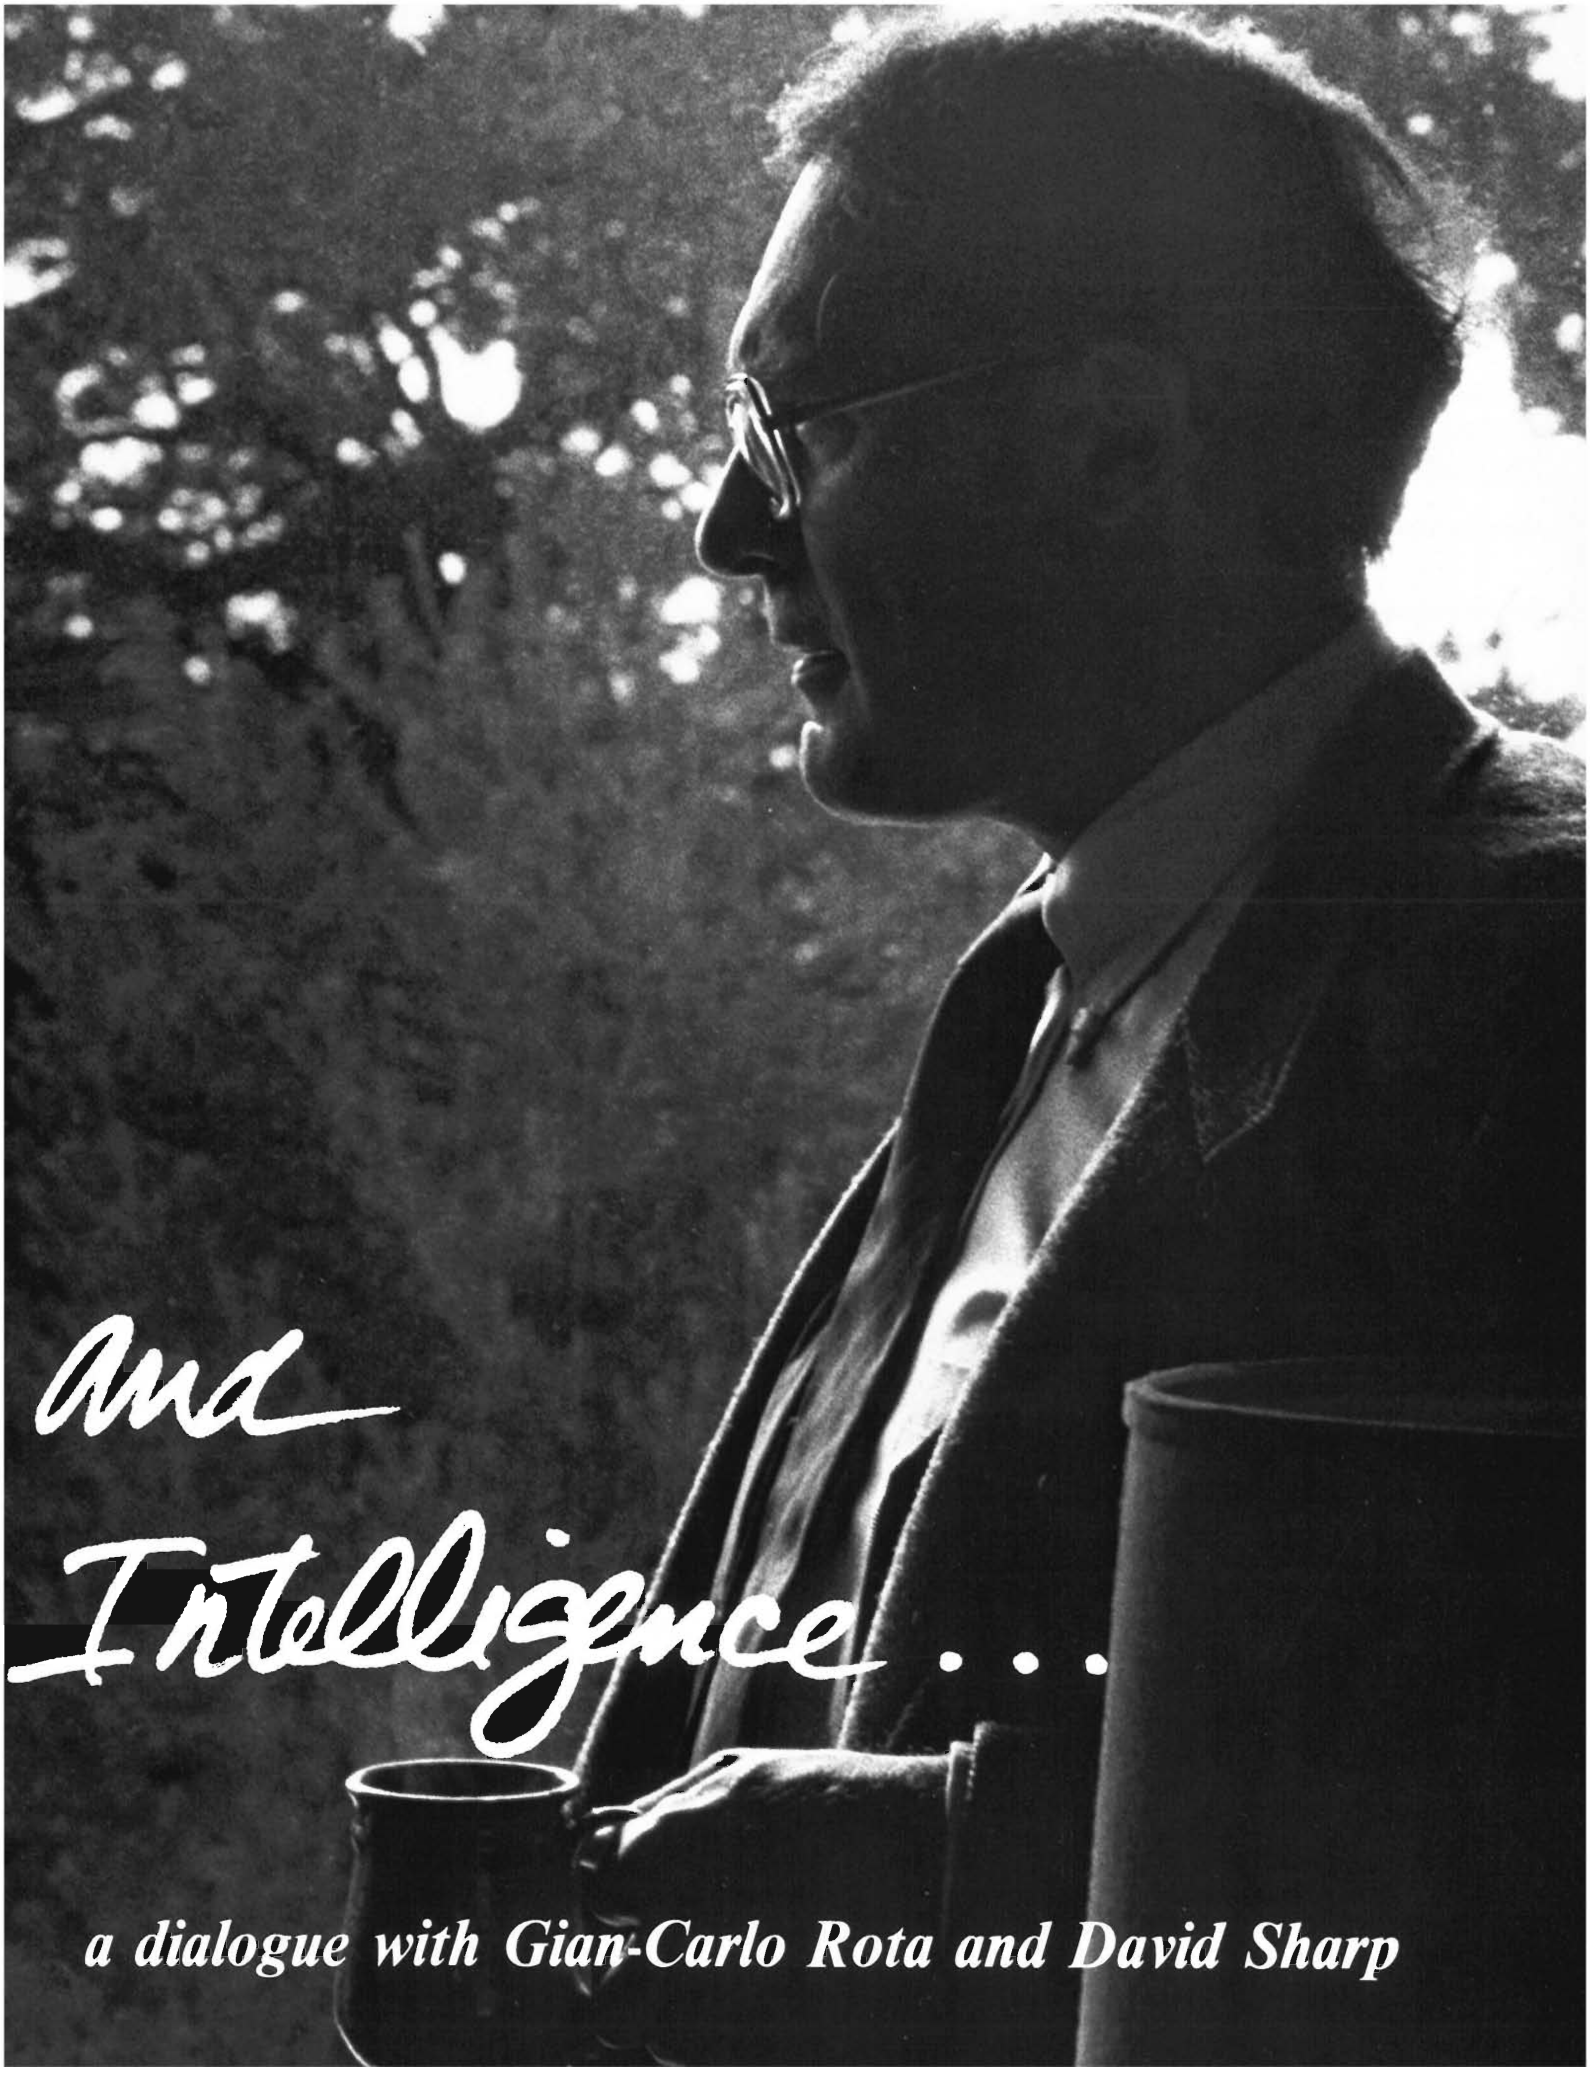

and Intelligence . . .

a dialogue with Gian-Carlo Rota and David Sharp

The Talmud and the Tao are good trainers of scientific minds.

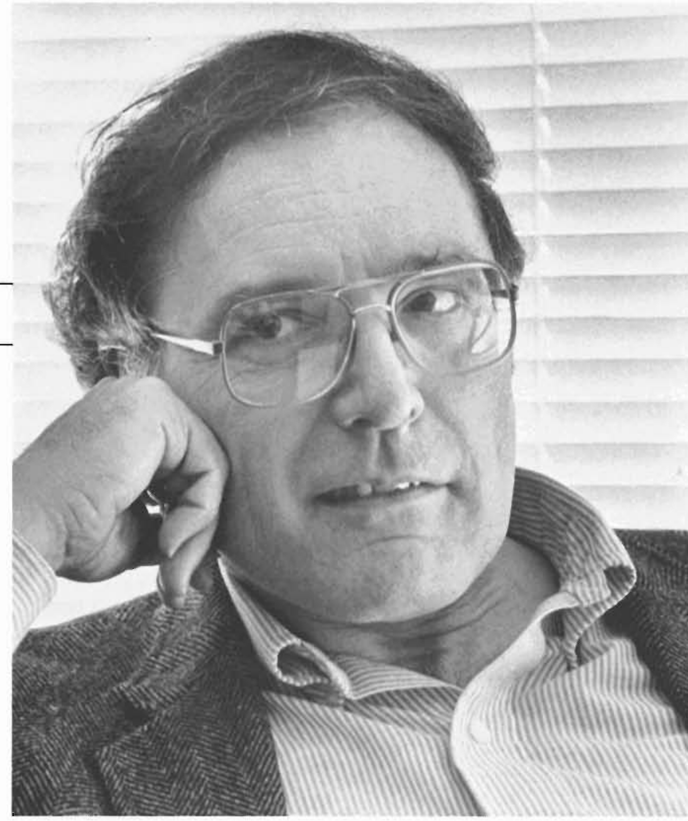

DIALOGUE

SCIENCE: What's going on in science? Does anyone have the big picture, or are people too stuck in their narrow specialties?
ROTA: I'm afraid I can talk only about mathematics. Until a few years ago, you found everywhere a strong tendency to specialization; now you see much straddling between branches of mathematics that used to be tight, separate compartments. People who can work in more than one field, and who are able to see analogies between seemingly disparate concepts, will be the ones who set the pace for the future. The younger generations are not wasting any time in picking up this hint.
SHARP: In physics, too, the trend toward synthesis is evident. The recent marriage of particle physics and cosmology is leading to an understanding of the early universe. Nonlinear dynamics now cuts across hydrodynamics, many-body theory, and plasma physics.
SCIENCE: What about biology?
SHARP: A good example of scientific interaction is neural networks. The theory was started by a neuroscientist, Warren McCulloch, and an applied mathematician, Walter Pitts, caught the eye of a pure mathematician, von Neumann, and was worked on by a mathematical logician, Steve Kleene, and eventually by an engineer, Claude Shannon. The subject borrows from neurobiology, artificial intelligence, and classical logic. It purports to offer a synthetic approach to the old and difficult scientific problem of the working of the mind.
ROTA: As you were saying, a new unit in science is being formed that remains to be named. It will include the best of theoretical computer science, neurophysiology, molecular biology, psychology, and the mathematical theory of information. It will be important to name it properly. As the Latin proverb says, "*Nomen est omen*."
SHARP: The name is a presage.
ROTA: We need to name it glamorously, to make it into a respectable profession, and to bring people together to work on it. This is beginning to happen. It is the central drama of today's science.
SCIENCE: Has it something to do with understanding the idea of intelligence?
ROTA: It has to do with muscling into processes of behavior, whether human or machine.
SHARP: Including intelligence.
ROTA: But not specifically intelligence. As is true in the early stages of any science, it's hard to tell the crackpots from the geniuses. They are mingled together, as they were in the beginning days of physics when, for example, Kepler thought he could classify the distances of the planets from the sun by the lengths of the sides of the five regular solids inscribed in a sphere. Newton himself believed in magic.
SHARP: Not to mention alchemy.
ROTA: Nevertheless, he discovered the laws of mechanics. Similarly, we could view the new science as a melting pot of good guys and bad guys, of con men and serious people. It is hard to sort them out. Sometimes the same person is both.
SCIENCE: Why do some scientists remain active, while others burn out?

Gian-Carlo Rota, survivor of seventeen years as a consultant to Los Alamos, is a leader in the field of combinatorics and a popular professor of applied mathematics and philosophy at MIT. He was born in Milan, raised, as he says, by the iron hand and given charge at the tender age of eight of his father's impressive library, a place where he developed a liking for the smell of printed matter. Erudite, worldly, and immensely generous to his friends and students, Gian-Carlo shares his penetrating insights with the accent on humor. Asked about his role at Los Alamos, he replied, "I wish I knew. I manage to snoop around and every once in a while pop into the Director's office and have a chat with him. I can be outspoken—no one is offended because next Sunday I leave anyway. By the time I return everything's forgotten and we can start all over again."

ROTA: I have a one-word answer. It is the word *Kultur*. A broad cultural background, the learning at an early age in the family of the value of things intellectual, is the main factor that keeps people from burning out. That's the reason many first-rate intellectuals come from Jewish or Chinese backgrounds, where they were exposed to intellectual values at an early age. The Talmud and the Tao are good trainers of scientific minds.
SHARP: There is more to it than that. You go into science because you like it. That liking carries you through the ups and downs of your career, through the plateaus you reach from time to time in your work. While on one of those plateaus, it is important not to be afraid to try to learn new things.
ROTA: Call it intellectual chutzpah.
SHARP: If one gets caught doing a small extrapolation of what someone else has already done, the work will be neither daring nor

The farthest reaching contribution of modern psychology . . . is the realization that intelligence is not a single monolithic faculty.

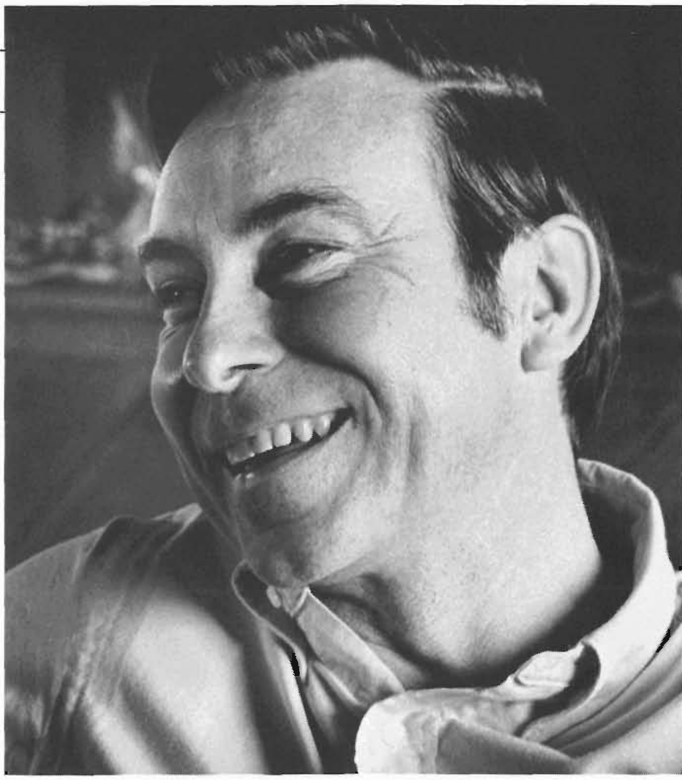

David Sharp, a Fellow at Los Alamos, shares with Gian-Carlo an undergraduate background at Princeton and a capacity for working in several fields at once. Always intense, enthusiastic, and persevering, David switches from neutral networks to Rayleigh-Taylor instabilities to quantum field theory with an agility reminiscent of his figure-skating youth. He is also working with Gian-Carlo on the theory of stochastic processes, and on weekends they moonlight on the philosophy of mathematics.

exciting. Once the excitement is gone, you can kiss good-by to creative contributions.

SCIENCE: Who are the daring scientists of our times?

SHARP: Feynman is one. Near the end of his career, what is he doing? He is not working on particle physics where he made such enormous contributions. He's working on computer science, on tiny computers based on quantum-mechanical principles; he's now helping to teach a course at Caltech on computer science.

ROTA: My Ph.D. thesis advisor, then at Yale, Jack Schwartz, is another example. He worked on functional analysis during my graduate-school days. Today he's deeply engaged in computer design and robotics—quite a switch.

SHARP: That brings to mind another heavyweight, Jim Glimm, who has made truly outstanding contributions in mathematics and mathematical physics, from quantum field theory to fluid dynamics to C*-algebras. Unfortunately, the department structure of today's universities discourages people from change.

SCIENCE: You said that a new science is emerging. Are young people being trained in it? Is something serious really happening, or is it all phony?

ROTA: Everyone agrees that we must reform our programs at the universities. But it is hard to change the sluggishness of an academic environment, and tension is building. We need imaginative new programs in mathematics; we need daring departures that straddle mathematics, physics, and biology.

SHARP: Mathematicians are concerned with the decline in the number of students in graduate programs, as well as with the decline in support for mathematics. For some decades there has been a split in the mathematics community. Some feel that mathematics is enriched by contact with nature, with physics and biology. Others take an aesthetic, self-contained view and argue that the internal logic of mathematics will dictate its development. This point of view is sometimes identified with the Bourbaki school. In the old days Hermann Weyl and David Hilbert were deeply involved in the development of the theoretical physics of their time. In our time computers have caught up with physics in providing the external stimulus for mathematicians.

ROTA: For mathematical logicians most of all. Witness the exodus of mathematical logicians out of their field into computer science. Thanks to logicians we have sophisticated programming languages and superior software. Computer science has become too important to be left to engineers. Fortunately, physicists and mathematicians are switching to computer and hardware design in great numbers, attracted by the higher pay. Professors of computer science, even those who do not know how to read and write, make double the salary of even the best mathematicians.

SCIENCE: Isn't it true that mathematicians switch to computer science when they're burning out? When do you know a mathematician is through?

ROTA: Von Neumann used to say that a mathematician is finished by the age of thirty. As he got older, he increased the age to thirty-five, then to forty, to forty-five, and soon to fifty. We've inherited from the 19th century a misleading notion that mathematicians have to do their work early or they're finished. That's not true. The kind of work a mathematician does as he grows older changes. An older mathematician will work on questions of wider scope, whereas a younger mathematician will choose to work on a single difficult problem. A variety of talents are required for the scientific community to thrive. What is perhaps the farthest reaching contribution of modern psychology to human welfare is the realization that intelligence is not a single monolithic faculty that can be measured on a linear scale. You may be smart at working out math problems, but stupid at doing everything else. The old IQ test was very effective at testing one kind of intelligence, what we might call quiz-kid smarts, but it did not test any of the other kinds.

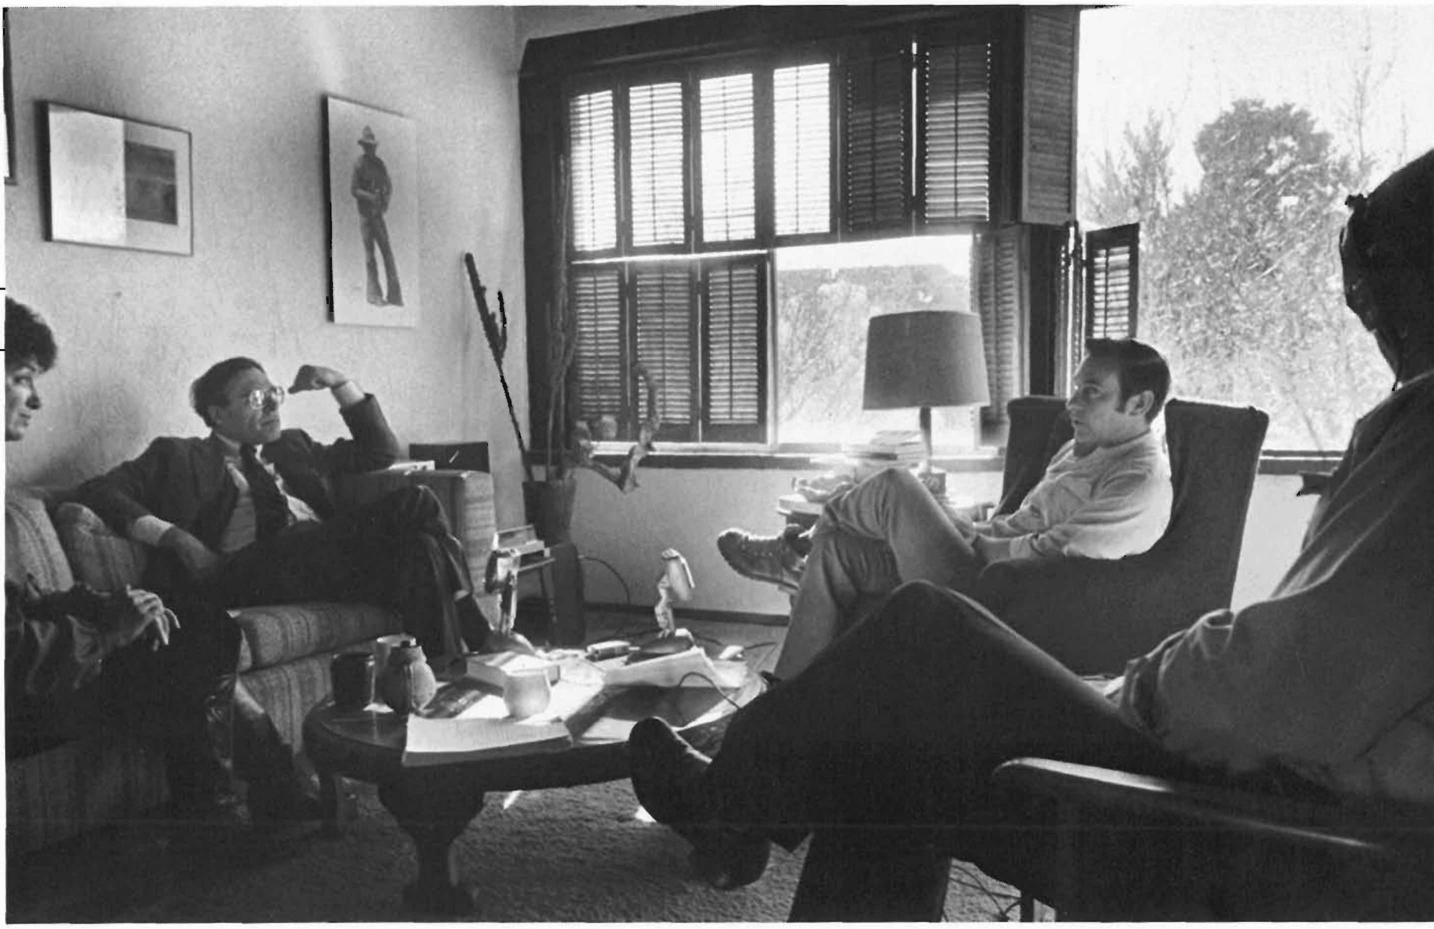

SCIENCE: Are you born a mathematician, or can you be taught to be one?
ROTA: Most people have some talent or other that could make them into reasonable mathematicians, given the motivation. But there is also such a thing as raw mathematical talent, just like there is talent in music. Unfortunately, we do not know how to measure "raw talent," though we can recognize it when we see it.
SHARP: Psychology is now beginning to discover the subtleties of learning. People are either geometric and visual, or verbal and logical, or kinesthetic. Psychological studies of math teaching have shown that different teaching techniques work for different people, depending on whether they are verbal, visual, or kinesthetic. Unfortunately, these subtleties will take a long time to seep down into grade-school teaching.
SCIENCE: Most people teach themselves anyway.
ROTA: But a teacher can be effective in discovering his students' talents and in encouraging those talents. A good teacher does not teach facts; he teaches enthusiasm, open-mindedness, and values.
SHARP: Somebody once said that teaching is not efficacious except in those few cases where it is superfluous.
ROTA: A good teacher brings out the best in his students. What young people need the most is encouragement. Left to themselves, students may not know how to decide what is worthwhile. They may drop an original idea because they think someone else must have thought of it already.
SHARP: Students need to be taught to believe in themselves and not to give up easily.
SCIENCE: Is there a conflict between letting students follow their own crackpot ideas rather than learning what someone else has already done?
ROTA: There is a ratio by which you can measure how good a mathematician is, and that is how many crackpot ideas he must have in order to have one good idea. If it's ten to one then he is a genius. For the average mathematician, it may be one hundred to one. You must have the courage to discard lots of attractive ideas. This is a feature of creative thinking that the man in the street fails to realize.
SCIENCE: And you have to try out all those ideas?
SHARP: Pretty much, but you mustn't become infatuated with the sound of your own words. You have to be ruthless in throwing out your own bad ideas. You have to constantly weed your own garden.
SCIENCE: Coming back to mathematical talent, how would you characterize it?
ROTA: An outstanding characteristic of mathematical talent is the ability to spot analogies. Another, one of the rarest, is the talent for applied mathematics, the talent for picking out of a maze of experimental data the two or three parameters that are relevant, and to discard irrelevant parameters. This is rare, because it is taught only at the shop level.
SCIENCE: How is mathematics applied?
SHARP: Most people, even some cultivated scientists, think that mathematics applies because you learn Theorem Three and then Theorem Three somehow mysteriously explains the laws of nature. That doesn't happen even in science fiction novels— it is pure fantasy. The results of mathematics are seldom directly applied; it is the definitions that are really useful. Once you learn the concept of a differential equation, you then see that in biology, say, there are differential equations all over. This you cannot see unless you take a course in abstract differential equations. Here, as everywhere, what applies is the cultural background you get from the course, not the specific theorems taught in the subject. If you want to learn to speak

There is nothing deadlier for a mathematician than to be . . . instructed to lay golden eggs.

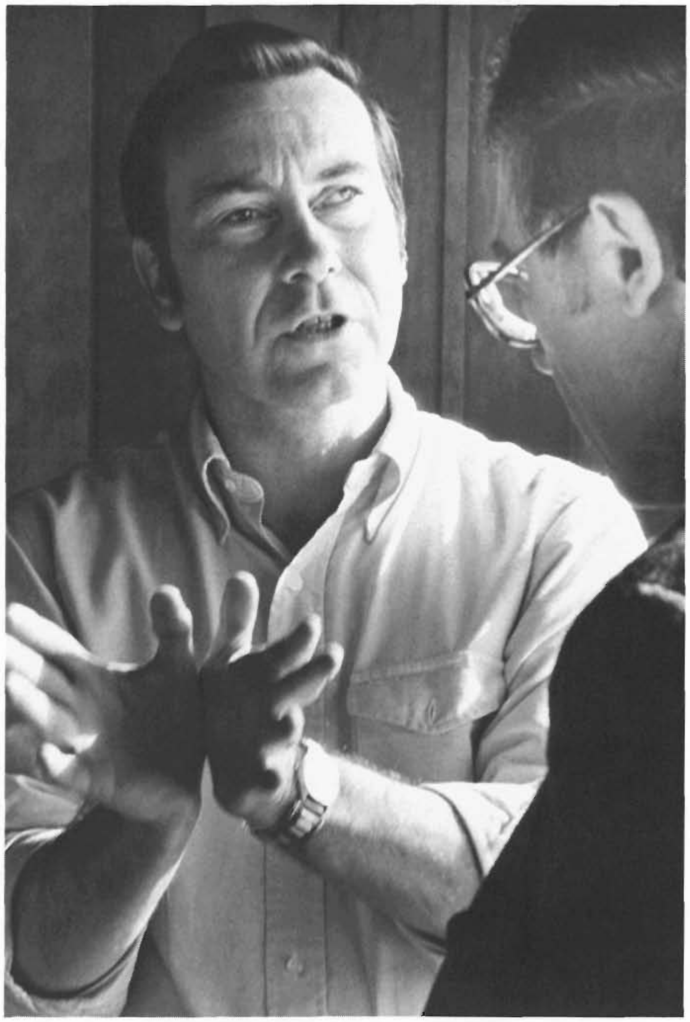

French, you have to live the life of France, not just memorize thousands of words. Similarly, if you want to apply mathematics, you have to live the life of differential equations. When you live this life, you can then go back to your molecular biology with a new set of eyes that will see things that you couldn't otherwise see.

ROTA: I once naively entertained the thought that biologists could tell me what their mathematical problems were so that I could think of solving them. That's ridiculous. Biologists seldom have the mathematical view that is required to spot problems in the mathematics of biology that may be right there staring at them. A biologist will go on doing experiments all his life and never see anything deeper than binomial coefficients. It is not that the problems aren't there; rather, biologists by and large don't have the view that comes only from a solid education in mathematics.

SCIENCE: Once you have isolated a problem, how do you go about solving it?

SHARP: No mathematical problem is ever solved directly. You must manage to look at the surroundings, at the company it keeps. Eventually, with the assistance of the Holy Ghost, you might be able to see through your problem.

ROTA: All creative mathematicians and physicists I know work like this. They have constantly in their minds a list of a dozen of their pet problems. Anything they hear, they automatically test against their dozen problems. Every once in a while, there's a hit, and then people cry out, "He's a genius! How did he know that this was the right trick to apply to Problem Three?" Little do they know that the guy has been testing for years everything he hears against Problem Three. That's the way a lot of discoveries are made.

SCIENCE: What makes a creative mathematician?

SHARP: Rule One: Don't ask him to be creative. There is nothing deadlier for a mathematician than to be placed into a beautiful office and be instructed to lay golden eggs. Creativity is never directly sought after. It comes indirectly. It comes while you are complaining about too much routine work, and so you decide to spend half an hour on your secret pet project. Those are the occasions when you get good ideas. Or while getting ready to lecture your undergraduates, you realize that the textbook for the course is lousy, and that the subject has never before been properly explained. While you work on explaining some old material, lo and behold, you get a great new idea.

ROTA: Creativity is a bad word. Unfortunately, we must leave it in the books because the people in power believe in it with sanctimonious credulity. It is a dangerous and misleading word.

SCIENCE: We recently sponsored a conference on creativity at Los Alamos.

ROTA: Look at the list of participants. It raises your eyebrows. You cannot bunch together creativity in one field and creativity in another. It's like matching producers of shoes with producers of meat loaf, because they're both producers. It is an error of logic. A friend of mine, a well-known painter, was looking at a copy of a painting by Velázquez. I watched her reactions. She started by saying, "How funny, this stroke is going down! Normally, we brush this way, but he brushed it that way." Then, "This is a combination of colors I've never seen," and so on. She said nothing about Velázquez's creativity. It is demoralizing to children to hold up Einstein or Beethoven as examples of creativity to be imitated. The idea of genius, elaborated by German romantics, is destructive; it is a flight into fantasy. There is reason to believe we've killed classical music because of that idea. People think that they will be either geniuses like Beethoven or nothing. But look at the Baroque Age—there were hundreds of little Italians who wrote good music and didn't give a hoot about being creative.

SCIENCE: Are there fashions in mathematics as there are in the arts?

SHARP: Today, mathematics is returning to the 19th century, to concrete computations, after seventy years of very abstract

Philosophers are needed today more than ever to tell the AI engineers some unpleasant truths.

mathematics. The latest fashion is 19th-century mathematics. Some of the best work in mathematical physics is based on a constructive approach.

SCIENCE: What does constructive mean?

SHARP: Instead of proving an abstract existence theorem, you produce an algorithm that delivers the solution. It's a powerful methodology.

ROTA: Concreteness is the word of the day. Now that we have learned to be abstract we can afford again to be concrete. Today's mathematics is more concrete than the mathematics of twenty or thirty years ago.

SHARP: In the fifties at Princeton you couldn't hold your head up if you weren't working on algebraic topology. Combinatorics had never been heard of, except possibly by a couple of dazed statisticians. Now combinatorics is a flourishing field. Gian-Carlo, what gave birth to the field of combinatorics?

ROTA: The time was ripe for it. Combinatorics is an honest subject. No adèles, no sigma-algebras. You count balls in a box, and you either have the right number or you haven't. You get the feeling that the result you have discovered is forever, because it's concrete. Other branches of mathematics are not so clear-cut. Functional analysis of infinite-dimensional spaces is never fully convincing; you don't get a feeling of having done an honest day's work. Don't get the wrong idea—combinatorics is not just putting balls into boxes. Counting finite sets can be a highbrow undertaking, with sophisticated techniques.

SCIENCE: What about the practical side of combinatorics? Is it as important for computer science as people are saying?

ROTA: The area of combinatorics that has found substantial application in computer science is the invention of efficient algorithms. Algorithms are instructions for performing a task. Even the fastest computers need good algorithms. A telephone company that succeeds in implementing an even slightly more efficient algorithm in, say, its switching network may well save a billion dollars. The payoffs are staggering. Today much effort in combinatorics is going into developing a theory of algorithms that will tell the optimal speed at which problems can be worked out. Some problems, unfortunately, cannot be worked out at any reasonable speed. We want to be able to tell when a problem can be solved efficiently, or whether it must be reformulated.

SCIENCE: What sort of problems?

ROTA: Take sorting. The combinatorial problem is the following. Given a sequence of numbers, how can you rearrange them in increasing order with a minimum number of transpositions? After years of research, we now know the most efficient algorithm for rearranging numbers in increasing order. We know that we have achieved the maximum speed.

SCIENCE: What kind of a proof can one give for something like that?

ROTA: The hard part comes after someone thinks he has found the fastest algorithm and wants to prove that it is the best possible. At the beginning someone will prove that N numbers can be rearranged at a speed, say, of N^2. Then someone else will modify the procedure and show that the same task can be performed with a speed of $N \log N$, and so on, until finally someone will cleverly prove that no one will ever do any better. That is the difficult part. In the theory of algorithms, one of the unfortunate turn of events was the discovery of NP-complete problems, namely, problems that can only be solved by algorithms that grow exponentially and therefore cannot be worked out on computers in a reasonable time. How to get around NP-complete problems is a frontier of combinatorics on which the best people are working.

SCIENCE: Working on what?

ROTA: On how to change an NP-complete problem into one for which a workable computer algorithm can be invented, by leaving out irrelevant cases or by taking a carefully chosen subset of the problem.

SHARP: Some problems are NP-complete only if one asks for the exact solution. If you are content with an accuracy of 2 or 3 percent, then an NP-complete problem may become tractable.

SCIENCE: Let's turn to highbrow combinatorics.

ROTA: Much combinatorics of our day came out of an extraordinary coincidence. Disparate problems in combinatorics, ranging from problems in statistical mechanics to the problem of coloring a map, seem to bear no common features. However, they do have at least one common feature: their solution can be reduced to the problem of finding the roots of some polynomial or analytic function. The minimum number of colors required to properly color a map is given

Research is not so much discovering something new as becoming aware of prejudices that stop us from seeing what is in front of us.

by the roots of a polynomial, called the chromatic polynomial; its value at N tells you in how many ways you can color the map with N colors. Similarly, the singularities of some complicated analytic function tell you the temperature at which a phase transition occurs in matter. The great insight, which is a long way from being understood, was to realize that the roots of the polynomials and analytic functions arising in a lot of combinatorial problems are the Betti numbers of certain surfaces related to the problem. Roughly speaking, the Betti numbers of a surface describe the number of different ways you can go around it. We are now trying to understand how this extraordinary coincidence comes about. If we do, we will have found a notable unification in mathematics.

SHARP: The ultimate motivation for these developments was the Riemann hypothesis, which remains unproved. When this hypothesis is proved, it will give the best information about the distribution of prime numbers. Remarkably, this information is also coded in the zeros of an analytic function. The Weil conjectures set up an analogous function to the Riemann zeta function for a simpler case.

SCIENCE: Gian-Carlo, tell us your contribution to combinatorics.

ROTA: The one contribution of mine that I hope will be remembered has consisted in just pointing out that all sorts of problems of combinatorics can be viewed as problems of location of the zeros of certain polynomials and in giving these zeros a combinatorial interpretation. This is now called the critical problem. Over the years people have added examples of critical problems, though we still haven't gotten any idea of how to solve it in general. I'd like to see someone get such an idea before I die. The four-color conjecture—that with only four colors you can color every planar map so that no two adjacent regions have the same color—is one of these critical problems.

SHARP: I thought that had been settled by a computer proof.

ROTA: Not really. What we want is a rational proof. It doesn't help to have a brutally numerical answer spewed out by a computer. A problem is interesting only when it leads to ideas; nobody solves problems for their own sake, not even chess problems. You solve a problem because you know that by solving the problem you may be led to see new ideas that will be of independent interest. A mathematical proof should not only be correct, but insightful. Although, as Erdös says, nobody gets blamed if his first proof is messy.

SCIENCE: Will all these abstract ravings have some impact on how we view the world?

ROTA: It's a domino effect. You start with an abstract idea, and pretty soon it turns our world upside down. Leonardo da Vinci said, "Theory is the captain and application is the soldier." That's the practical side. If we take a deeper look, we see that nature imitates mathematics.

SHARP: Mathematicians look for relationships between fields of mathematics that hitherto were thought to be unrelated.

ROTA: Mathematics is the study of analogies between analogies. All science is. Scientists always want to show that things that don't look alike are really the same. That's one of their innermost Freudian motivations. In fact, that's what we mean by understanding.

SHARP: You often hear that the purpose of a scientific theory is to predict. That's not correct. The purpose is understanding. Prediction is one way to test whether our understanding is correct. Simplicity, scope, and beauty are as important as prediction in judging whether a theory leads to understanding.

ROTA: May I phrase what you just said in philosophical terms? Science turns paradoxes into trivialities.

SHARP: Gian-Carlo, there's interest, activity and, most of all, talk about artificial intelligence today. You've been following these developments. What is your candid opinion of AI?

ROTA: There is an old *New Yorker* cartoon of a wine factory in California. The director says to some visitors to the factory, "We

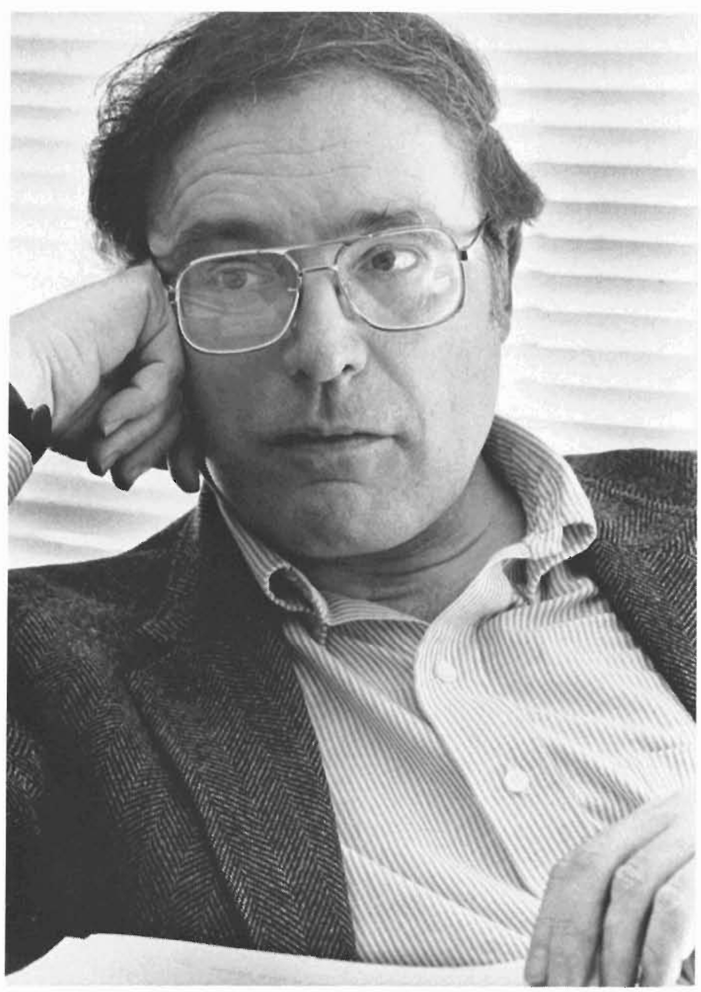

Any description of vision that omits the function of purpose and expectation will be ineffective for purposes of AI.

used to employ humans to press wine, but now we use up-to-date methods." In the background, you see a huge mechanical foot going up and down. To my mind that's the present state of the art in AI.

SHARP: Remind us of the objectives of AI.

ROTA: The accepted objective of AI is to build machines that perform tasks now performed by humans. Do not confuse AI with the more conservative computer science. Computer science and artificial intelligence are now distinct fields, and practitioners of either are at loggerheads with practitioners of the other.

SHARP: So AI research attempts to simulate human behavior. What progress has been made?

ROTA: We had to swallow a bitter pill and realize that tasks that were at first thought to be easy to simulate by computer turned out to be hard. Whenever a large memory is required, a computer will perform better than a human being, but any task that requires the slightest bit of recognition is awfully hard for a computer. Try to build a computer that recognizes when a man is wearing a uniform. Cognitive psychologists, computer scientists, specialists in artificial intelligence, and neurophysiologists are now figuring out a new breakdown of the basic components of human intelligence. Their research is beginning to reveal that the basic talents are not at all what we always thought they ought to be. The building blocks of the process of perception are turning out to be completely unexpected. Their discovery is bringing about an enormous advance in our thinking, such as has not happened since Plato.

SHARP: The enemy, here as elsewhere, is wishful thinking; it is the strangling power of prejudice. Much research in AI today consists in making scientists aware of unverbalized prejudices about thinking and speaking processes.

SCIENCE: Can you comment about the expert systems?

SHARP: The excessive claims made about these systems are creating a delusion.

SCIENCE: What claims?

SHARP: For example, the claim that expert systems will do away with physicians because computers will better diagnose disease.

ROTA: The first expert system to gain wide acceptance, in my opinion, will be the one designed for lawyers to look up cases. An expert system is basically the exploitation of the concept of a questionnaire, brought to its ultimate and gory conclusion by the computer.

SHARP: Let's consider how an expert system for medical diagnosis is constructed. First you sit down with a bunch of doctors and get them to tell you how they arrive at their diagnoses. You find that the better the doctor, the more difficult it is to characterize his diagnoses with an algorithm. But a basic element of their trade is collecting symptoms and looking up the possible diseases associated with those symptoms. If the symptoms don't uniquely pinpoint one disease, the next step is to get more information from the patient. The combinatorial model for this process is a tree-search algorithm supplemented by a set of rules for producing if-then statements. Although an expert system constructed in this way works pretty well, it has a built-in flaw because a patient can seldom cough up a complete and accurate list of his symptoms. A good diagnostician has a feeling for a patient's condition that is difficult to implement algorithmically.

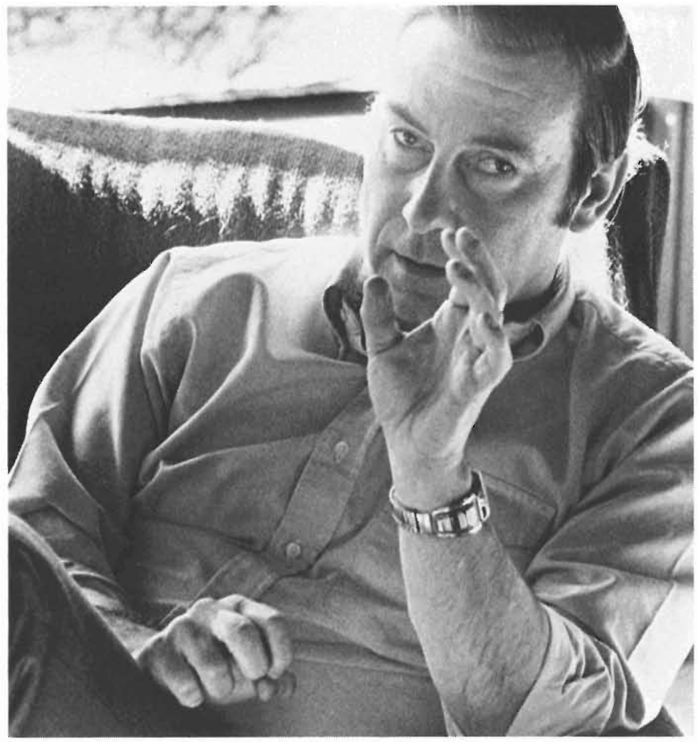

SCIENCE: Is that what you call context?

SHARP: Yes, human beings have the ability to make use of context-dependent features. We do not know at present how to program context-dependent behavior in computers. The issue of context came up with a vengeance in computer programs for chess. In the fifties computer scientists thought that they could write powerful chess-playing programs. They coded text-book openings, and they wrote programs that would look ahead a couple of moves. The results were mediocre. The programs never learned from experience, and once a human opponent found out their weaknesses, he could consistently beat the computer. Now, faster computers look six or seven moves ahead, and programs play almost at the master level. But they still don't learn, and a very good player can eventually beat them. The issue here is context. A chess player has a feel for a strong position. This contextual feel has not been implemented in chess programs.

SCIENCE: Might it be that the algorithm will be too complicated?

SHARP: Worse. We don't know the principles. In physics we know the principles, and therefore we can write down equations that describe complicated situations, sometimes even too complicated to compute. For instance, we believe that the Navier-Stokes equations

The element of anticipation is essential to all grasp of meaning.

describe turbulence, even though we're not able to compute turbulent flow. In the problem of context-dependence, we are still missing a clear formulation of the basic principles.
ROTA: That's where philosophy comes in.
SHARP: That's radical!
SCIENCE: What's radical?
SHARP: To say that philosophy should become an active partner in the creation of the new science of artificial intelligence. Usually philosophy is a Johnny-come-lately that gives perspective on a developed science. It is usually not a leader.
ROTA: Philosophers are needed today more than ever to tell the AI engineers some unpleasant truths. The philosopher's role has always been that of stating facts that may have been on everybody's mind but that no one dared state clearly. Eventually, engineers will reluctantly acknowledge that what the philosopher says is the truth, but they will then get rid of the philosopher.

Let me give you an example. In the early fifties engineers attempted to build a machine that translated from one language to another. They fell flat on their faces because they had unclear ideas on language. A good philosopher would have said, "You must begin by realizing that language is not what you thought it to be. You must bring out your unverbalized prejudices and observe language objectively without a screen of preconceived ideas on what it ought to be like." Research is sometimes not so much discovering something new as becoming aware of prejudices that stop us from seeing what is in front of us. For example, a naive view of words states that, by and large, they have a fixed meaning. Contemporary philosophy stresses instead the variety of possible contextual senses. The problem of meaning is the problem of describing the nature of the interaction between the inherited meaning of a word and the variable contextual senses it may have. For example, when someone utters a sentence, you can understand it, because of your anticipation of what comes next. This element of anticipation is essential in all grasp of meaning. It's easy to write poems about it, but try to write down the formal rules! This is precisely the task contemporary cognitive philosophy has set itself.
SHARP: The beginnings of this formalization can be traced to Chomsky and his context-free grammars. The almost mechanical rules of these grammars turn out to capture more of the structure of natural language than anyone thought possible. Formal grammars were effective for the inventing of computer languages. Eventually, they were enriched with context-sensitive grammars. But language depends on context in a way that we don't know how to express with formal grammars. We need a new idea that is yet to come.
ROTA: The problem of context-dependence is not limited to language. It has to do with sense-making generally. Consider the following example. You are at the airport. A gate opens and people come out. What are they doing? They are arriving. Think of a robot that can tell the act of walking from the act of arriving. A child can tell, but not a machine. The difficulty is that the act of arriving is purely context-dependent. The same difficulty occurs in all AI problems. Here are a few other examples that display the difficulty of formalizing context-dependence. How can you tell birds flying from birds migrating? Or someone goes to Oxford, visits the colleges, is shown the classes, and has dinner at High Table. Then at the end of the day, he asks his guide, "I've seen all these wonderful things, but where is the university?" Why is it difficult to tell where the university is? Because a university is a contextual construct. Phenomena that we believed to be physical are revealed to be actually contextual.
SCIENCE: Are you saying that all sense is context-dependent?
SHARP: Meaning is inextricably context-dependent. Context-dependence displays a variety of different layers of description that seem dependent on each other, and yet cannot be reduced to one another. For example, one can give a detailed neurophysiological description of the brain, but that does not describe the mind, although the mind depends on the brain's physiology to function. The reductionist mistake is to think that context A is reducible to context

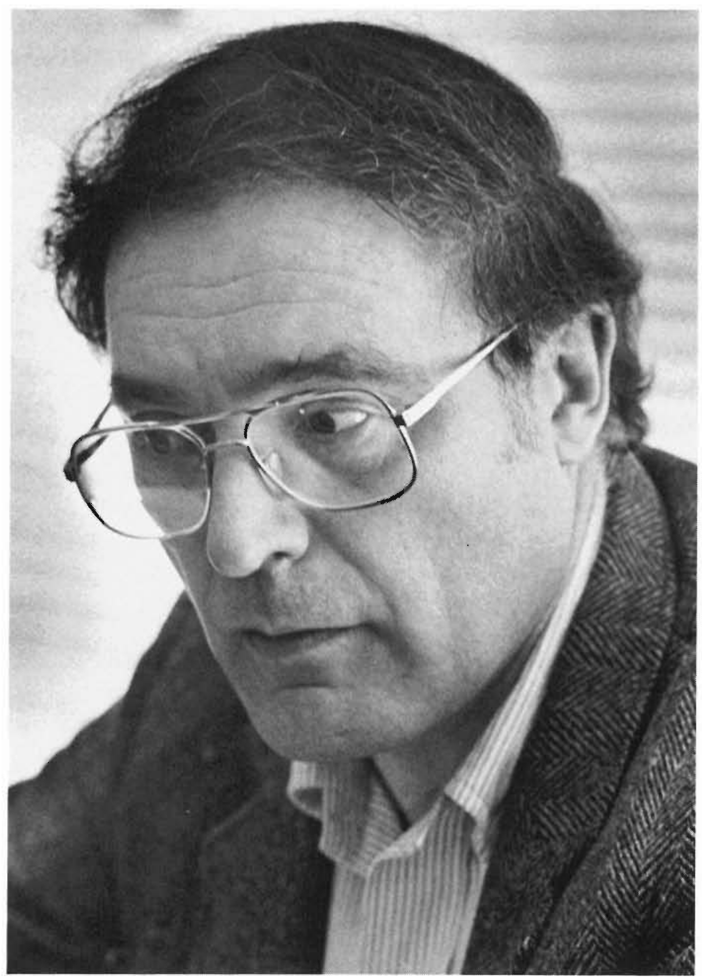

When you search deeply into . . . the Western mind, you discover the craving that all things should be reduced to one.

B, just because A cannot exist without B. Here is a naive example of a reductionist explanation. Johnny, having been asked to explain his presence in the girls' dormitory after curfew, answered, "I obeyed Newton's equations of motion." A neurophysiological description of the mind is a similar reductionist mistake.

ROTA: The laws of biology ultimately depend on the laws of physics. But from a knowledge of the laws of physics, you cannot infer the behavior of living organisms. Let me give a simple example of context-dependence, due to Ulam. You want to define a key. No amount of staring will be able to identify this object as a key unless you already have an unverbalized grasp of the contextual function of keys.

SHARP: No molecular analysis of this piece of metal will lead you to grasp the contextual function of keys and locks.

SCIENCE: Can you break down the notion of context into its fundamental elements?

ROTA: The act of perceiving a key requires an implicit background and an implicit foreground. These two elements are part and parcel of any perception. If you want to build a machine that sees, then this machine must have, programmed into it, some sense of purpose. Every act of seeing is inextricably bound with an unexpressed foreground of anticipation.

SHARP: We recognize visual scenes by virtue of our expectations of seeing certain events depending on particular contexts. The notion of expectation is the most fundamental in solving the problems of pattern recognition. If we are to build machines that see, we must find a way to encode in a machine the expectation of certain patterns. Without built-in expectation, there is no recognition. Any description of vision that omits the function of purpose and expectation will be ineffective for purposes of AI.

SCIENCE: That seems like a very tall order—to build a machine that has expectations!

ROTA: It hasn't been done, but it may not be as mysterious as it sounds. One cannot perceive anything, whether it is seeing the blue sky or grasping another person's mood, outside of a context where what is perceived plays a role as part of an organized project. What you perceive when you perceive is a function, not a thing. Being a key, a cup, a book are functions, not things. You may think that this is obvious, but people who are trying to build machines that see are still to some extent wedded to the old theory of perception, the one that pretended that the act of perceiving an object, a key for instance, was some sort of comparison with a little key inside the brain. The mortal blow was dealt to this simplistic theory by AI. As soon as scientists conceived the project of building a machine that perceives, they realized that the little-picture-in-the-brain theory did not make sense, and they became aware that an essential component of all perception is an act of choice that must to some extent remain arbitrary. Such a choice is not determined by physical data, but by your expectations of what the function you are perceiving is meant for.

SHARP: What you are saying is that all sense is functional, and function is the interplay among expectation, physical presence, and purposefulness.

SCIENCE: How does this philosophical sermon affect research in AI?

ROTA: Let's take a simple example, the act of following a rule. If you reflect on the word "following," with the idea of simulating it on a machine, you realize that you do not know the functional sense of "follow." You try various thought experiments: A follows B in a dance, or a blind man touches someone who sees and follows him. Do these examples lead us to an understanding of following in general? We are not able to pin down the meaning of the word "following" in a general framework. This problem keeps coming up in AI. AI is relying upon the received meanings of words. This is valueless for scientific purposes, as if we were trying to discover chemistry by taking earth, air, fire, and water as our basic concepts. Nowadays, the need to write computer programs that work forces us to perform investigations in philosophy that philosophers were formerly loath to do. For example, I now read the word "horse" on this page. What's really there is some ink. The word "horse" is, strictly speaking, not on the page. It's intentional. How is it that I decide to see the word "horse" rather than some black dots on a white sheet? At present, we are barely beginning to understand, thanks to the possibility of computer simulation that forces us to face these philosophical puzzles.

SCIENCE: What do you mean by "intentional," and how has the computer helped you understand this notion?

ROTA: Intentionality is the key word used to denote the new theory of perception. When you look at a printed page and see the word "horse" instead of seeing a collection of meaningless ink marks, you are selecting one function of what you see against another. This selectivity is the basic component of perception. It is now realized

Einstein was the last genius of oneness.

that intentionality is the notion to be formalized in AI. The practical needs of AI act as a powerful stimulus to this formalization.

SCIENCE: It looks like AI is helping us replace the old philosophical subject/object question with the notion of intentionality. So what?

SHARP: You see, now that we begin to see the importance of intentionality, we will analyze it in limited and rigorous environments. Computer programs performing limited cognitive functions will help us to make a science out of AI. Progress in science relies on using simple cases to test basic principles of wide applicability. To give you an analogy, solid-state physics is based on Coulomb's law of electrical interaction between charges. Coulomb's law itself is simple, but the law manifests itself in condensed matter in extraordinarily complicated ways. However, Coulomb's law operates in a simple fashion in the hydrogen atom. Similarly, by considering limited instances of perceptual problems, by developing computer programs that can solve simple problems, we are beginning to have examples that will eventually reveal the basic principles of perception. We can now test our ideas on relatively simple computer examples without grappling with the full complexity of neurobiology or human experience.

If you want to build an airplane, it might seem strange to begin with the study of the flight of birds. Boeing Aircraft won't gain much from numerical studies of feathers. Nevertheless, there are ideas common both to the flight of birds and to the flight of aircraft. These common ideas are the laws of aerodynamics. Similarly, finding the common laws for both human and machine perception is an exciting task of our time for which neurophysiology and AI are supporting each other.

ROTA: In both contexts the problem is to understand how sense-making arises out of staring at some physical data that, taken in isolation, are devoid of any sense. This is the major problem of philosophy. How does this thing called sense happen? How can you make sense while physically looking at an object or hearing a spoken word?

SHARP: In formal systems we haven't found a set of syntactic rules strong enough to do away with the need of a previous tacit knowledge of semantic sense.

ROTA: Dave is putting his finger on the problem of the inadequacy of formal logic to deal with contextual matters. Present-day logic cannot even distinguish between "and" and "but." Present-day logic claims that "and" and "but" have the same meaning, but anybody who talks uses "and" and "but" in a quite different sense.

SHARP: The actual meaning of "and" and "but" is not captured at present by their syntactic description.

ROTA: The notion of axiomatic description that we have today is inadequate to render this difference. All axiomatic descriptions of "but" have failed, not to speak of those of "nevertheless."

SHARP: Or "meanwhile."

ROTA: The language we speak is at odds with logic. It used to be thought that formal logic is a rendering of our reasoning process, but after the AI experience we have realized that formal logic is just another flight from reality.

SHARP: Formalization hasn't made a dent in our understanding of context. Someday maybe, the notion of context will be incorporated in a greatly extended logic that will delve deeper into the foundations of our thinking.

ROTA: Our ideas of foundations are changing. What we mean by foundations today is quite different from what Hilbert or Russell or Whitehead meant. Old Bertrand Russell would take it for granted that there is such a thing as foundations in mathematics and that such foundations are needed. "Mathematics shall be derived from its foundations," ruled Bertie; it must follow from a simple and consistent set of axioms.

SHARP: That's an idea with a long and tortuous history. Such a simple set of axioms has been found for simple theories like the first-order predicate calculus but not for richer branches of mathematics, such as set theory.

ROTA: There is an old anecdote about Frank Ramsey asking Wittgenstein in Cambridge sometime in the late twenties, "Look at the bridge over there. Now suppose a contradiction were to be found in set theory. Do you think that bridge would fall?"

SHARP: Gödel's incompleteness theorem did away with a lot of those pretensions of logicians, by showing that in some systems there are true statements that can neither be proven nor disproven. This was a blow to Hilbert's vision of mathematics following from a

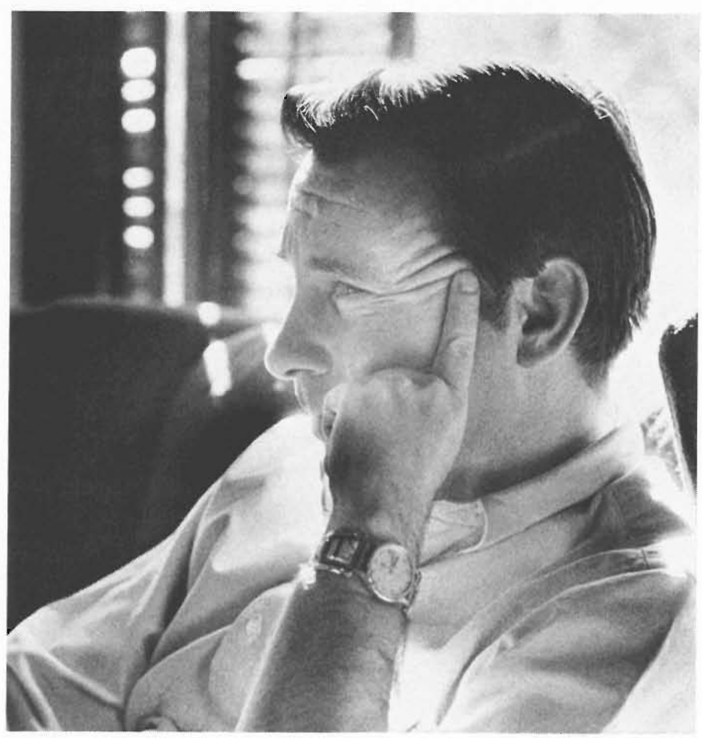

Truth is not to be held onto like a teddy bear.

simple set of consistent axioms. Gödel's theorem is upsetting because we want definitive truth, with a capital T.
ROTA: We must first look for the unstated wishes that we have in the back of our minds when we ask for foundations of mathematics or science. When you search deeply into the cravings of the Western mind, you discover the craving that all things should be reduced to one, that the laws of nature should all be consequences of one simple one, that all principles should eventually be reduced to one principle. It's a great Jewish idea. One God, one this, one that, one everything. We want foundations because we want oneness.
SHARP: Isn't this another instance of reductionism?
ROTA: This craving for reducing all physical laws to one law may be a delusion. Einstein was the last genius of oneness. Maybe he is right. But more probably, we will have to get used to several sets of laws of nature, existing together and irreducible to one another. The laws that describe living systems, if any, will not be reduced to the laws of physics or to the laws of cognitive behavior, if any.
SHARP: Physics is based on the paradigm that one analyzes physical processes in terms of the concepts of space, time, energy, velocity, and so on, not by metaphysical intervention. In the 17th century the understanding of what were the right questions to ask caused a fundamental reorientation in the way we looked at the world. Nowadays, in artificial intelligence, one needs a similar reorientation, one that points to the right questions to ask. The questions will be suggested by an enlightened analysis of behavior.
SCIENCE: How would you summarize the direction of present-day philosophy in relation to AI?
ROTA: We need to understand how sense-making arises out of staring at physical data that by themselves are meaningless. How can you get sense by merely looking? The philosopher's role is to tell the AI specialist that he doesn't know what he's talking about, to put it bluntly.
SHARP: From 1650 well into the 20th century there was a hue and cry about individual sciences dropping out of philosophy and becoming independent—first physics, then biology, and now AI. Some of these sciences might well gain from an excursion on the shores of philosophy.
ROTA: There is danger that a new profession will come into existence, the AI specialist. People who work with computers will have to hire an AI expert, like they keep a physician, a lawyer. It will be a self-perpetuating profession lacking an adequate scientific base. These people might come to control a great deal of power under the aegis of the computer. Beware!
SHARP: You see, that's what philosophers are for. ■

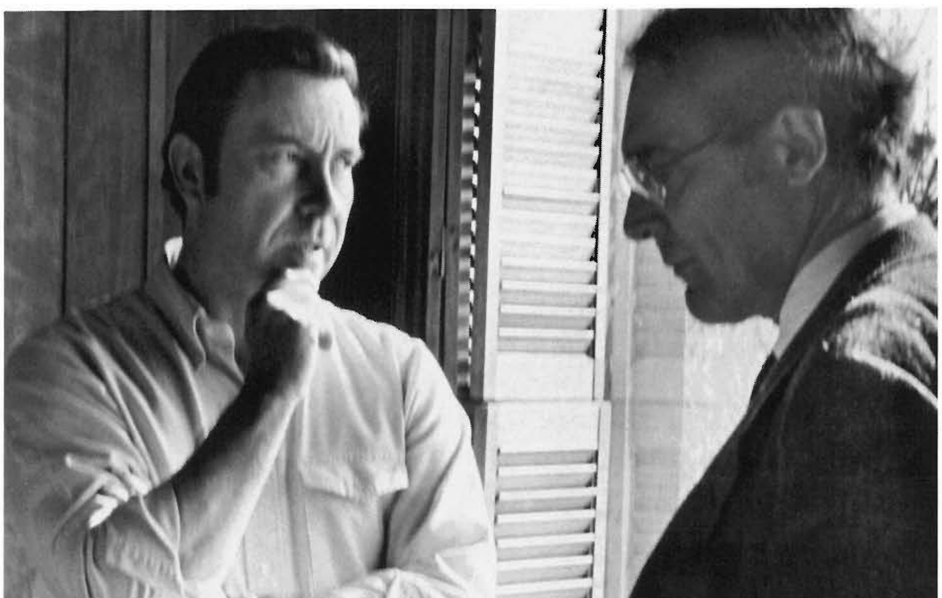

Gian-Carlo Rota *is Professor of Applied Mathematics and Philosophy at the Massachusetts Institute of Technology and Director's Office Fellow. He has been a consultant to Los Alamos for seventeen years.*
David H. Sharp*, a Fellow in the Laboratory's Theoretical Division, is known for his work in quantum field theory, hydrodynamics, and the theory of gravitation.*

Photos by LeRoy N. Sanchez