

Gödel's Doom

A WORK OF FICTION
BY GEORGE ZEBROWSKI

"SO what are you going to beg time for now?" I asked as Witter slid in across from me in the cafeteria booth. A thin, hyper type, he folded his hands in front of my coffee and said, "It's an experiment I want to run on the new AI-5." He spoke very precisely, very insistently, as usual. "I've been haunted by it all my life, but now it can actually be done."


"What do you mean?" I asked, picking up my coffee, afraid that he would knock it over.

"Well, previous Artificial Intelligences were too slow and not capable of complex inference. The question is how much time can you give me?" He brushed back his messy brown hair.

"How much do you need?" I sipped the coffee, sensing his restrained excitement. Witter had always been a valuable worker, so I had to listen and try to keep him happy, within reason, despite his nervous enthusiasms. But he was never satisfied with merely testing equipment and programs for industrial applications.

George Zebrowski is the author of over 50 short stories and essays. His most recent novels are *Macrolife* (Harper & Row, 1979), *The Omega Point Trilogy* (Ace, 1983), and *Sunspacer* (Harper & Row, 1984). The *Monadical Universe* (Ace), a collection of science-fiction stories, will appear this spring.





FICTION: GÖDEL'S DOOM

"I don't know," he said cautiously. "A lot maybe. More than a couple of days."

I put down my coffee, irritated. "You don't know? Can't you estimate?"

"Nope. I'd better explain."

"Go ahead."

"You know about Gödel?"

"I know Gödel's proof, but tell me from scratch. You might be doing some illegal reasoning."

He leaned forward as if he were going to tell me a dirty story. "Well," he said, lowering his voice, "you're familiar with the conclusion that no machine-like entity that proceeds by clearly defined mechanical steps can complete any system that is rich enough to generate simple arithmetic—that is, make it a consistent system in which we could not come up with new, true, and still unproven propositions, in fact ones that would be unprovable in the system, yet clearly true."

"I know, math can't be mechanized."

"Not completely mechanized. We've done it to a remarkable degree..."

"What else is new?"

"Well, if Gödel's proof is true, and human minds can regularly generate true but unprovable propositions in any potentially self-consistent system, then mechanism, or determinism, does not apply to us."

"But what is it that you want to *do*, Witter?" I was only half listening. It was late in the day. The cafeteria was nearly empty and the newly polished floor was a large mirror; our booth seemed to float on it.

"Well," he said, "I want to give the new Artificial Intelligence the command to complete mathematics."

"What?" I suddenly saw what he was getting at.

"Don't you see? We can do an experiment that might settle the nature of the universe—whether we live in a hard determinism or a soft one in which free will is possible."

I smiled, feeling superior. "But we know Gödel was right. Math can't be completed. He gave a powerful formal proof, one in which you can't have it both ways."

Witter, who had been looking away as we spoke, turned his head half around and fixed me with one glassy brown eye. "Come on, Bruno. Why not run the experiment anyway?"

I shrugged and sat back, looking around. "As you said, it might take a long time—forever, if Gödel's right."

"Maybe," he said, finally looking at me with both eyes. The combination of the blue and brown eyes had always given me the creeps. "According to Gödel, the computer

will crank out mathematical statements forever, and we'll never know if the body of the system is a complete one. But if it is complete, then our AI will finish it off in some finite period of time. It's the fastest system ever developed, able to do involved operations that might take centuries otherwise."

"No matter how fast it is, we won't disprove Gödel. He proved that independently of all need to do experiments! Now I know why you want a lot of time. We won't live long enough to learn the result, even if you're right, which you can't be." I started to get up.

"Look," he insisted, "why not do the experiment? If we live in a hard determinism, as so many believe, then it's already true—the AI will complete math or any system we give it. But if Gödel is right, the AI-5 will run on forever, unable to complete."

"We don't have forever. You've gone bonkers."

"Why don't we do it? We *can* do the experiment! Look, for the first time an experiment involving pure logic and math may yield knowledge of the world outside."

That part appealed to me, but I saw a way of being perverse. Was he presenting me with a choice or dictating that I authorize the experiment?

He smiled, anticipating my thought. "Either there's free will, or you're fated to let the experiment be done."

I sighed. "But there's still the matter of how long it will take, Felix. AI-5, no matter how fast it is, may keep running and we won't be able to tell whether it's an un-completable process or just a very long one."

He shrugged. "Aren't you willing to take the chance?"

"This just doesn't make any sense to me at all."

He smiled again. "But it gets to you, doesn't it? My point holds. Why not do it? Just to see. How often in the history of math or logic has there been a chance to do pure theoretical work that might reveal something about the real world?"

"But it's doomed to fail!"

He nodded. "Probably, Bruno, I'll grant you that. But even so, the experiment will be historic. Purely mathematical and empirical at the same time."

"Romantic mathematics, I call it."

"Or Kant's synthetic a priori!"

I'd read some of that metaphysical junk, and he seemed to be stretching it. Sure, synthetic meant acquiring new knowledge, and a priori meant that it wasn't derived from experience, strictly speaking, but from reasoning. Our experiment would give us new knowledge of the universe through nonempirical means. "But you're cheating," I said. "Whatever you call it, using the AI

“Well,” he said, “I want to give the new Artificial Intelligence the command to complete mathematics. Don’t you see? We can do an experiment that might settle the nature of the universe...”

means only doing an empirical experiment."

He cocked one eyebrow and gave me a crazed stare with his blue eye. "Would you say that it would be more empirical if we did it by pencil and paper? That's all Gödel had to work with."

"Okay, I guess I'll have to say that there are no purely a priori activities. Even using the mind alone is a way of reaching out into the universe. What we call experiments are merely corroborations. Einstein himself said that if the experiments didn't come out as he expected, then he'd pity the God who made the universe that way."

"Okay, Bruno, I know you know more than most section chiefs, but are we going to do it or not?"

So we ran the experiment, if you could call it that. Witter was right about one thing. If Gödel's proof was somehow wrong, and we could complete even one system on our fast AI, then a lot of people would have to do a lot of rethinking in the groundwork of logic and math.

But I knew damn well that Gödel couldn't be wrong. Formal proofs do not fall easily. It would be a mistake of some kind if our AI-5 showed that completeness in a significant system was achievable.

All right. We both wanted to see what would happen if we tried it. We pieced the time together from a dozen other projects when people would be away or on vacation.

It was Friday night, after hours. We would be alone until Monday. I sat down at the keyboard and tapped in the command. Witter was sitting next to me, staring up at the bank of screens.

The AI began its run, building arithmetic up out of baby talk. Soon it was all going by in a blur, but the AI showed no sign of slowing down.

"There is one danger," I said as we sat back and waited. "If the AI can't complete arithmetic, it will sift through larger and larger banks of information..."

"It can handle infinite amounts of data," Witter replied.

"Yes, but the power needed for that, Witter, the power! The cost!"

He shook his head. "Don't shout. That won't happen. It will all be over in a few hours at most."

But the AI-5 kept running. An hour went by.

"It's not going to stop, Felix. It can't. Gödel was right. But even if he was wrong, it may take more than our lifetimes to prove it."

"Take it easy, Bruno. Go polish the floor, or something." He was too serene.

Another hour went by. Witter stared at the screens, hypnotized by the blurred flow. Rivers of reasoning ran

from their headwaters to a new ocean of well-formed propositions, and still the ocean was not filled; it would never be filled.

As I looked around at the clean right angles of the room, at the symmetrical terminals and easily accessed units, I began to think that maybe Witter was slightly stupid, that he didn't understand simple logic or the idea of a proof. Gödel's paradoxical conclusion could not be broken, unless it wasn't a double bind to begin with, because you can't have it both ways. Something was very wrong with Felix Witter.

And yet, I wanted him to have a point. This was an experiment, a recourse to more than personal opinion; it could do more, in principle, than reasoning, prediction, or guesswork. Set a powerful genie to do the impossible—not because you think the genie can do it, but because you can ask, and it has the power to do all that's possible. So why not ask, just to see; human beings have always been suspicious of mere reasoning, no matter how powerful. Suddenly I wanted to see Gödel fall, to see the pride and arrogance of mathematicians crumble.

But as we watched the AI-5 chase the mirage, there was no sign of an end, no slowdown at all.

"I'm hungry," I said. "Want a pizza?"

He nodded without looking at me. I got up, went out into the hall, and called it in from the wall phone. Then I alerted the security guard downstairs and asked him to leave it out on a cart in front of our workroom.

"We may have to stop it," I said hours later, "even if it's close to completion." Though the pizza had been very bad, I thought as I eyed the empty boxes on the cart, a full stomach had taken some of the romance out of what we were doing. "We can't tie up all this power and time indefinitely. It's using more every minute, and it'll be my ass if we can't justify it."

"No!" Witter shouted maniacally. "It may be very close."

I burped, waiting for my heartburn to subside. The AI-5 hummed along.

"We can continue from this point onward at another time," I insisted.

"Be quiet!"


I reached over to stop the run. Felix grabbed my hand and pinned it to the panel.

"What's wrong with you?" I demanded.

"Just a few minutes more," he said, fixing me with his mismatched eyes. "We're at the edge of a major discovery!"

"Felix, this can't be done." I struggled to free myself, but his strength was that of a true believer.

"Be still, you fool," he said harshly. "Don't you see? This will be the culmination of our careers. We'll never



match this no matter how hard we work. Gödel is one of the supreme monuments of mathematics, marking the limits of human minds. If we topple him..."

"You may not like what you get," I said, twisting my arm. "If his proof is right, then mechanism is false and minds are not machines. They escape the completeness of the purely mechanical. But if Gödel is wrong, then we're automatons! I'd rather not know."

He shook his head. "There's even more to it than that, Bruno."

"What?" I was breathing very hard, unable to free myself.

"We're opening up the very vitals of reality."

I had to laugh. "By manipulating man-made symbolic structures? You need a bucket of cold water to soak your head in. Let me go!"

"Completion may be only a few minutes away. Do you want to stop and then wonder what might have been?" He tightened his grip.

"But you can't know how far along it is."

He let go of my hand and seemed to cool down, and I found I didn't have the heart to reach over and stop the run.

"You're right," he said, "I'm sorry. It probably is all for nothing."

I massaged my hand. The AI continued its work run. "Don't feel too bad about it," I managed to say. "It was a nice idea, but it had to confirm Gödel. I'm glad we're not machines."

He was shaking his head. "You don't understand. There's no reason to fear that. It's not a problem."

"What isn't?"

"Free will," he said as the AI-5 stopped its run.

Witter and I looked at each other, then at the main screen. It read:

SYSTEM CAPABLE OF GENERATING
ARITHMETIC COMPLETE

"It's a mistake of some kind," I said. Something strange seemed to pass across my eyes. I sat back, expecting to lose consciousness as the tension got to me.

"Maybe," Witter was saying, "but we can test to see if it's a mistake."

"How?" I heard myself ask, even though I knew the answer.

"By trying to make a true statement that is not provable in the system. As long as the AI can show us that we can't make such a statement by proving it, then the system is complete."

The room went black for a second. "But maybe we can't make such a statement," I said.

"We can try," he answered.

We tried for the next 12 hours. I was relieved that our prime AI was no longer running a huge power draw. Witter brought a smaller AI on-line and had it question the alleged complete system achieved by the AI-5. It failed to come up with a single true proposition that was not provable in the complete system.

"There's no question about it," Felix said finally.

"There's only one thing left to do," I replied. "We've got to run the whole thing again."

Witter looked at me, smiled strangely, then sat down and gave the command.

As the AI-5 began its second run at Gödel, Witter turned to me and said, "Funny about determinism. I always think of it as stuff outside me, pushing at my skin. But I feel free inside. When that second run finishes, we'll be certain that we're living in a hard determinism. No choice is our own, if we've understood the word correctly. Even our decision to run the AI-5 again was not made freely. We're automatons. No avoiding the conclusion, Bruno."

He was baiting me, I was sure. "But we resist the notion. Doesn't that suggest something?"

He shrugged. "That we're free in our minds but not in our actions. We can envision alternatives, but whichever one we pick is determined, right up through an infinite future."

"Witter, I thought you were intelligent. There can't be such a thing as unconditioned freedom. There are always initial conditions—necessary and sufficient conditions for every choice. Otherwise we could perform miracles, make happen things that are uncaused. The existence of free will cannot violate causality."

He grimaced at me and I felt stupid. "Yeah, I know all that. But *do* we have the freedom to choose between alternatives?"

"I think we do. Physical conditions make us both the determined and determinators in our own right. Things affect us and we affect them. Determinism goes right down into us, into our consciousness and will, and we send it back out. I couldn't prove it to you without a physiologist, though."

The AI-5 was still running its second completion smoothly. If it succeeded, then it might be that we were living in a universe where even choice among alternatives was an illusion.

"It's a mistake of some kind," I said. Something strange seemed to pass across my eyes. I sat back, expecting to lose consciousness as the tension got to me. "Maybe," Witter was saying.

Witter looked at me suddenly. "I wonder if our running this program can have an effect on the universe we live in?"

"What are you talking about?" I asked. He seemed to have a mind like a break dancer.

"Maybe our attempting what Gödel said was impossible can change the universe?"

"I don't think so, Felix. But there are other things you might like to consider."

He took a deep breath. "What's that?"

"Well, we began with the idea that no finitary deductive system can complete a rich, self-consistent system. But what if the AI-5 is not a finitary deductive system? Assume it can work outside the limits of the human mind, which is all that Gödel may have charted. It was all he could demonstrate because he had only his own mind to work with."

Witter nodded. "I see what you mean. If our AI reaches completion, then it follows, perhaps, that it's not a finitary deductive system, and we can draw no conclusions about the nature of the universe."

I smiled. "Right. And we don't have to worry about being automations, or that our sense of inner freedom is a mirror trick of some kind. Free will is a special case of determinism. It's determinism from the inside. The means of determinism are also those of free will."

Witter was watching the screen with a worried look on his face, as if he now expected the AI to fail. It didn't matter one way or the other, if what I had said was true.

"Unconditioned free will would be omnipotence," I continued, "and that's an absurd state to be in. No law, no causal structures. It's just a conceptual extreme, like infinities."

"Something is working against us," Witter said softly.

"What do you mean?"

He gripped the panel. "It won't come out the same way twice," he replied.

"You're still mistaking the maps for reality," I said.

"Look at the time, you fool! It's almost as long as before. If the AI doesn't repeat its completion in the same time, it will run on forever."

"So what. We have the first completion in memory, step by step, for whatever it's worth."

He swiveled his chair and glared at me. His eyes were bloodshot and had dark circles around them. The whole experiment, I saw, was eating up his entire energy. "You don't see, do you?" he said. "You think in terms of tricks of language, ways of speaking...you can't imagine worlds dying and others supplanting them. You don't give a damn about anything except apportioning time and keeping other administrators happy."

"What are you talking about, Felix? I'm here with you,

and we're doing what you wanted. Have you lost your mind?" I almost felt hurt, as if he were questioning my loyalty.

He pointed to the clock on the wall. "Look, time's up and our AI is still running."

"So what? It was a fluke the first time, a mistake. You can't beat Gödel, and it wouldn't matter if you could."

He laughed. "You still don't see!"

"No, I don't."

The AI-5 was still running.

"It will run forever this time. Our decision to run the experiment puts us at a great juncture between possible universes. We collapsed the wave function reaching our minds."

"What are you saying?" I demanded.

"Proving that *our* universe was deterministic threw us into a freer one. Gödel proved his work in the wrong universe. Here the AI will run forever. But if we stop it and start again, something even stranger might happen."

"You're off the deep end now," I said, feeling sorry for him.

"We might be moving across a whole series of universes, drawing closer to the unconditional omnipotence that has the true freedom to

be everything..."

"Yeah, and can't become anything in particular. That's what I was saying. Witter, wake up. We have the other program. Go see for yourself. That system was completed. In this one there's obviously some kind of difficulty. Neither result means a thing. Get that through your stupid head!" Mathematicians were all idealists to some degree or another, always secretly believing in the literal existence of infinities, numbers, and tortured geometries. Witter was no exception.

He shook his head and smiled. "There's nothing in the memory, Bruno. See for yourself. Go ahead, punch it up."

I leaned forward and punched in the order. Nothing came up. I went into search mode. Still nothing.

"We've left that universe behind," Witter said.

"It's got to be here," I said.

The screen remained blank.

"You erased the memory!" I shouted.

"I did not," he replied softly, and I knew he was telling the truth.

I glanced at the food cart; it winked out of existence.

"Did you see that?" I asked.

"Bruno!" Witter shouted. "We've escaped a totalitarian cosmos. We're free!"

"Relatively," I said, shaken.

He was looking at me strangely, and I saw that both his eyes were now brown. As the AI-5 continued its endless run into a free infinity, I feared what we would find when we went outside... □

CONTINUED ON PAGE 168



Gödel's Theorem: The paradox at the heart of modern man

In the summer of 1930, at age 24, mathematician Kurt Gödel proved a strange theorem: mathematics is open-ended. There can never be a final, best system of mathematics. Any mathematical theory we can describe is incomplete; that is, any theory we invent, sooner or later, will run into certain simple problems it cannot solve. This is Gödel's Incompleteness Theorem, and when Gödel received an honorary degree from Harvard University in 1952, the citation termed his achievement the most important advance in mathematical logic in a quarter century.

For those of us who aren't familiar with higher mathematics, the import of Gödel's Theorem is somewhat hard to understand. In fact, when the original paper appeared, it probably was unintelligible to many professional mathematicians. But philosophically, it turned the ordered world of mathematics upside down. Never again would math be a closed system, ultimately provable by its axioms. To the pure mathematician, this meant infinite room for the queen of the sciences to evolve and grow. And eventually, the theorem's impact subtly changed the layman's view of his universe.

Generally speaking, a mathematical theory is formulated in a specialized language—a language including things like numbers, variable symbols, and so on. Some "sentences" in this specialized language are taken as the axioms, or basic assumptions, of the theory. By combining these axioms according to certain definite rules, we produce logical proofs. A proof is really just an orderly sequence of sentences, with the last sentence of the proof the thing that is being proved. Something that is proved in this manner is called a theorem.

Eventually, people realized that the notion of "a mathematical theory" is equivalent to the more concrete notion of "a program for generating sentences." Any mathematical theory can be turned into a certain kind of computer program that successively prints out all the theorems that the original theory would have proved. It is this connection between concrete machinery and abstract theory that gives modern mathematical logic its importance.

As an example of a mechanized theory, consider euclidean geometry. Alfred Tarski showed that it is possible to write a computer program that prints out, one after another, every theorem that can ever possibly be proved from Euclid's five postulates. What makes the euclidean geometry program particularly powerful is that it is complete. That is, given any sentence *S* about euclidean geometry, either *S* or *not S* will show up in the theorem list sooner or later. If the Greeks had known this program, then they wouldn't have needed Euclid!

Let's say you want to know if a triangle's three angle bisectors always meet in a point. All you have to do is sit down and watch the euclidean-geometry program's printout until it says either "A triangle's three angle bisectors always meet in a point," or "A triangle's three angle bisectors do not always meet in a point." Anyone who has worked with computers will realize that you may have to wait a very long time—perhaps

billions of years—but at least you will not have to wait forever. The fact that the geometry program is complete guarantees a finite wait.

Of course geometry isn't all of mathematics. Geometry talks about points and lines but not about numbers, which aren't so finite. As soon as a mathematical theory is broad enough to deal with the manipulation of natural numbers, then Gödel's Theorem becomes important. As numbers approach infinity we go beyond what is provable with axioms in a finite mathematical system. In other words, it is impossible to construct a mathematical system that contains all the axioms that can be formulated within the language of arithmetic. When we prove one, another awaits—ad infinitum. Whenever I create a program *p* to print out true sentences about numbers, the program will be incomplete. Let's take one of these sentences and call it *G_p*, the Gödel sentence for the program *p*. This true sentence states that neither *G_p* nor *not-G_p* (the negation of the sentence) will ever appear in the list of theorems.

If you think about it logically, you can see why *G_p* must in fact be true, yet not provable by *p*. We are assuming that *p* prints all true sentences, but that *G_p* is a true sentence whose property is that it cannot be printed by *p*. There's a paradox working here.

A number of philosophers think that Gödel's Theorem has something to do with the question of whether computers will ever "think" just like human beings. Can a machine understand our paradox? Perhaps not now, although we may be creating the conditions for such an evolution. And are we really just complex machines ourselves? In *The Freedom of the Will*, J. R. Lucas of Oxford argues that the very fact that we understand Gödel's Theorem means we are not machines. Briefly, his argument goes like this. Suppose someone writes a program *l* and tells Lucas that it is designed to print out exactly those sentences which Lucas will, sooner or later, say to be true. Lucas now claims that he can find the Gödel sentence *G_l* for *l*, and that he can see that *G_l* is indeed true. But based on Gödel's Theorem, the program *l* will never print *G_l*. So Lucas is in fact smarter than the program *l*!

Most theorists, myself included, think Lucas is wrong. The flaw in his argument is that the Lucas program *l* will be so complex that he will never actually find the sentence *G_l*. Perhaps what Gödel's Theorem really tells us is that we may indeed be sophisticated machines, in the broad sense, but that our programs are so complicated we'll never understand them.

In any case, Gödel's Theorem is a rare example of a mathematical result that matters. Mathematics is richer and more complex than anyone before Gödel had suspected. Mathematics may be on the brink of solving some of the oldest philosophical puzzles about free will, consciousness, and the nature of reality.

In the final analysis, we may be open-ended, too.

—Rudy Rucker



Rudy Rucker is a novelist and mathematician. His book *Software* (Ace) deals with intelligent robots, and his popular-science book *Infinity and the Mind* (Birkhauser) contains one of the best popular discussions of Gödel's Theorem.