Roundworm Cells and Cancer Genes

By studying this... simple, experimentally tractable micro-organism, my lab has been able to help out in the big problem of trying to understand what happens during cancer.

by Paul W. Sternberg

My major obsession in life is to understand animal development—how a single cell divides and generates the many specialized cells that form the adult organism. And, as it happens, by studying this process of development in a very simple, experimentally tractable microorganism, my lab has been able to help out in the big problem of trying to understand what happens during cancer. I'll start by introducing the current concepts of what happens in the early stages of cancer, and then I'll tell you about the roundworm we've been studying, and then at the end I'll bring it all together.

A cancer arises from cells that escape their normal growth control and divide continuously. Eventually the cells acquire the ability to invade surrounding tissue—that is, metastasize—or commandeer a blood supply, or both. Imagine a nicely organized tissue—say a layer of cells such as your intestinal wall. The cells are slowly dividing to replenish themselves. Say you get a mutation—a change in a gene in a particular cell that gives it different properties. That mutation, in some instances, might cause that cell to divide faster than its neighbors. Soon the faster-dividing cells are encompassing more and more of the layer. They start to take over, in other words. Then another mutation might cause the cells to grow even faster, and lose their ability to maintain their nice, sheetlike formation. They might start forming a lump. Then there might be a third mutation that divides even faster and has other properties, for example the ability to crawl around and invade nearby tissue. It's by a series of such mutations that most cancers progress.

Graduate student Gregg Jongeward watches roundworms through a stereomicroscope, while his inflatable friend appears to be preparing for a doctoral candidacy exam. (Apologies to Edvard Munch.)

Typically, it's more than three mutations, and they don't happen very fast, which is why some tumors can take 10–20 years to develop. A "genetic predisposition" to cancer often means that the cells have one such mutation to begin with, which shortens the chain of mutations needed for the cells to become cancerous. Certain mutations make the cells pretty sloppy at replicating themselves, increasing the rate of mutation. Normal cells replicate their genetic material very accurately, so a mutation in the machinery that insures this accuracy would quickly lead to more mutations. A recently discovered colo-rectal cancer-predisposition gene might be of this type.

There are two kinds of genes that can mutate to cause cancer. Oncogenes—that is, cancer-causing genes—are one type. This class of genes was discovered about 20 years ago. An oncogene's normal function seems to be to stimulate cell growth and division, so that mutations activating these genes inappropriately would likely lead to cancer. The other kind of genes, discovered over the last 10 years, are called tumor-suppressor genes. These genes tend to inhibit cell growth and division. If such a gene is eliminated from a cell, that cell will grow and divide when it shouldn't.

To understand how these changes can affect a particular gene, we need to review how a gene directs the synthesis of a protein. Proteins are the building blocks of the cell—the structural components that form the cell's architecture, the enzymes that form the cell's machinery, and the messengers that regulate the cell's activities.

A protein consists of hundreds or even thousands of small building blocks, called amino acids, linked together like beads on a string in a very specific order. The genetic instructions of every organism are encoded in very long molecules known as DNA. Particular segments of that DNA, called genes, are transcribed and copied into another nucleic acid called messenger RNA. Each gene typically contains the instructions for one protein. The messenger RNAs are then translated into proteins by some very specific and exquisite machinery in the cell. The machinery is a complex of perhaps 50 to 80 proteins and several pieces of RNA. The machinery also does proofreading, making sure that each amino acid is put in the right order. The proteins then fold up and form three-dimensional structures determined by their sequence of amino acids, and these structures do the work of the cell.

Some mutations decrease or abolish a gene's activity. For example, the transcription of DNA into RNA could be blocked, or the translation of messenger RNA into protein could be blocked, or the folding of the protein could be abnormal, or the protein could be made but wouldn't work. Or the gene could just be deleted from the genome. Other mutations cause the protein to be more active than normal, or make the gene direct the synthesis of too much protein. All of these things occur in nature. So a mutation could inactivate a tumor-suppressor gene and prevent the synthesis of an inhibitor, which would lead to more cell growth and division, and lead to cancer. Or a mutation could activate or make more of an oncogene, leading to cell growth and division and cancer. Our task is to identify all these genes—and people think that there are at least 100 of them—and figure out what each gene's protein does, and how all these genes and proteins are linked together to form the circuitry that controls what the cell does.

A surprisingly high percentage of cells die during normal development—they either commit suicide or they're murdered.

The normal role of the genes that, when mutated, lead to tumors is to determine a cell's fate during development. A developing cell has to make many choices. It has to decide how many rounds of cell division to undergo—does it not divide at all, or does it generate a million progeny cells? If it divides, what kind of progeny does it produce—skin, nerve, muscle, liver, or what? Does the cell survive, or does it die? A surprisingly high percentage of cells die during normal development—they either commit suicide or they're murdered. And finally, the cell must choose whether to stay where it was formed, or to crawl to another location in the organism, like the neural-precursor cells that Associate Professor of Biology David Anderson studies

Cellular induction: Whether the individual cells in the bottom row become type A or type B cells depends on whether they are within range of a signal from the green cell above them. The black lines are generic tissue structures.

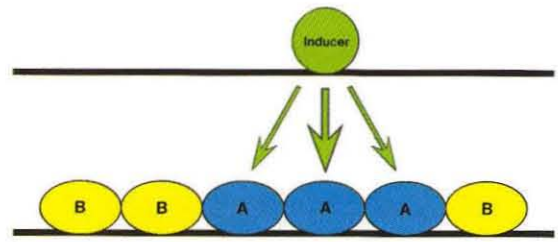

[*E&S*, Spring 1990]. The problem that I set out to study 10 or 15 years ago is: How are the instructions for the fate of particular cells coded in their DNA?

Now, in most organisms, what a cell does depends on signals from its neighbors. In the simplest possible case, consider a type A cell, colored blue in the drawing above. The fact that this cell is a blue A cell as opposed to a yellow, or B, cell, depends on a signal from a neighboring green cell, which I'll call an inducer cell. We can demonstrate this by surgically removing the green cell, and when we do, the cell that should be an A is instead type B. Or we can get rid of the A cell; then its neighbor, which is normally a B cell, becomes an A. So we conclude that the A cell becomes an A by virtue of the fact that it receives a signal from the inducer cell, and the B cell can't become an A because it doesn't get the signal. The signal is a chemical—usually a protein, in the examples that I've been studying—that is secreted, or released, from the inducer cell and interacts with a protein on the surface of the A-cell-to-be and directs its development.

The organism I spend most of my time working with is *Caenorhabditis elegans*, one of the nematodes, or soil roundworms. Nematodes are as common as the dirt under your feet—there are perhaps a hundred of them per cubic inch of soil—and they literally stick to your shoes as you walk through the grass. But you're not in a constant state of being grossed out by this because they're so small that they're almost invisible to the naked eye. At right is a worm in its normal habitat in the laboratory. It's crawling

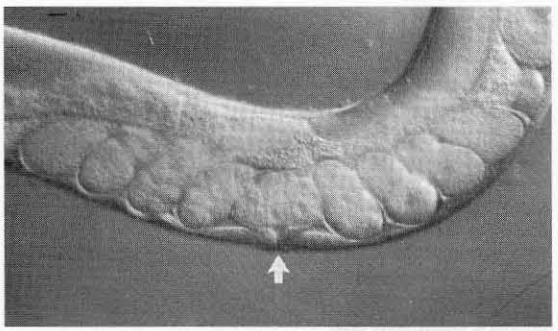

Right: A full-grown, one-millimeter-long roundworm takes its constitutional on a petri dish. The wavy lines are tracks left by other worms. The dark blot below and to the worm's left is an egg; below that is a baby worm.
Far right: The vulva (arrow). The line of nine spudlike objects above and flanking the vulva are fertilized eggs. The dimples in the eggs are cell nuclei; thus the egg directly above the vulva has already divided into at least eight cells.

on a petri dish, in a slurry of the bacteria it eats. These small creatures have a number of advantages as lab animals. They're very easy to raise. They're also easy to handle—we can pick them up with very small, sterile platinum wires, and move them from petri dish to petri dish. And they grow very rapidly, going from an egg to an egg-layer in three and a half days. We get two generations a week for genetic studies, so we can do a lot of experiments. One worm on a petri dish will give rise to 300 progeny in, say, four or five days. Of course, there's a slight disadvantage in that you have to look at the worms daily to follow their growth, as opposed to most other organisms, where you can ignore them for a week at a time because things don't happen very fast.

The key to our technique is that the animals are transparent, so that we can actually watch individual cells as they grow in the intact organism, and follow what becomes of them. (This approach was developed in 1976 by John Sulston at the MRC Laboratories of Molecular Biology in Cambridge, England.) We put the worm under a microscope that magnifies it about a thousand times, and as the worm goes about its business crawling all over the petri dish, we twiddle knobs under the microscope stage to move it around and keep the worm in our field of view. This skill takes some practice—it takes most students several weeks to acquire the knack—but it has the added advantage of making us tough opponents in the video arcade. We can also remove a particular cell by focusing a laser microbeam through the microscope's optics onto that cell, boiling it. Furthermore, roundworms only have a small number of cells. Excluding the germline—the eggs and the sperm—the hermaphrodites have 959 cells, and the males have 1031. (Hermaphrodites are females that make sperm as well as eggs.) The number isn't completely precise, because occasionally a worm is plus or minus one cell. So, after years of study, we now know all the cells in the organism as individuals. In many cases, we know what the cell is going to do before it does. We can tell by its position that a cell is going to make skin instead of a vulva, for example, yet we can show by doing the sort of microsurgical experiment I described above that the cell hasn't yet made the choice itself.

My lab has been studying the process by which the vulva is formed on the belly of the developing worm. The vulva is easy to study, because it develops rapidly—in just a few hours—and it involves only a handful of cells, making it easier to track their individual fates. And since the vulva is not vital to the worm's growth or reproduction, we can easily grow viable mutant strains that have inborn (hereditary) defects in vulval development. The vulva is the organ that gets the eggs out of the animal. Once eggs are produced in the ovary, they get fertilized in the gonad by the worm's own sperm, or by sperm from a male worm. (These eggs are quite small—1,000 would fit on the head of a pin.) The fertilized eggs start dividing. Once an egg has divided into a 20-cell embryo, it is forcibly ejected through the vulva and onto the petri dish to make room for another egg. The vulva is actually a specialized piece of skin, as Sulston discovered. In the embryo's developing gonad he found

Below: Although they don't yet know it themselves, the cells indicated by white arrows are fated to become vulval cells, while those marked with solid black arrows will become skin. The anchor cell (dotted black arrow) is the divinity that shapes their ends.

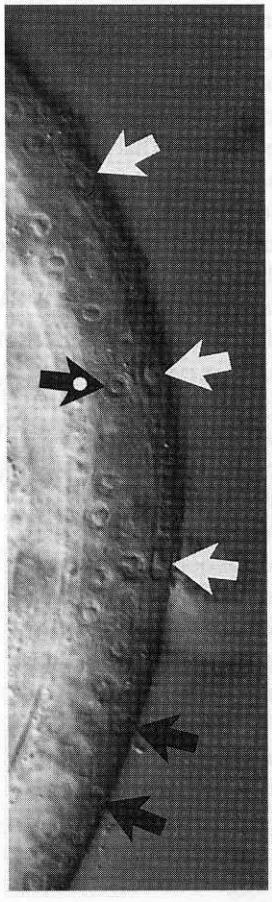

Right: A few hours later, precursor cells P5.p, P6.p, and P7.p have each given rise to a family of cells, as shown by the black lines. These cells, which look like sunny-side-up eggs, are now moving inward to form the vulva, visible as a dark, arrow-shaped indentation. The letters indicate the cell's mode of division: Longitudinal, Transverse, or Non-dividing. The anchor cell is labeled "ac," and is surrounded by the developing uterus.

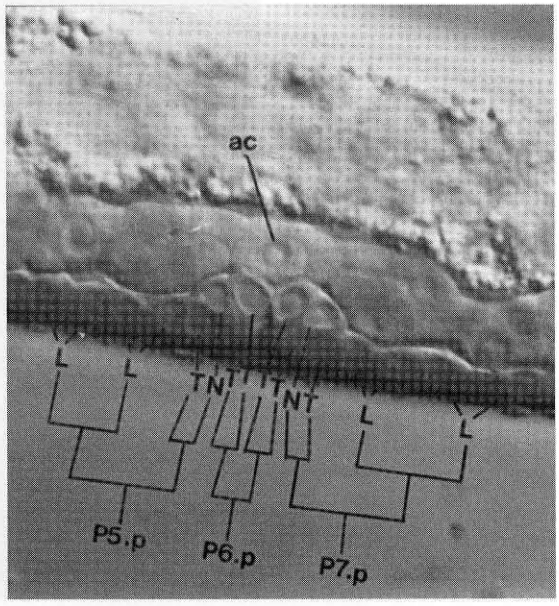

one particular cell, called the anchor cell, that signals three precursor cells in the skin to divide an extra time, start moving into the worm's body, go through a complex series of shape changes, divide again, and generate the cells of the vulva.

The problem my lab is working on is this: How do these cells know to become specialized and make a vulva instead of remaining nonspecialized and making just skin? In the smooth belly of the adolescent worm at left, the three white-arrowed cells will give rise to the vulva, and the two black-arrowed ones won't—they'll just become skin. But if given the chance, they would make a vulva. The signaling cell, shown here with a dotted black arrow, produces a signal that reaches the three nearby cells but not the more distant cells. If we destroyed those three cells, there would still be a vulva because the outer cells would move in and make one. So these cells really have two choices—they can make a vulva or skin.

And the beauty of roundworms is, we can see it all happen. If we sit in front of that microscope for eight hours, we can actually watch these three cells divide, move into the worm's body, and connect up to the developing gonad and form the vulva. This technique allows us to do a variety of experiments with unparalleled precision, because every animal is the same, and we know all its cells. We get the same reproducible effect from the same perturbation, a level of precision that you rarely get with more complicated animals.

We've found lots of mutations that affect vulval development, and I'll give you examples of two classes. One class we call vulvaless. All the cells are present—the signaling cell, and the cells that normally respond to the signal—but no vulva is formed. There are two things that could be going wrong here: The cells could be failing to respond to the signal, or it could be that the signal is not being sent. The vulvaless class contains examples of both kinds of malfunction. In the other class of mutations, called multivulva, not only do the three normal cells make the vulva, but the other three more distant cells also try to make vulvas. These mutants are easily recognizable to the trained eye, because they have lumps on their bellies formed by cells in the wrong location that are trying to make vulval structures but can't quite do it. (They also have a normal vulva, so they can still lay eggs.) One of the really interesting properties of these multivulval mutants is that, even if we get rid of the signaling cell, all the cells still make vulvas. The cells act as if they are constantly getting the signal. I'll explain why shortly.

These mutations allow us to identify the genes involved in the signaling process, but our real goal is to understand the order in which they act. Over the last few years, genetic studies have told us that these genes make proteins that act like switches. That is, the proteins can exist in two states—active and inactive, or ON and OFF. On the opposite page is a simplified model of three of these switchlike proteins acting in series. The *lin-3* signal activates the *let-23* protein, which impinges on the next switch (*let-60 ras*) and turns it on, which in turn throws the third switch, *lin-45 raf*, and that switch then makes the cell turn

Below, left: A multivulva mutant worm. The three growths on the worm's right (i.e., belly) side are vulva wannabes. The normal vulva is also visible, midway between the lower two growths.
Right: This simplified signaling pathway consists of three switchlike proteins acting in series to decide a cell's fate.

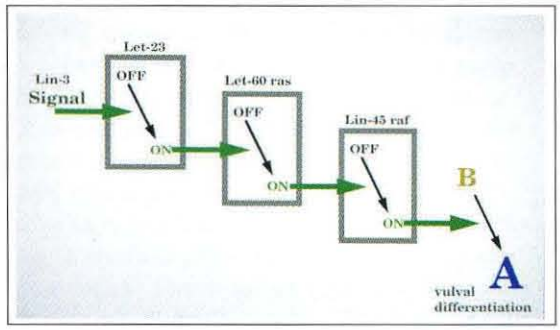

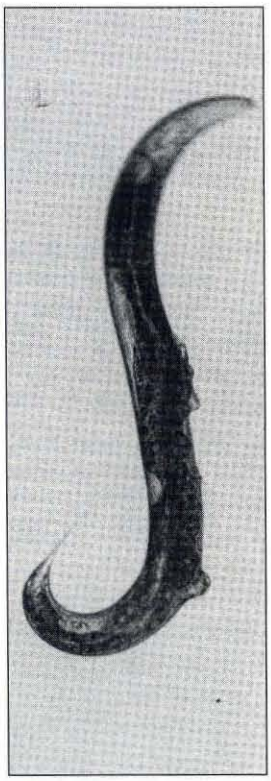

from type B into type A, which differentiates into the vulva. (These genes' arcane names come from abbreviations describing what their proteins do—*let* stands for lethal, for example, and the number 23 indicates it was the 23rd gene discovered that, when eliminated, causes the worms to die.)

Mutations can affect this process in several ways. For example, if we make a mutation that eliminates the activity of the *lin-45 raf* gene, the third switch is now broken, locked in the OFF position. The *lin-3* signal comes on, and turns the first switch on, which turns the second switch on, which tries to turn the broken switch on, and nothing happens. The cell stays as type B. There's no vulva formed. This worm is one of several strains of vulvaless mutant worms we've made. Other mutations that cause a particular protein to be much too active—locked into the ON state—cause multivulval worms. If, say, the second switch (*let-60 ras*) is always on, it will turn the third switch on, and make the cell become an A, even if there's no *lin-3* signal. Because the switch is stuck in the ON position, it doesn't need anything beforehand to turn it on. In some cases, like the *let-60 ras* gene, we have one mutation that locks it ON and another that locks it OFF, so we can set the switch in whichever position we want.

So the key experiment is, if we have one mutation that locks one switch ON, and another mutation that locks another switch OFF, what happens if we put both mutations together in one animal through a simple genetic cross? Which mutation wins? There are two possibilities: Say the switch that's stuck in the OFF position acts after the switch that's stuck ON. The signal comes in and turns on switch number one. Switch number two is broken in the ON position anyhow and is already trying to turn switch three on, but can't because number three is stuck OFF, and the cell stays a B. Switch number three wins. Alternatively, if the broken OFF switch is earlier in the pathway, say at switch number one, when the signal comes, nothing happens at number one, but since number two is stuck ON, it will turn on number three regardless, so number two wins. Either way, the mutation farthest downstream prevails.

By doing many such experiments, we can come up with the order in which the genes act. (In fact, all Caltech biology majors are required to take a worm-genetics lab where they make such crosses and try to deduce a pathway.) There are considerably more genes involved than just these few, and tracing their interactions is much more complex than what I've just described—for example, some genes are inhibitors that send a signal downstream that tells another gene *not* to turn on; an inhibitor gene stuck on ON acts like an ordinary gene stuck on OFF, but that's the idea. I started working on this pathway about a decade ago, and we've probably only figured out one-fifth of it.

But what does this have to do with cancer? It turns out that all four of these genes have counterparts in humans. Raffi Aroian, Min Han, Andy Golden, Russell Hill, and Jane Mendel in my laboratory have demonstrated this in two ways. First, recall that every protein consists of a particular sequence of amino acids that

.	.
.	.
.	.
N	S
G	G
G	S
Q	Q
Q	Q
Q	V
Q	E
Q	Q
P	P
T	T
G	G
S	S
I	V
L	L
W	W
M	M
A	A
P	P
E	E
V	V
I	I
R	R
M	M
Q	Q
D	D
D	N
N	N
P	P
Y	F
T	S
P	F
Q	Q
S	S
D	D
V	V
Y	Y
S	S
F	Y
G	G
.	.
.	.
.	.

Left: A portion of the amino-acid sequence for the *lin-45* roundworm protein (left) and the human *raf* protein (right). Right: Postdoc Wendy Katz maneuvers a roundworm cell into the line of fire of the laser—the thin vertical black box sitting atop the microscope.

determines the protein's chemical structure. There are 20 different amino acids, each of which we represent by a one-letter code. Thus A stands for alanine, G is glycine, R is arginine, S is serine, T is threonine, and so on. One part of the roundworm gene, *lin-45*, tells the cell to assemble this sequence of amino acids: TGSILWMAPEVIRM. When we look at the sequence of the human *raf* gene, we find almost the same sequence—TGSVLWMAPEVIRM. In fact, all the letters shown in red at left are the same, and this degree of similarity extends over several hundred amino acids. So when we see two proteins this similar in sequence, we have a pretty good bet that the proteins have a similar structure. We can then hypothesize that they have the same function in the cell, which generally turns out to be true.

This is a very important hypothesis, and it is supported by our second proof: We take a gene from one organism and put it in the other and see if it still functions in the same way. This technology, called transgenic technology, is now available for a variety of simple experimental organisms, including worms, fruit flies, mice, and yeast. We inject into the ovary of the worm the human (or whatever) DNA that has the gene we want to compare. Some of that DNA gets incorporated—we don't know the details of how it happens, we're just lucky that it does—into some of the worm eggs. These worms, called transgenic worms, now have acquired that injected gene, and they'll pass it on to their progeny. We don't want to have to take a DNA sample from each worm to find out if it has the gene, so we put a marker on the gene we inject. One nice marker is a mutant worm gene called roller, which causes the animal to have a twisted body. When roller worms try to move forward, they twist to the right and end up rolling. It's very easy to spot them—all their siblings are crawling around very nicely, and the rollers just go in circles.

I'll give you a particularly spectacular example of such an experiment, using a rodent gene that's about 80 percent identical to a roundworm gene that's part of a different pathway we're also studying. If you took five normal, or "wild-type," worms and put them on a microscope slide and waited for five minutes, they'd all be gone—they'd have crawled out of your field of view. But if you put five mutant worms there and waited 40 minutes—eight times as long—all the worms would still be in your field of view. Some of them might not have moved at all. The worms have a neurological defect that makes them have serious trouble moving forward or backward. Junho Lee and Gregg Jongeward in my lab have identified this gene and the protein it makes. It turns out that they're very similar to a gene and protein that other people have discovered in rodents and humans, where it is involved in moving membrane proteins around inside the cell. So we asked ourselves, if we take the rodent gene and inject it into our mutant worm, will the mouse gene function? If it does, the worm should move normally. There are subtleties, of course—we have to make the mouse gene into a hybrid so that it stands a chance of functioning in the worm, yet have most of the protein encoded by the mouse's DNA. We've done this experi-

ment, and the "rescued" worms crawl just fine. So this mouse gene will function in the roundworm, and we can confidently say, to a first approximation, that the two genes are the same.

We know from similar examples that each of the genes involved in vulva differentiation in the roundworm has a counterpart—or several counterparts—in humans. Thus the protein encoded by the *lin-3* gene looks like human EGF, or epidermal growth factor, protein. And just as *lin-3* is a signal between roundworm cells, EGF and related proteins act as signals between human cells. Then, on the responding cell, there's a protein that acts as the receptor—in the worm it's *let-23*, which resembles the EGF receptor protein in humans. This receptor binds to the signaling protein and controls what that cell does in response. Inside the cell, the signal is somehow transduced, or changed in form, by other proteins—switches like *let-60 ras* and *lin-45 raf* in the worm, and their human twins, genes called simply *ras* and *raf*. The transduced signal travels down pathways that many research groups are just beginning to explore, and eventually reaches the cell's nucleus. There the signal controls what genes are turned on to make the cell proliferate, or change shape, or otherwise choose its fate.

Above: Grad student Junho Lee pulls a worm out of a petri dish. The worm is impaled on the tip of the stainless steel probe in his right hand.
Below: A three-day's supply of fresh, nutrient-laden petri dishes for the Sternberg lab.

Since we can draw a one-to-one correspondence between the worm genes and the human genes, we can say, "If the genes work in this particular order during this particular process in worm development, then we predict that in humans, these genes will act in the same series to control cell growth." The genes' actions may have different effects because they are triggering other switches that the worm doesn't have, but we expect the order of their triggering to be the same. And this prediction turns out to be correct. So we can use the simple genetics of one organism—the worm—to learn about some really important genes in an organism that we care a lot more about—ourselves. And all of the human equivalents are known oncogenes. In fact, *ras* is a particularly infamous oncogene—it's the one most frequently mutated in colo-rectal cancer.

But this isn't the whole story. If it were, we could probably solve the cancer problem in a few years. Unfortunately, there are a lot of genes still to go. For example, there are at least two other proteins between the EGF receptor and *ras*. Just in the last few months, it's been discovered that *ras* interacts physically with the *raf* protein. Then, after *raf*, but before cell growth, there are a lot more genes. We still need to figure out what they are, and the order in which they act, and then we need to know the details of what controls them and how they function. That's the level of understanding we're going to need in order to look at a tumor and say what went wrong. And *that* knowledge will enable people who are good at that sort of thing to design ways of intervening—that is, to come up with therapeutics or new drugs.

There are two ways to eradicate cancer: One is to prevent it from happening in the first place. We can all stop smoking; we can get rid of a lot of environmental carcinogens. We know that most agents that lead to cancer are either mutagens that mutate the DNA or tumor promoters that stimulate cell proliferation. And the more cells divide, the more likely they are to mutate and cause cancer. That's something we can take care of without any fancy science—we just have to use common sense. The other way, to eradicate cancers that have already started, is to come up with the next generation of very specific anticancer drugs. The drugs we have now essentially kill any and all dividing cells. This has nasty side effects, because the cells that line the stomach, and the cells that make hair (not to mention the ones that do a host of other things) also have to divide. You wind up killing them, too, which is why chemotherapy patients suffer nausea and hair loss. But as researchers discover which protein binds to which receptor to send a signal, they can try to come up with drugs that interfere only with those specific interactions. No one's done it yet, but it's promising—last year, a number of biotech start-up companies formed to take advantage of the knowledge we've gained about the signaling pathways in these oncogenes. The point is, the basic understanding of the mechanism will lead to large-scale efforts to come up with drugs based on those mechanisms. □

Paul Sternberg chose biology as a major because "I couldn't get an appointment with the economics advisor." Sternberg earned his BA in biology from Hampshire College in 1978, and his PhD from MIT in 1984. He came to Caltech as an assistant professor in 1987, and was promoted to associate professor in 1992. Sternberg holds a joint appointment with the Howard Hughes Medical Institute in Pasadena, where he was appointed assistant investigator in 1989, becoming associate investigator in 1992. This article is adapted from the Seminar Day talk he gave in May.