[music plays]

[ANNOUNCER:] From the Howard Hughes Medical Institute, the 2013 Holiday Lectures on Science. This year's lectures, "Medicine in the Genomic Era," will be given by Dr. Charles Sawyers, Howard Hughes Medical Institute investigator at Memorial Sloan-Kettering Cancer Center, and by Dr. Christopher Walsh, Howard Hughes Medical Institute investigator at Boston Children's Hospital. The fourth lecture is titled "From Cancer Genomics to Cancer Drugs." And now a brief video to introduce our lecturer, Dr. Charles Sawyers.

[music plays]

[DR. SAWYERS:] DNA sequencing technology has been revolutionized in the last several years, not only in the price. It's no longer cost-prohibitive to dream of any kind of sequencing project, but also the depth at which once can sequence. It's now possible to sequence, for a specific region of a gene that you're interested in, to sequence what's called a million-fold deep, meaning if there are a million cells in that sample, you could sequence the gene in each cell. This is a kind of mind-boggling possibility. It's also now possible to sequence DNA from a single cell, which you can capture from patients' blood.

So I envision over a five-year time frame at a minimum, but certainly within 10 years, serially tracking the course of a cancer by sequencing DNA from the blood of these patients and understanding, in real time, how the tumor is evolving in response to a drug regimen. I think the question of whether this approach will result in cancer being like HIV, where a three-drug cocktail has taken a lethal disease and converted it to a chronic disease. The answer is yes, cancer will be a chronic disease. We will, through understanding mechanisms of resistance, be able to give cocktails that will maintain the cancer in a remission.

But there's growing evidence that, if you give two or three drugs together at the beginning, you can increase the magnitude of the initial remission to levels much deeper than we currently measure. And by doing so, I think it's, in principle, possible to cure the cancer.

[applause]

Well, it's great to see all of your smiling faces again today. If you remember from yesterday, we talked about the fact that cancer is a genetic disease. It's caused by mutations in our own genes, and I showed you this very compelling example that's just 10 years old of a type of leukemia whereby knowing what that genetic mutation is, we can make a drug that blocks the mutant protein made by that mutation and completely change the lives of these leukemia patients. You had a lethal disease that would have killed them in five years, to now living with a chronic disease by taking a pill every day that has minimal side effects.

So what I want to tell you about today is how that example, and a few other ones, just 10 years old, catalyzed a very important conversation in our community that asked the following question: If all human cancer is caused by mutations of specific genes, and if by knowing the names of those genes

and proteins we can develop new drugs, then it seems incredibly compelling that now, we should know the names of all the cancer genes in all types of cancer.

Now when we started this conversation, I want to remind you that we had stumbled upon the cancer genes by accident, one by one, anecdotal laboratory studies or the animal tumor viruses, such as the Rous sarcoma virus that I told you about yesterday. So here, we're talking about a very different kind of science, something that I would call "big science," in which large groups of people and large amounts of money need to be spent on one project. There was a bit of a debate initially about whether this would be a good idea. The price of DNA sequencing when we first started talking about this was not as cheap, as you've heard from Dr. Walsh, that it is today. But we felt that the potential answers that could come from this and the impact it would have on the lives of patients with cancer was so compelling that we needed to start. And of course, while that conversation was going on, the price of DNA sequencing kept dropping and dropping, so it became much more tangible to imagine this project on a very large scale.

So what are we talking about? What's the actual experiment? The idea is to sequence the DNA from a patient with cancer, sequence the tumor DNA as well as the nontumor DNA, to identify all the mutations that are just found in the tumor DNA and not in the patient's matching normal DNA, and then to filter out the silent mutations and see what's left. So as we saw just a few minutes ago in Dr. Walsh's talk, we're just going to rely on basic molecular biology, the triplet codon, and this example of tyrosine. We want to filter out mutations, such as the one on the left, which lead to a nucleotide change but not an amino acid change. But we want to keep track of the mutations, such as the one in the middle or the one on the right, that change the protein: either through substitution of a new amino acid, in this case, tyrosine to cysteine, or the introduction of a stop codon.

We've also learned, as we've delved into this sequencing a little further, that there are other alterations in tumor DNA that we hadn't previously appreciated and we call these indels, which is short for insertions or deletions. An example of that is shown here, where the normal sequence in the middle differs from the tumor sequence by the deletion of four nucleotides on the top, or the addition of four nucleotides on the bottom. So just adding any nucleotides is going to change things. If you add four, you're going to completely throw off the triplet codon reading frame of a protein that's made by that gene and produce a drastic change. In addition, through whole genome sequencing, we can now discover translocations.

So yesterday, I told you about the Philadelphia chromosome translocation, which was discovered more than 50 years ago by looking at chromosome spreads of tumor cells from patients with that leukemia. And that was possible because in order to look at a chromosome spread, you have to be able to grow the tumor cell in culture and capture it in mitosis. But not very many forms of cancer can be easily grown in that way and analyzed in mitosis. But with whole genome sequencing, you don't have that obstacle. You can determine whether translocations are present just directly from DNA isolated from the tumor, and as we now know, translocations, which we used to think were only found in liquid tumors, like leukemia and lymphoma, are actually found very commonly in solid tumors as well.

So what's the scope of this big project? Well, as of about eight months ago, we had the results from sequencing 4,000 human cancers, from 23 different types of cancer, and some of those are listed here. This is an example of a project that was felt to be so important that everyone needed to work together and so big that it could not be accomplished without everyone working together. And it was really gratifying to see how well people interacted. We actually set up some rules for cooperation as this project started, such that when the sequence data would come off the machine and go through its first pass of computer-based analysis, to get the first look at what the sequence was telling us, it was immediately posted on Internet-accessible data sets, that could be accessed by all the investigators in this project, as well as other scientists interested in looking at the data in real time. So a fantastic example of how scientists can work well together.

So I want to pause for a minute and just get you to think about what are the possible questions that can be asked and answered, once you have this avalanche of data. And one of the first ones that comes to mind is, in each individual patient's tumor, how many mutations actually are there? So for the first time, with this much data, we could ask that question. Yesterday, you saw that in that leukemia, chronic myeloid leukemia, we only knew about one: the BCR-ABL translocation. But we'd never looked in the rest of the genome, so now we have. And what do we see?

So this is a graph where the number of mutations is on the Y axis and I've ordered each of these different kinds of cancers from the one on the far left that has an enormous number of mutations, to those on the far right with the lowest number, and just naturally, by that order, you see a very striking result. And that is, cancers in adults have more mutations than cancers in kids.

So how many mutations are there? Well, in this group that I've blocked here, which represents a large swath of adult cancer, each person's tumor tends to have roughly 50 to 100 different mutations. Now let's think about that for a minute. Does that mean it takes 50 to 100 mutations to cause the cancer? Well, let's look at the kids. The kids' cancers only have about four or five mutations per cancer. So my interpretation of that result is, in fact, it probably … it's not 50 to 100 but more like 4 or 5. Otherwise, why would these kids have developed cancer?

So what does that say about the increased number of mutations in adult cancers? We think a lot of these mutations, although present and somatic, meaning they're in the tumor only—they're not in the patients' germ line—we're now calling those "passengers." We think they are mistakes that happen during the cell division and replication of the DNA that get carried along as the tumor propagates. We now have a very compelling challenge and that is, how do we separate out the real drivers, the real cancer genes, from the passengers? And I'll show you that in a minute.

Now I want to also point out a couple of other interesting patterns that emerge by revealing the code underneath what some of these cancers are. So, in the box on the left, I'm showing you lung cancer that developed in smokers, and in the green box on the right, lung cancer that developed in nonsmokers. And I hope, if there's one thing you leave this room doing, it's if you smoke cigarettes, you're going to stop right now, because look at what you're doing to your lungs.

Now the other one that's high, next to that pink box, is another take-home point. This is melanoma. This is cancer that's caused by exposure to UV light, sunlight. This is all about sun block, so two very preventable forms of cancer. And we know why those environmental hazards, smoking and UV light, cause cancer. They're inducing mutations. So here's the absolute proof of that fact.

So, how are we going to interpret this avalanche of mutation data and figure out which ones are the true causes of cancer? Well, we have to use statistics. That's the best way that we can do it today. There may be other ways, using more biological studies that will help us in the future. So how do we do that? Well, we lay out a map of the entire human genome and one way to do that is shown here.

So let's line up all the chromosomes from 1 to 22, and then, your X or Y on the right. And then, deep into the board, let's just lay out each chromosome from beginning to end. And now, one by one, let's populate this grid with a cancer patient's mutation profile. And as we start to build up that map, you see some patterns. You see, first of all, to get on this map, you have to be seen in more than one patient, in fact, several patients. We want to be confident when we call these genes cancer genes, that they are truly cancer genes. And we assume that if we see them recurrently in more than one patient, they have a higher probability of being causal, as opposed to passengers. And you start to see that there's some real skyscrapers on this plot—the most commonly mutated cancer genes—and there's others that are important cancer genes but less frequently mutated.

So I haven't told you the names of any of these genes. We did this analysis without looking at the names. The next question is, have we ever heard of any of these genes? We've been doing 30 to 40 years of cancer research the old way. Are we discovering anything that we already knew? And the answer, I think, fortunately, is yes. Many of these genes are famous proto-oncogenes, some of which we heard about yesterday. So it's very gratifying that we have used this technology and this strategy … this statistical strategy of assigning a gene the name "cancer gene" successfully. Some people in the field, having spent quite a bit of money—the federal government and other government funds from other countries to do this project—did we learn anything new? We certainly did. We discovered some new cancer genes that no one could have imagined, and I'll give you a few examples. But I think, perhaps, more importantly, we've begun to put some boundaries around the size of the problem that we're facing.

So how big is that problem? Well, as of a data analysis from earlier this year, the number of cancer genes is 140. And the balance, a ratio, amongst those 140, between the dominant oncogenes and what we think are the recessive suppressor genes, is 60 to 80.

Now I don't want you to leave and go home and say guess what? We know there are only 140 cancer genes, because I guarantee you, a paper will be published sometime in the next few months, an update of this, in which it won't be 140. But I promise you, it won't be 1,040. It will be 150, or maybe just south of 200. So we know the size of the problem now. We know it's not exponential. We can really start to envision complete understanding of all the genetic causes of cancer.

So with 140 cancer genes, can we now categorize them into different buckets and see some patterns that are emerging, and the answer is yes. So in this circle, I'm showing you three different categories

that this community of cancer scientists has agreed upon to lump these into. And again, don't hold me to this as only three. I suspect, and I actually even know, that there'll be little tiny additions to this, perhaps new categories that emerge, but they're not going to be dominant causes of major percentages of cancer. So let's look at the 71 cell growth and survival genes. What are these? This is over half of the cancer genes. These are ones that you're familiar with, and many of which we talked about yesterday. These are the proto-oncogenes. These are genes that are involved in signaling pathways in cells that sense signals from outside the cell and regulate the normal growth of various tissues—a classic example being the epidermal growth factor receptor, which we'll talk about in some detail in the second half of the talk. Another group of proteins or genes that is part of the cell growth and survival set, those 71, are the genes that regulate the cell cycle.

So yesterday, we talked about the two genes on the right-hand side, P53 and the retinoblastoma gene. These are tumor suppressor genes which function as brakes on the cell cycle, and you have to have loss of two copies of each in order to get cancer. But on the left-hand side are some examples of additional genes that have come out from this analysis, as well as from other cancer biology studies, that are positive regulators of the cell cycle, that receive signals from the growth factors from outside the cell, and cause cells to enter mitosis. And in those classes of genes, just like we talked about yesterday, you only need one mutation in the gene. Those are dominant.

Okay. So what about this category at the bottom, genome maintenance? There are nine genes that we've put into that category. What does that mean? Well, I didn't go over it when I first showed you but hopefully, some of you saw on this graph that there's this group of cancers on the far left that has a very high number of mutations. If you look carefully at that Y axis, you'll see I've put a little break in it, because these are patients whose tumors have 1,000 mutations per tumor, so an order of magnitude greater.

So how could that possibly happen? Well, there are, as you probably know, when the DNA is replicated in a normal cell, DNA polymerase is a really good enzyme, but it's not perfect and there's mistakes that are made.

About every billion bases, there's a mistake made. And in order to correct the mistakes, we have proofreading enzymes, which go back and read the new sequence and fix the errors. So here's an example of a mistake that was made. There's a G here, which, as you obviously know, should have had a C but there was a T put there by mistake so a proofreading enzyme cruises across the DNA, stops, recognizes that mistake, and fixes it, and moves on and scans the rest of the DNA for other errors. So this proofreading system is pretty darn good, but as I said, one in a billion bases can be mutated and gets missed. Well, imagine what would happen if you had a mutation in the proofreading enzyme itself. Suddenly, you would accumulate mutations at an extremely rapid rate, and that's exactly what happens in these patients with 1,000 or more mutations per tumor. And the types of tumors that have that problem is a form of colon cancer. It's not all colon cancers, but that's where this was first discovered. So repair genes are important in limiting mutations and loss of function, recessive mutations in those repair genes leads to this phenotype.

Now the third category, which is a big category—60 genes as of this update—I've called "cell fate." So what do I mean by cell fate? Well, I'm talking about the process by which a stem cell from a tissue differentiates, gives birth to progeny, progenitor cells, which then further differentiate into mature cells of whatever tissue that is. And eventually, those mature, fully differentiated cells tend to die. Dr. Walsh showed you a terrific example of that in the video that showed the population of the outer cortical neurons from the stem cells, the neural stem cells at the base of the brain. So that's a classic example of a cell fate process in that tissue. But how does a mutation in such a gene that regulates that process cause cancer? Well, it doesn't change the proliferation rate of the cell. It just blocks the ability of the cell to undergo one of those steps in differentiation. So, if there's a block, there's a pileup of cells that are trying to get through that block. And that pileup of cells can form a tumor. And a really nice example of this is also in colon cancer but a different kind of colon cancer than the one that has the high mutation rate.

Colon cancer develops from cells that originate as stem cells in the gut. So this is a view of the inside of the gut, and there's a cross section of that sort of tower, which is known as a villus. And the villus is lined by cells on the outside that are the fully differentiated cells that play an important role in absorbing nutrients whenever you eat. Those fully differentiated cells come from the undifferentiated stem cells that reside at the bottom, where that arrow is, in what is known as a crypt. So if you could run the video now, we will see how this process works.

So we're looking now from the inside at one of these villi, and now we're going to go through the cross-section view and look down at the crypt. This is a video developed from real mouse studies in which you can turn a stem cell blue color through a genetic trick, and then you can watch … every cell that's born from that stem cell will be blue. And you can watch it develop. The cells move up the crypt. They get into the lumen. They differentiate along the way. They'll get up to the top of the villus. They'll then complete their life cycle of absorbing nutrients, and then undergo programmed cell death, or apoptosis. This is also happening in all of our colons right now—very rapid rate. We're still before lunch. I think you're probably okay with that.

So what happens if you have a mutation in the stem cell? There's a common mutation in colon cancer in a gene called APC that's important in cell fate. So the cells that have that mutation are now in that greenish-brown color, they are going to have trouble making it through that final stage of differentiation and they're starting to pile up. And therefore, a tumor starts to develop. So this … this is how colon cancer develops in a large percentage of patients.

So now that we have these three categories—140 genes that are impacted in those 3 categories—how do we think about doing something to help patients? Do we need a different drug for each one of these mutations? The answer may be yes, but we can simplify this a bit further because these different categories can be subdivided. So, in the case of cell growth and survival, those 71 genes fall into different pathways that regulate different signaling pathways or processes in cell death, or apoptosis. Similarly, the cell fate genes fall into different pathways. I know there have been questions about epigenetics. Turns out through this cancer genome sequencing project, for the first time, we found mutations in genes that we call chromatin modifying enzymes. These are the proteins that maintain

the epigenome. So we know that epigenetics is very important, and through this project, we know that mutations in the proteins that control the epigenome occur. I'm presenting this idea of simplifying the problem because it raises the possibility that if we can reduce this into separate pathways, perhaps we can limit the number of drugs that we actually need. So if one of these pathways is the RAS pathway, and we can have mutations in either EGFR, RAS, BRAF, or MEC as the cancer gene, perhaps we could treat all four of those cancers with the same drug, which would be a drug that blocks at the end of the pathway. And we have some pretty good evidence that this is a possibility and … we'll see. So let me pause now for any questions about the first part of the talk and the results from this new cancer genome project. In the tie?

[STUDENT:] Given that we have 140 genes, we know there's … where they are and mutations that develop, would it be possible to study the mechanisms that lead to the mutations that you could target them before the cancer even develops?

[DR. SAWYERS:] So … well, as I said, there's that small category of genome maintenance where actually the problem is, we think the initiating event is a mutation in a gene that leads to many more mutations, so that would perhaps be the best example of one where, if you could stop that process, you'd really shut down the process in its tracks. But I'm not that optimistic about being able to prevent mutations generally, unless we could somehow improve the proofreading ability of the normal … of our normal proofreaders. I don't have a good insight into how we might do that. Here in the middle?

[STUDENT:] Is it only after the final differentiation that the cells are able to go through programmed cell death?

[DR. SAWYERS:] So the question is whether you have to be fully differentiated in order to die, and the answer is, in normal processes, that's exactly what happens. And many tissues do that, like the colon or the blood cells, very rapidly, normally. There's a lot of work trying to see if we could activate programmed cell death in tumor cells, to get them to die when they're not fully differentiated. The challenge there is finding a therapeutic strategy that would make that happen just in the tumor cell but not activate programmed cell death in the normal tissues. In some ways, it's the same challenge that we face with chemotherapy, where yes, tumor cells are killed by chemotherapy but so are normal cells because the strategy is attacking a process that's essential to both tissues. Over here?

[STUDENT:] When you're trying to inhibit a process from occurring and that could stop a whole range of cancers, couldn't that also affect other cells, as you were saying? How are you going to, in fact, keep it from doing that to other cells instead of …

[DR. SAWYERS:] I think what you're asking is … we're trying to be very clever, using this term "targeted therapy." We're making a drug that attacks the root cause of the cancer, but the drug also would inhibit the proto-oncogene form of that. I'll give an example with the epidermal growth factor receptor, where, sort of out of serendipity, the drug is more potent against the mutant form of the epidermal growth factor receptor than the normal form and therefore, we can get away without causing toxicity to the normal tissue, but if we use a higher dose of that drug, we will see exactly what you're talking about.

So I need to move on to the second half of the talk now. And, I want to now address the question, since we've now defined the size of the problem, can we make drugs against these cancer genes? And I want to tell you a story from lung cancer that has happened extremely quickly, beginning in about 2004 to now, and has led to a dramatic change already in the way people are treated with lung cancer throughout the world.

So, how do we diagnose lung cancer? Well, for decades, we've done it by looking at the tumor cells under the microscope, and we traditionally classify lung cancer into four different types, based on the appearance of the tumor cells under the microscope. The largest fraction of lung cancer we call adenocarcinoma. And there are three other types called squamous, large cell, and small cell, and they have the histologic appearances that you see on this slide. But now that we have all this genetic information, what if we put that information on top of this histologic information?

So if we just look at the adenocarcinoma side of this, we can create this pie chart of all the different genetic subtypes of adenocarcinoma. And you can see there are many different types, and up in the top right, you can see this one that we've talked about several times called EGFR, which accounts for roughly 15-20% of all adenocarcinoma. There's another biggie that we'll get back to called KRAS, but then, as you look along the right, there are many smaller pieces of the pie. But fortunately, they have gene names that we recognize. Almost all of the genes on this list are kinases. Same class of cancer oncoproteins as the BCR-ABL protein, as I showed you yesterday. We learned from that example that we can make these competitive antagonists of ATP binding and inhibit kinases.

So that has been a tremendous set of activity in this field, to make inhibitors of each of these different kinases. And as of today, we have FDA-approved drugs for all of the ones that are still in bold. So how do these genes cause the lung cancer and how do these inhibitors work? So you've seen the signaling pathway before. Everything circled in red is an example of a protein in that normal signaling pathway for which we have examples of lung cancer patients who have mutations in that gene. The normal lung cell will see growth factors outside the cell, such as EGF, which will bind the receptor and initiate this signaling cascade, down these different branches of the pathway, and lead to cell proliferation in a very regulated and controlled fashion. In patients who have a mutation in the EGF receptor, this signaling pathway is turned on at a much higher level. Imagine turning up the volume. Still requires the EGF outside the cell, to a little bit of a degree, but the volume is on much louder, and the proliferation continues and the cell becomes a tumor.

So what happens if we give one of these EGF receptor kinase inhibitors? It binds the mutant EGF receptor and as the last questioner asked, it actually binds this mutant receptor much tighter than it does the normal one on the left, so that lets us get away with a better therapeutic index. After the drug binds, the signaling pathway is shut off. It gradually decays over a matter of hours. And there we have it.

So what happens when you now move to a clinical trial of this drug, which, so elegantly designed to attack this problem, is it still going to work in a patient? Well, the first clinical trial of the first EGFR inhibitor was conducted by some of my colleagues in Memorial Sloan-Kettering in New York. Just over

10 years ago was one of the first patients and this is a chest x-ray of a patient with metastatic lung adenocarcinoma who happens to have this mutant EGFR in her tumor. And you can see … I know you're not used to looking at chest x-rays, but let me just outline in red, this is the lung fields on the left and the right. They normally should be quite clear and all that fluffy white infiltrate that you see—that's lung cancer.

So this patient is quite sick, having trouble breathing, because obviously you need clear lung fields to oxygenate your blood. Within one week of taking this pill once a day, this is what the chest x-ray of this patient looked like. Just a very dramatic, complete clearing of the tumor cells by taking this one inhibitor. This drug is now the standard treatment for patients with EGFR-mutant lung cancer.

So how well does it work? Well we now … I showed you yesterday how we track the progress of patients taking a leukemia drug by looking at their blood counts. Well, patients with lung cancer, what we do, is we either look at their chest x-ray or we do CT scans to look at the size of the tumor. And then we track how big the tumor is, how much does it shrink over time. And we show that data in the form of these graphs, where each one of these bars is a patient and if the bar is above zero, it means that the tumor grew during the treatment period. If it's below zero, it means it shrank. We actually now call these plots "waterfall plots" because we think it looks a little bit like a waterfall.

So you look at this and I can tell you, having seen waterfall plots from the chemotherapy era for these patients, this is a complete game-changer. The lives of these patients has been transformed in a way almost as dramatically as the leukemia patients I told you about yesterday. But you can also see that not everyone benefits, and not everyone benefits to the same degree.

So why are these patients not responding and why are the ones in the middle not going all the way down to -100, which would be no detectable tumor? Well if you remember, from this plot, we learned that cancers don't have just one mutation. They have, at a minimum, in children, three to four. And in these patients—these mutations actually occur in nonsmokers—at least 10 to 20. I think this probably explains an interesting result that Dr. Walsh presented yesterday. When you remember the film about Dante, there was a mutation … Patients with that form of hemimegalencephaly have mutations in the AKT3 gene. That's a cancer gene, but he did not get cancer. We think that you need multiple mutations to get cancer. Therefore, we may need multiple drugs to inhibit the additional mutated genes to make sure all of the patients have a complete remission.

So let's go back to that pie chart. If you remember, this is a pie chart of adenocarcinoma of the lung, and then looking at what percentage of adenocarcinoma of the lung is caused by different mutations. I've circled another one there. It's not a common mutation in adenocarcinoma of the lung, but it's a gene called BRAF, which is one of the genes in that signaling pathway. Now it turns out that BRAF is one of many types of cancer genes that is not just mutated in lung cancer but in many different kinds of cancer.

So let's ask the question in a different way. Instead of thinking about diagnosing the cancer based on where it came from, what if we flip the order and diagnose the cancer based on what the mutation was? So I'm now flipping the pie chart the other way and showing you, of all the cancers that have

BRAF mutation, how do they segregate? Well, it turns out that roughly half of patients with melanoma and thyroid cancer have BRAF mutations, and smaller fractions of colon, lung and prostate cancer. But it's very interesting that this one gene, and in fact, it's the exact same codon change in all of these cancers, is a cause of that cancer. Turns out, we have a very good drug that inhibits BRAF. It's also a kinase and this is a kinase inhibitor.

So how well does it work? Well, we've tested it in lung cancer and this is what the waterfall plot in one of the early trials for lung cancer looked like. Very dramatic result, like the EGFR result, very exciting, transforming the lives of these patients. What about in melanoma? Same thing. It's now approved for melanoma and has changed the lives of these patients. What about colon cancer? Well, … we tried very hard. We've given the right dose. We've made sure patients were taking their drug. The levels were perfect. No responses at all. So what's going on?

So let's go back to our signaling pathway diagrams and see if we can figure it out. So I showed you earlier that when the mutation is in the epidermal growth factor receptor, it starts up at the top of the cell and the signal is propagated down. This time, the mutation is in BRAF, a little further downstream and this is what the signal looks like. You give the inhibitor. The signal decays, and then the cell dies. Proliferation is shut off. What about melanoma? Same thing. The signaling pathway looks like it's wired exactly the same way. Cell is proliferating. We give the drug. The signal decays. Proliferation stops and the cell dies. Let's look at the colon. Well, at first glance, from laboratory studies, it looked like the wiring was exactly the same, but as we've learned from playing around with these cell lines from these colon cancer patients, when we give the inhibitor, we learned that there was something we had actually neglected. It turns out, in colon cells, there's something called a negative feedback loop. So for reasons that presumably are important in maintaining the normal physiology of the intestinal cells, but aren't so important in the normal physiology of the lung cells, evolution has selected for this feedback group.

So why am I telling you this? Well, let's think about what happens now when you give an inhibitor of BRAF. Just like we saw in the lung and the melanoma, the downstream signal decays. But what's going to happen to this feedback loop? Just as that downstream signal has decayed, the feedback loop, the negative feedback loop, is going to go away and now, the number of copies of the EGF receptor is going to rapidly increase. Well, EGF happens to be floating around in the colon all the time. And so as a result, it can bind this new population of EGF receptors, and the signal can be reinitiated. Now the part of the signal that's going down the RAS pathway is still blocked when it gets to BRAF because our BRAF inhibitor is still there. But there's a parallel pathway on the left-hand side that can stimulate growth and proliferation and cause the tumor to keep growing.

So what would you do next? Well, I showed you that we actually have a drug that inhibits the EGF receptor. What if we gave these two drugs together? And just in the last six months, the clinical trials of this experiment have been done and are still underway. And when you give both drugs, first of all, in the lab, you can shut off the signal. And in the patient, you can now cause the waterfall plots to go in the right direction. So very, very exciting that we can overcome that problem, but a very important lesson about the question that I was raising and that is, can genetics replace the way we diagnose

cancer? I think the answer is no. I think we need to know both things. We need to know where did the tumor come from and what is its genetic profile and that will help guide us to the best combination of drug therapy.

So where do we stand? We have 140 cancer genes in 12 pathways. How many drugs do we have today? I think it's a difficult answer. We have 38 approved targeted therapies today when we only had 3 or 4 ten years ago. But what percentage or fraction of cancers can we actually treat with these drugs? It's not an easy question to answer, but I'll go back to lung cancer to point out the other side of this pie chart.

So I've gotten you excited about the success stories on the right half of this pie chart. The ones in bold here are also kinases. We don't have approved drugs for these kinases but they're in clinical trials and I'm reasonably confident that we can put those in the taken-care-of category over the next couple of years. Of course, the gray category is a frustrating one. We actually do not know what's driving these lung cancers, and this raises some of the questions that Chris's talk raised about, maybe there's epigenetics playing a role here, or other causes, … or maybe it's multigenic. But KRAS is a very important one. It's arguably one of the most common drivers of this form of lung cancer and like BRAF, it's an example of a cancer gene that's involved in multiple different kinds of cancer. Some would argue … it's one of those skyscrapers on that plot. So we desperately could use a KRAS drug. But KRAS is not a kinase. It's considered an undruggable target. So there's been a lot of frustration in the field and even neglect of this as a target because there's not a tractable strategy to inhibit it.

So what is KRAS? Just to remind you that in the signaling pathways, it's just one step upstream of BRAF. It causes lung cancer, shown on the right, by driving the BRAF pathway, but also, it can talk to the left-hand side and drive the other half of the pathway. So we can't just use a BRAF inhibitor and take care of it. So the fact that it's undruggable to me is a lame excuse. If you remember yesterday, I told you, when we started talking 15 years ago about making inhibitor of BCR-ABL, most of the field told us it's a kinase, it's undruggable. But I think you've seen that that's not true. So in the same way, I don't think we should give up on KRAS. In fact, I'm extremely excited to tell you that just two weeks ago, one of my colleagues, another HHMI investigator, published a paper in *Nature* showing … identifying an inhibitor of one mutant form of KRAS that looks very good in the laboratory and in tissue culture studies, through a very clever, novel chemical approach, but it just points out that clever, curious scientists need to keep pushing on this problem.

So I want to close with a completely different strategy for treating cancer that takes advantage of all of the mutations I've told you about. And it has to do with using the immune system to our advantage. So we all know that the purpose of the immune system is to protect us from foreign invaders, from viruses and bacteria and so forth. What I'm showing you here on the top is a normal cell that gets infected by a virus. The viral proteins are now made in that cell. The virus replicates itself but the viral proteins that are made get degraded into peptides and they actually get presented on the surface of that infected cell. And that's one of the ways the immune system recognizes the infected cell and can eliminate it. I want you to just now imagine, on the bottom here, this is a cancer cell which has a BRAF mutation. That BRAF mutation creates a mutant protein and therefore, when BRAF is degraded

normally into peptides, there's going to be a mutant peptide. So this is what that would look like and just like the cell infected by the virus, the cancer cell has the same protein machinery to present these peptides on the cell surface. So on the top, you've got a normal cell infected by virus and viral protein presented on the outside. On the bottom, you've got a cancer cell with a mutant peptide, the mutant piece of BRAF on the outside of the cell.

So the question is, can the immune system see these things? And the answer is yes. T cells in the immune system are highly specialized to very specifically recognize these mutant peptides through the T-cell receptor, and when they make this match, they bind. In the case of the virus-infected cell, that T-cell receptor match leads to the stimulation of the T cell. It grows and grows and eventually kills the virus-infected cell and the infection goes away. Well, does that happen in a cancer cell that's presenting a mutant peptide? Well, we actually know the answer because cancer grows in patients. For some reason, the T cell is not able to take care of it. We actually know, in patients, that the T cell can make the connection, make the recognition, but they don't kill the cell. Well, there's part of the story on the top that I left out and that is that the immune system, just like the cell cycle, has a go signal—the T cell receptor. It also has a brake called a checkpoint. And we need that break because if we don't turn off the immune system, the hyperactive immune system in a patient with an infection can cause damage. So there's this set of molecules known as immune checkpoint molecules that have to be engaged in order to turn off the immune system. So that engagement happens normally in the setting of an infection. The immune system goes down and everything is okay. So in a very nefarious way, cancer cells have learned how to activate that checkpoint at the beginning, so they engage this immune checkpoint molecule and stop the T cell in its tracks before it can proliferate and destroy it.

So immunologists have figured out the molecular biology of this recognition system with great precision. In the same way that I told you about making inhibitors of kinases, you can make inhibitors of these checkpoint molecules. And just in the last year, one of those molecules was approved for the treatment of cancer based on this strategy. So what happens here is that the immune system is held at bay. What if you gave a checkpoint inhibitor? Can you increase the number of T cells, and can you then kill the cancer cell? So the very dramatic demonstration of this is with this example of melanoma. So this is a patient with melanoma that's metastasized to the liver. And you can see all those gray circles in the liver. Those are tumor deposits. In fact this is a CAT scan that was taken 12 weeks after the patient began the therapy with the checkpoint inhibitor, and you might say, why are you showing me this slide? Because the patient has been on the drug for three months and nothing good has happened. With the kinase inhibitor examples, we see those effects in weeks, but the immune system, it takes time for it to rev up, and what's really exciting is if you now look at 36 weeks, all of the tumor cells in this patient's liver are gone. Really amazing result. This drug has been tested in hundreds of patients now. It has a waterfall plot in melanoma that looks like this, and is now an FDA-approved drug.

So I want to finish with the idea that we've clearly moved to a very exciting new era of treatment of cancer. We still use surgery and chemotherapy and radiation, but the targeted drug therapy, the examples I told you about yesterday in leukemia, and today in lung cancer, melanoma, and colon cancer, are really proliferating at a fast rate. But just in the last couple of years, this new idea of immunotherapy, teaching the immune system to wake up and recognize the cancer cells as foreign …

there's definite proof of concept, and there's early data now coming out that the tumors that respond the best are the ones, ironically, that have the most mutations, because they present the most mutant peptides to the immune system.

So, in finishing, we've entered a new era of understanding of cancer. This big, large cancer genome project has put the boundaries on the size of the problem. At a minimum, we now have a much better, more precise way to diagnose cancer, a combination of the tissue diagnosis and the genetic diagnosis. The genetics is also teaching us, as I told you yesterday, how to anticipate and understand resistance to these drugs and develop rational combination therapies. And what we really need now is more drugs and treatment options, which we're going to count on for you to do. Thank you.

[applause].

[DR. SAWYERS:] So, in the back.

[STUDENT:] My question is that, would this be an acceptable alternative to replacing skin grafts in the patients who use that as their treatment for cancer?

[DR. SAWYERS:] Use grafts? You mean, like bone marrow transplants or something?

[STUDENT:] Yeah.

[DR. SAWYERS:] Yeah. So, bone marrow transplant, which is using the immune system of a donor to fight off some of the cancer, has worked. In fact, ironically, the best example of that is chronic myeloid leukemia. Which was the reason that was a successful therapy in the past. But the toxicity of that treatment is enormous and the percent of patients who are actually eligible to tolerate that kind of toxicity is quite low. You have to be quite young to undergo it. There is another form of immune therapy that's not yet approved but there's evidence of it working, which is a form of gene therapy, where you take the T cells out of the patient, and you engineer a gene that will express what's called a chimeric T-cell receptor. The outside part of the T-cell receptor is re-engineered to recognize a protein on the surface of a tumor cell but the inside is the same T cell receptor. So when that cell sees the tumor, it proliferates and kills the cell. It's worked in a form of leukemia and we'll see if it can work against other cancers. Over here, yeah.

[STUDENT:] I was wondering, what were the effects, on like culture and just people with cancer in general who aren't a part of these clinical trials, with the mass production and usage of these drugs that you talked about today?

[DR. SAWYERS:] That's a fantastic question. I have to say that the excitement of this is … it happens first at research centers and then gradually rolls out in to the larger community. One of the first challenges is the technology to actually sequence the DNA of the tumor and to perform a diagnosis in an FDA-approved fashion. It's taking time, much longer than I have the patience for, but there is important regulatory hurdles that need to be addressed for important reasons. We don't want sloppy testing done. In the US, however, most patients with lung cancer who have one of those mutations I think are getting the right drug. In fact, the FDA has made it mandatory for a drug company that makes

the drug, in order to get approval to use the drug, they have to get a diagnostic test up and running that's also approved. That's called a companion diagnostic, and I think that's a very good way to make sure that these drugs get out quickly and safely to the larger population. Here in the front.

[STUDENT:] Do mutations occur more frequently in stem cells than others? If so, is it more dangerous to target those stem cells with targeted drug therapy than to treat other non-stem cells?

[DR. SAWYERS:] That's a fascinating question. The answer is, we don't know for sure, in most cancers, whether the mutation is in the stem cell or a progenitor, but we know in selected cases, it can be in either place. The mutation does not have to be in a stem cell. In prostate cancer, in fact, it's in a almost-differentiated cell, or almost fully differentiated cell, and what the mutation can do is actually make the cell sort of reprogram and go a little bit backwards in differentiation, and then become a tumor. Now, the second half of your question is, if the mutation is in the stem cell and we have a drug against that, I think what you're asking is, will we wipe out the normal stem cells? And the answer is no, because the mutation is only present in the stem cell of the tumor and not in the normal stem cell. Yeah, behind you, in the glasses, yeah.

[STUDENT:] In immune therapy, if the checkpoint is inhibited, how do T cells stop when it is necessary?

[DR. SAWYERS:] That's a great question, so … such a brilliant strategy to disable the brakes of the T cell, and get the remission in the cancer patient, but … aren't the brakes there for a reason? Well in fact, the major side effects of that drug is problems with the colon, and problems with some other tissues, that are in a category of diseases we call autoimmune disease, something sort of like ulcerative colitis or Crohn's disease that you may have heard of, that are caused by an overactive immune system, an immune system that is making a mistake, sort of a proofreading problem. It sees normal tissue as foreign and inappropriately attacks. So it's not a horrible side effect in these patients. It's treatable. I mean, there's some newer checkpoint inhibitors that don't seem to show that side effect as dramatically. Alright. Thank you very much.

[applause]

[music plays]