nytimes.com

From Gene Editing to A.I., How Will Technology Transform Humanity?

Siddhartha Mukherjee

40-50 minutes

“A geneticist, an oncologist, a roboticist, a novelist and an A.I. researcher walk into a bar.” That could be the setup for a very bad joke — or a tremendously fascinating conversation. Fortunately for us, it was the latter. On a blustery evening in late September, in a private room at a bar near Times Square, the magazine gathered five brilliant scientists and thinkers around a table for a three-hour dinner. In the (edited) transcript below — moderated by Mark Jannot, a story editor at the magazine and a former editor in chief of Popular Science — you can see what they had to say about the future of medicine, health care and humanity.

MARK JANNOT: For years, many pregnant women have undergone amniocentesis to test for rare metabolic disorders and other fetal issues. And couples who use in vitro fertilization can screen the embryos for genetic abnormalities. What sorts of advances in genetic screening and manipulation are coming, and where do you see that taking us?

CATHERINE MOHR: When I was pregnant with my daughter, my husband and I were joking, “Well, if she gets the best of both of us, she’ll be a superhero, and if she gets the worst of

both of us, she's not going to make it out of first grade." And so we were rolling the genetic dice, which you do when you choose to have a child. It's not totally random, of course; there's all kinds of great things about your mate — that's why you chose them — and hopefully there's some pretty good things about you, too. But the temptation to engineer what you think of as the best combination, as we become more capable of doing it, I think it's going to be irresistible for a lot of people. You're investing so much of your life into this little being, and you're going to love this child, and you want to give them every advantage in life. We are already screening for diseases to avoid passing on our "bad" genes, but this same technology will let us start screening for our "best" genes — the ones we really want to pass on. As screening becomes cheaper, easier and more reliable, and more people are using assisted-reproductive technologies, I see us, as a society, sliding down that slippery slope pretty far, one couple at a time, each trying to do what's best for the child they are hoping to bring into the world.

SIDDHARTHA MUKHERJEE: It's certainly a tempting path, toward a potentially terrifying slope. But that only works if you do in vitro fertilization and create a pool of testable embryos. Then you have to biopsy those embryos-in-dishes, sequence their genes, identify and interpret the gene variants that you want to select (Variant A and B and C and D) and implant the "desirable" ones.

GEORGE CHURCH: Or we may turn to gene editing. If, for example, you have a dominant-allele disorder, like Huntington's disease or Marfan syndrome, and you want to have children, you could edit the sperm, change that allele so

that all sperm are healthy and your offspring will be fine. All sperm come from spermatogonial stem cells in the man's testes. You can use editing tools and work on stem cells in Petri dishes so that you're removing the bad allele and replacing it with DNA that has been designed and synthesized on computer-controlled machines. And then you can implant a pure population in which you've checked that the edit is what you wanted it to be, with all cells with only the desired "on target" changes. This has been done in mice. It's a great opportunity. It's only one time, and they're good for life. In principle.

JANNOT: And why is that not being done now?

CHURCH: Until recently, we didn't have good methods for doing gene therapy that we could apply to editing stem cells, sperm cells.

JENNIFER EGAN: How hard is it to edit genes?

MUKHERJEE: Well, that's one of the surprises, is how extraordinarily easy it is. There are still technical challenges, and some of them may be hard to surmount, but the protocol is quite simple. We recently edited a gene in human blood stem cells to enable therapy for some forms of leukemia. We've sequenced the genomes of the edited cells and have not found a single "off target" effect thus far, although we are still looking. For other genes, off-target effects have been reported, so it seems that it's case dependent. But over all, the fidelity of the system seems quite remarkable.

CHURCH: At this point, there's nothing published in the literature demonstrating successful editing of human sperm stem cells, the germline. But if you want to edit the DNA of,

say, pigs, it's very easy with Crispr, which is a set of editing tools that uses enzymes, guided by RNA and proteins, to make a change at a precise location in your DNA. You're injecting a small thing in that changes as little as one base pair out of six billion, in each cell. So it's nanosurgery — very precise and automatically in many cells at once.

MUKHERJEE: It's like taking a massive encyclopedia and saying: Go to Volume 7, Section 8, Page 240, Paragraph 5, and change the word "this" to the word "that." I'm simplifying, of course.

MOHR: And to use your encyclopedia analogy, everyone who is unlucky enough to have their edition of the encyclopedia printed with "this" gets sickle-cell anemia, and everyone whose edition has "that" doesn't. But, George, while you are saying we can't quite do gene editing of the germline cells for producing genetic-disease-free children, editing genes in the adult — gene therapies aimed at altering all of the mature cells in an already-formed organ or a cancer — you're saying that's closer?

CHURCH: Some gene therapies involve adding missing genes, others involve subtracting toxic versions of genes and some involve precise editing. And yes, it's getting closer; there are some gene therapies that are already approved for human use.

MUKHERJEE: At least one that is approved is for retina diseases. Not gene editing — changing the native genes in the genome — but introducing new genetic material into human cells. That's because introducing viruses carrying new genetic material into the eye is easier. You can inject viruses

because the immune system does not seem to be as active in the retina, and the injected virus doesn't spread all over the body.

But the ones that involve gene editing are on their way, they're in the pipeline. There's a lab at Stanford that's doing gene editing on blood stem cells for sickle-cell disease. Then you can transplant those blood cells and replace the diseased cells, and the sickle-cell disease should be cured. We'll get comfortable with it, and by comfort I mean not just becoming comfortable with technically how to do it, but realizing it doesn't all of a sudden cause some horrible cancer, or some terrible disease, which, if you ask me, I think is quite unlikely. But at some point the decision will come down to the F.D.A. and other organizations; they'll have to say, let's go forward. Bottom line, our capacity to become more comfortable with the consequences of gene editing will come from diseases where the stakes, as it were, are more simple and higher — especially with a disease like acute myeloid leukemia, where there's an extremely high mortality rate — and then we'll backtrack our way into reproductive technology.

CHURCH: I think it's more likely we'll be using gene therapy first in childhood diseases, based on the realization that many diseases make permanent damage by the time the child is born. Like blindness, for example — if you don't correct it very early in life, you can "cure" blindness in the sense that they can see photons, but they can't really process them into an image.

MOHR: Blindness is an interesting one in this context. It isn't life-threatening like the leukemias Sid was talking about, but the problem is an absence of function, which seems in some

ways less risky to tackle. If you fail trying to fix it, you haven't made it worse — the person is still blind — but if you do succeed, there is only upside. I can imagine these are the kinds of deficits we'd be most eager to try to address because of the way we as people think about risk: We're O.K. with risking that things will get better, but not too happy doing it when there is a chance they'll get worse.

CHURCH: Then there are the diseases that won't affect people until late in life, but they could be treated with gene therapy very early in life. This may be the case with Alzheimer's. We already know that the alleles that are highly associated with Alzheimer's are something called APP, for early-stage Alzheimer's, and the ApoE e4 variant, for late-stage. We could change them in the sperm cell to an allele that already exists in the population. And you're changing it essentially 100 percent because it's going through this bottleneck of a single stem cell. And you're not trying to change it to a gene that no one's ever tested before; it's a gene that's been "tested" millions of times in the millions of people in whom it occurs naturally.

MOHR: So, in the same way that a woman might take folate before and during pregnancy to prevent neural tube defects in the fetus, you'd have your partner take the gene therapy to do some allele substitution. "O.K., honey, I love a lot about you, but we're going to need to edit out that cystic fibrosis variant and tweak those Alzheimer's alleles of yours before we start thinking about kids."

EGAN: Speaking as someone who is terrified of Alzheimer's, engineering it away is an appealing prospect. But I wonder: Who exactly would have access to this technology? Even

basic reproductive technologies like I.V.F. are expensive, so less possible for poor people. One unintended consequence, it seems to me, could be a small number of extremely healthy genetically engineered elites and a large and comparatively ill and genetically challenged underclass.

CHURCH: But all of these technologies are constantly getting cheaper — look at what happened with the cost of sequencing the genome, from billions when we first did it to a few hundred dollars today. I think these therapies would end up similar to preventive medicines like vaccines. Vaccines are enhancement relative to our ancestors, and they've been able to be made ubiquitous. Our ancestors lived in mortal fear of all these diseases, and we just take it for granted that we're immune to them.

EGAN: I'm struck by the tremendous confidence with which you talk about these things, almost as if they had already happened. You're thinking forward to a point when all of this will be a matter of course, but I'm still back at the point where it all sounds so speculative. I find myself thinking, Whoa, what about operator error? I mean, nothing technical works simply or perfectly, ever. And yet so much of what we take for granted now — flying in airplanes, for example — would have struck me as equally hubristic in the planning stages. And of course it is catastrophic when a plane crashes, but that's an extreme rarity.

REGINA BARZILAY: We're working with a complex system that we are only beginning to understand today. It's well known from selective breeding of domestic animals that selecting for one target trait often brings along many other undesirable and often unexpected traits. Let's say you guys

identified a genetic fix to a problem. How likely is it that changing “this” to “that,” following your analogy, is going to bring some other, unexpected side effects that we cannot control?

CHURCH: Well, in some of these things, you’re literally changing a gene to what is healthy. For instance, in the case of sickle cell, changing a particular gene variant to what everybody else has is probably pretty safe as long as you can be sure that’s what’s actually happening. So the probability of unexpected consequences seems quite low. Once we go forward, as we get more and more confidence, we will start taking bigger and bigger steps; then we might end up with something that has unintended consequences. You know, eliminating smallpox from the entire world could have had negative consequences. We rolled the dice and figured that we could back up if there were some problem. To think that genetics is irreversible is no more likely than that eradicating smallpox is irreversible.

JANNOT: What are the most interesting applications for A.I. in medicine right now?

BARZILAY: This is a great question. Companies like Google and Facebook track every action you take online and use that to build a model of your preferences. They then use this model to personalize the complete user experience, the content you see, the products they recommend to you, the advertisements they show you. In some ways they know more about you than you know about yourself. But if you go to any clinic, for cancer, heart disease, you name it — there is no A.I.

I learned this in a very personal way. When I was 43, I went in

for a routine mammogram, and all of a sudden I was diagnosed with breast cancer. This was a big shock because, to the best of my knowledge, nobody in my family had ever been diagnosed with cancer. At every point in my treatment, I had many more questions than my doctors had answers to. I remember I did my mammogram, and they said, “Your cancer is really tiny.” I said “Great!” Then we went to M.R.I., and suddenly they see cancer all over. Then they did a biopsy, and they discovered it’s actually small; the M.R.I. was a false positive. How can we have this high-resolution M.R.I. modality and still not know that this is a false positive?

For me as a computer scientist working in artificial intelligence, it seemed obvious to train a machine to make these kinds of predictions. If you look at what was happening in computer vision, A.I. systems could already identify very subtle distinctions between images, at a level of detail that’s hard for the human eye to differentiate. Why do people need to undergo unnecessary procedures and live with months of uncertainty while the technology that can fully resolve the situation already exists?

And this was just one of many steps in the treatment pipeline where I saw how artificial intelligence could transform cancer diagnosis and treatment. As an A.I. researcher, I was stunned to see all these opportunities to help patients squandered. From a patient’s perspective, it felt cruel. We’re talking about well-understood technology commercially deployed in other industries, not brand-new research. And this is a general trend. It doesn’t matter what your disease is; today, A.I. is not yet part of clinical treatment.

MOHR: This is a problem that really affects providers also —

patients' medical data are kept in all of these separate systems, so it's hard to get all the data about even one patient if there are multiple doctors involved in the care, let alone being able to compare the data on many different patients. It evolved this way because we used to have paper records with narrative descriptions of each patient's condition, and our privacy laws never anticipated the tools we would have today — and what we could do with the data.

JANNOT: So what needs to happen?

MOHR: Revamping our practices and regulations around medical data while maintaining individual privacy will be essential both for patients like Regina and for A.I. researchers like Regina. It's likely to be slow, but it is starting.

BARZILAY: For my part, when I finally came back to my work at M.I.T., my experience as a cancer patient had totally changed my perspective, and I could not just go back to my old research. I started asking: What is the best way to spend my time, my mental energy? I could not forget the suffering and pain I saw in the hospital. I wanted to use data to provide answers now. It took me a while to find like-minded clinical collaborators and zoom in on specific questions that were meaningful to me but also could be implemented in the clinic.

Ultimately that brought me to two areas. One of them relates to something very basic in clinical research — extracting relevant information from patients' electronic records. Even though every hospital sits on a gold mine of data, it's severely underutilized by care providers and clinical researchers, because the records are mostly in text. Unless they're specifically trained, machines cannot read these stories; they

expect a database where information is properly structured. And so, today, if you as a patient want to know how patients like yourself responded to treatment in your hospital, you can't find the answer. Even in the most prestigious journals, almost all the studies that use past patient data do that data extraction by hand, which is expensive and slow and dramatically limits the scope of these studies.

In my core field of research, natural language processing, we've developed lots of tools that can automate this task. And so we applied those tools to create a database of more than 100,000 patients with breast disease from Massachusetts General and other partner hospitals that spans decades. Now with one simple query you can find a cohort of patients with the same disease features and study it over time.

Another thing I'm working on relates to reading mammograms. Today the risk models used in clinical practice are very imprecise. Our ability to predict who is going to get cancer is very, very low. Our idea was to let the machine algorithm look for patterns in the raw mammographic image: If it looks at the mammogram, from five years earlier, of a woman who went on to develop cancer, can it detect patterns?

The first step was to work with Connie Lehman, head of breast-cancer radiology at M.G.H., to use radiologists' best judgment to train the model. And that did improve the predictive results, but we felt that it didn't fully reach the goal. We wanted the machine to utilize *all* the information in the image, not just the things that radiologists are trained to spot as disease markers. We trained the machine to look at the whole image, and we fed in all the data about outcomes, and we said: What is the likelihood that this person is going to get

cancer in a certain time? This system worked way, way better than any risk models currently in clinical practice.

We are now thinking of expanding our work to prescreen for lung and pancreatic cancer. Imagine how it can change the game if these diseases, which are now diagnosed late, when they are largely uncurable, could be detected early — how many lives can be saved. That is the way that A.I. can transform medicine. It will identify patterns far too subtle for humans to identify.

MOHR: Regina is talking about a very specific kind of A.I. — machine learning and natural language processing, rather than what we think of in popular culture, robots in the movies who walk and talk and crack jokes. We'll have lots of beautiful analysis capability like Regina is talking about long before we have C-3PO.

In surgery, we're also starting to use the same sorts of tools that Regina is applying to radiology images and natural language analysis of medical records, but we're doing it with surgical videos and data from operations, data that we can readily harvest from surgical robots. These are machines that surgeons operate as extensions of themselves, enabling them to perform extremely delicate surgeries, through small incisions, and watch what's going on inside the patient's body via a video feed. They can actually see better than if they had cut the patient open. And the machine records every movement made and captures that video of the operation.

It is amazing how much a trained human can tell from just looking at a single frame of a surgical procedure. A well-trained surgical resident can walk into an operating room

where a surgery is underway, and can glance up and with one look at the screen know what kind of procedure it is, what step you are at in the procedure — they know what's going to happen next, and they can tell if it's going well or not, using clues like if you've got a lot of blood in the field, or from looking at the body language of all the people in the operating room. Is the surgeon stressed out? Has the music been turned down? Are people still talking? What are they saying? There's all kinds of clues.

We can use the data in those videos, use machine learning and natural language processing to train an A.I. to be able to pick up on all these same clues and to be able recognize the same things the resident can, and then ideally to be able to help you with what might be the best next step. It would be like providing every surgeon with the perfect surgical resident.

To achieve this, it isn't just recognizing what is in the picture or the sounds; these algorithms need to understand the context, where you are in the procedure, what's going to happen and what should ordinarily happen next. To do all that, we need to train them on a lot of data, looking at how a thousand different surgeons do exactly that same step, and what best practices are, and maybe clustered into five different styles of doing this particular surgery so you can tell which step to recommend next. The key is that by turning surgery into data, we can now start to use these remarkably powerful machine-learning tools to analyze and learn from these data. But first you need data. We're lucky with our robots, but in many areas of medicine it is hard to get your hands on the kind of data you need.

JANNOT: So, George, as you mentioned earlier, we've seen exponential decreases in the cost of sequencing a genome. I

imagine cheap genome sequencing leads to ubiquitous genome sequencing, which leads to a superabundant new stream of data to plumb for insights and new health advances.

CHURCH: That's right. We've gone from it costing almost $3 billion for a clinically unacceptable genome in 2004 to less than $1,000 in 2015 for a high-quality genome that precisely analyzes the DNA you inherited from your mother and father. I just started a company called Nebula Genomics, whose intention is to make it zero dollars or less. At this point everyone should be getting paid to sequence their genomes. Because the system could save something on the order of a million dollars every time we save a single child from a rare genetic disease. That million dollars should then be spread out to all people who participated, including the 95 percent of people who didn't get any bad news.

MUKHERJEE: In terms of what will drive future advances, there is the whole aspect of the genome, and then there's the whole aspect of what people have called the phenome — things that we do, things that we express, environmental things that happen to us, how we interact with the environment. Both are data sets. One of them is now a highly accessible data set, and with Nebula it will become a zero-dollar data set. The other one is not a zero-dollar data set, yet. But very soon you can imagine carrying some kind of GoPro, in which data becomes so cheap that you can start really monitoring that second data set, what you do, what you eat, whether you run, how much you run, the number of Fitbit steps, etc. Imagine the density of individuated information that comes from all this.

One implication is that 25, 50, 250 years from now, we

become a kind of clinical-trial society in which empirically driven decisions are constantly popping up. But by clinical-trial society, I mean all sorts of questions, because the information net becomes so rich — and the capacity to understand or deconvolute that information, because of computational power and because of A.I.-dependent algorithms, becomes so rich — that we begin to subject aspects of human behavior, human selves, that were previously considered outside the realm of assessment to a kind of deeper clinical assessment.

MOHR: The natural extension of that is, we have some kind of personal doomsday clock. And each action that we take is either extending it or decrementing it. So, I put something bad in my mouth and I start to eat it, and I see that that dropped my doomsday clock a little bit. I go out for a run and see that it bumps my doomsday clock up a little bit — I can see the immediate projected effect of all of the actions I take. If we could measure all of those things, people would be carrying their doomsday-clock algorithms around.

EGAN: What about privacy? If every fact about my body can be known, and if my knowledge of those facts depends on corporations helping me to track and measure the data, I will not be able to control whose hands that information falls into. As to what we do and think and express, social media is already quantifying our behavior, in exchange for giving us a platform and access. We pay a price for opening ourselves to corporate data systems in exchange for information; ultimately, anyone will be able to know anything about anyone, and that's a vulnerability.

MOHR: Privacy is at the heart of the problem around availability of medical data for training the machine-learning

algorithms that we were talking about earlier. Those of us who look at the data and see all the good it could do have a hard time imagining hurting people with that same data, and yet the possibility exists that the very things that teach us how to help people who have a condition will allow others to discriminate against them or victimize them because of that condition. These are hard problems, but we should try to figure out how to get the greatest societal good out of this data without putting those who donate it at risk — the benefit to us all is so potentially great. To shy away from it because it is "hard to do" has victims, too — someone who dies when we didn't know how to help them, knowledge that would have been available if we had been able to pool our data — that person is worth figuring out how to save. We're already figuring this out first in the diseases like cancer because patients are very motivated to share their data.

MUKHERJEE: Yes, and it begins to raise the question of *too much* information. With cancer we are already micromonitoring through blood tests, visual tests, etc. The crucial bar that we have to cross, for cancer, is whether those tests actually have an impact on saving lives or not. Ultimately the question is whether we end up detecting cancers that are clinically relevant, invasive, aggressive, likely to kill you — or will we be detecting thousands of cancers that aren't actually relevant and won't kill you and cause all sorts of economic consequences. This phenomenon is called "overdiagnosis," and it's a real concern among those who create cancer-detection tests. My opinion is that we will eventually find ways to discriminate one from the other. But there are people who are skeptics in the field who feel that we will be overrun with

useless information.

MOHR: It's all about feedback loops. If you're trying to control something and you want a specific outcome, you want to be measuring continuously, and measuring in a way that allows you to immediately tell the effects of each thing you do, because the thing you're trying to change is behavior. We can already do continuous glucose monitoring with a patch that just pierces the skin.

CHURCH: You might even have an inside/outside thing, where the skin is intact, but you've got something on the inside that's communicating.

MOHR: Well, in Sweden people are having RFID chips implanted in their skin so that they can pay, just with this thing in their skin. Like Apple Pay.

CHURCH: It's probably less invasive than tattooing.

JANNOT: What will it mean if we're going through our life getting constant feedback about our bodies now, our bodies in the future?

EGAN: I can only answer that as a fiction writer, because as a person, I don't live that way and I don't want to. Because I'm not a scientist, I'm interested in these things as they pertain to human inner life. And I come at it as someone who is uninterested in machines for their own sake. I think they're dull.

MOHR: For what it's worth, I don't think Jenny needs to be interested in her data for the monitoring of it to be useful to her at some point. We monitor our electricity use continuously. How often do you look at your electricity meter? You never

look at it. Unless you get an unusually high bill, or something flags it. Then you're glad it was being measured.

MUKHERJEE: I expect that those who are well won't look, but the ill will look. And the ill could be not just the physically ill; they could be the anxious, could be the mentally ill, could be those of us who have anxieties about our children, our futures, could be societies that are in peril.

JANNOT: What's this going to do to hypochondria?

MOHR: Yeah, that could be a problem. Imagine your body giving you "likes" from your measured parameters. Hypochondriacs would be like social-media addicts. Or maybe they'd just become extreme optimizers.

EGAN: There's a paranoid vision that comes right alongside it, which is: "There's a machine inside me doing something, and I have to get rid of it." It doesn't matter if a machine is there or not, that possibility is going to live in the minds of people who think that way.

BARZILAY: But would you get it implanted if you didn't want it?

EGAN: You might fear that someone else had implanted it in you. During the world wars, people all over the world worried that German spies were hidden around them. Imagine what it might be like to fear something that may be *inside* you. Think about how telecommunications technology has saturated our inner lives — our hyperemphasis on the visual, the curating and display of daily life, the constant monitoring of others. In the end, the technology seeps into our private experience. So when I think of someone installing a device inside his or her body to pay bills, I'm appalled. But as a fiction writer, I'm

ecstatic.

JANNOT: So, let's say that all this stuff works. We have a lot of monitoring, we have a lot of great data — what's the goal of it all?

MOHR: If I think about my goals for myself, it leads into why I have chosen this particular mission for my career — why everyone at this table has chosen to delve as deeply into the things they do — it's about improving the human condition, and also, not incidentally, making the science better for when we and our loved ones need it. It's why I build minimally invasive surgical tools. This is also why I keep up to date on my screening tests and think about better ways of monitoring the body: If at some point I get cancer, I want it to be Stage 1, and I'd like a surgical excision to be a cure in that situation, and I want a tiny incision. Using monitoring and technology to do small course corrections, rather than needing to do salvage when we are too far along in an illness.

CHURCH: When it comes to how we think about changing aging from our current normal, there are two major strategies here: One is extending longevity, and the other is aging reversal. The problem with longevity extension is, if you're not careful, you extend some of the weaker years of your life, which is not what we want. Aging reversal on the other hand sounds a little more speculative, but there are several examples demonstrated in mice where you can return old adult cells to embryonic stage by using a transcription factor to regulate certain genes. Another reason to do aging reversal rather than longevity is that it's hard to get funding for a long trial of a longevity drug, even for a veterinary drug, because if you say it's going to extend a dog's life by 10 years, that's a

10-year clinical trial. If you say that within five weeks it's going to make them stronger and more resistant to injury, then that's a five-week experiment.

MUKHERJEE: In terms of longevity, the diseases that are most likely to kill us are neurological diseases and heart disease and cancer. In some other countries, there is tuberculosis and malaria and other infectious diseases, but here it's the chronic diseases that dominate. There are three ways to think about these chronic diseases. One is the disease-specific way. So, you attack Alzheimer's as Alzheimer's; you attack cancer as cancer. The second one is that you forget about the disease-specific manners of attacking diseases and you attack longevity or aging reversal in general. You change diet, change genes, change whatever else — we might call them "trans factors," which would simply override the "cis factors" that existed for individual diseases. And the third option is some combination of that and some digital form of immortality, which is that you record yourself forever, that you clone yourself and somehow pass along that recording. Which is to say that the body is just a repository of memories, images, times. And as a repository, there's nothing special about it. The body per se, the mortal coil, is just a coil.

EGAN: I feel of two minds about longevity; on one hand, I want to live to be very, very old, partly because I had kids on the late side and I want to know their children as my mother — who had me at 24 — has known mine. But taking a step back, the mass possibility of extreme longevity has a selfish, devouring aspect. I mean, we're taxing the planet so hard as it is, the least we can do is not hang around forever!

JANNOT: And will we really want to? I mean, I realize this is a

fanciful question, but if this all works in, say, 25 years, will we be happier, will we have less sorrow in our society?

EGAN: I don't know, because we already confront so much less death than people did, say, before antibiotics. But does having fewer of those losses really make us happier?

CHURCH: After de-aging — or as part of it — we may set happiness itself as a goal. We have clearly set as goals simple measures like lowering cholesterol, but we're just beginning to study genetically engineering behavioral phenomena related to happiness.

MOHR: I'm not sure we really understand enough about sorrow and contentment to know. There was a book on people in extreme and terrible environments like concentration camps, and then also on people's just general malaise. The goals were looking at what were the characteristics of people who were psychologically resistant to tragedy. And what seemed to be most important were meaning, mastery and autonomy — feeling that there is some kind of meaning associated with things you do, working toward the acquisition of new skills and the ability to make choices for yourself. When you've got those three things, you are more resistant to tragedy. Maybe that is the secret to contentment.

MUKHERJEE: But if machines are doing all the work, then we'll have none of those things. We won't have mastery, we won't have meaning, we won't have autonomy.

MOHR: But we'll have art — art and mastery-oriented things like learning musical instruments.

CHURCH: But our future selves may not consider that rewarding — if our musical instrument is worse than the

machine's musical instrument, our chess worse than the machine's chess. If our mastery is lower, meaning is lower, because what does it mean to be able to be a poor imitation of a machine?

EGAN: Maybe a machine will be able to play the cello better than a human, but we go to the philharmonic to hear Yo-Yo Ma. Humans are more interesting than machines, plain and simple.

MOHR: Funny you mention cello, because that is the instrument I play. There are plenty of people, and even probably some machines, who can play the cello better than I do, but that doesn't take meaning away. I love the feeling of progression as I attain mastery — the beauty or the frustration in the moment. And it is my choice to keep trying — to keep creating. I think there is still great potential for humans to enjoy their lives in the time after menial work is done by machines.

BARZILAY: I actually believe that machines can help us achieve our goals better than we can do on our own. We are already using technology to expand our cognitive capacity — for instance, with machine translation we can read documents in foreign languages that we don't know. Why can't we expand this cognitive assistance to happiness? Happiness means different things to different people, but it is often linked to specific behaviors. Machines have immense capacity to remember our actions and predict our future behavior. This gives them the capacity to help us modify our behavior so we become our better selves. In my case, a simple heart-monitoring app changed the frequency and intensity of my running. The app gives points for achieving certain fitness

goals. When I first saw it, I just laughed and thought, Who can be motivated by these silly rewards? But guess what? Every morning at 5 a.m., I am running. Rain, M.I.T. deadlines, sleepiness — nothing stops me from getting my running points. And this change in my life has really made me happier.

MOHR: Exactly! You have clearly found purpose in getting better at running, and even though a car could drive you faster, that isn't the point at all. But both of our examples need bodies. Sid, in your vision of the uploaded consciousness, you're assuming that the body wears out but the mind can persist. I wonder if there isn't another ceiling beyond that in which the consciousness no longer wants to be conscious. Do you get immortality by uploading and then you feel this horrible sense of eternal ennui because you were uploaded and can no longer decide to learn to play the cello or go running along the Charles River?

MUKHERJEE: You're stuck being conscious.

EGAN: I think we're forgetting a basic truth about human life: Transience is what makes it precious. The inevitability of death infuses our lives with meaning and urgency. Hard to imagine sustaining those qualities in an eternally uploaded consciousness. You're left with just sensation. I'm not sure that's a gain in the end.

CHURCH: Well, if you have simple aging reversal, so you actually feel like, I changed from being 64 to being 24 — I can do everything I could do when I was 24 plus I have the experience of being older, and the open-ended explorations ahead of reading and writing our universe — I doubt that I'm going to have a serious case of ennui.

MOHR: You could even take up the cello.