MIT Technology Review

The precision medicine issue

Vol. 121 No. 6 | Nov/Dec 2018 | $9.99 USD $10.99 CAD

$2 million would save her life. Could you pay?

Should you?

Medicine is becoming hyper-personalized, hyper-accurate ... and hyper-unequal. p. 38

The AIs taking over from doctors

p. 22

Curing cancer with customized vaccines

p. 46

How to plan your digital afterlife

p. 76

12

0 74808 01308 1

Artificial intelligence is here.

Own what happens next.

MIT Technology Review

EmTech Digital

March 25–26, 2019
St. Regis Hotel
San Francisco, CA

Meet the people leading the next wave of intelligent technologies.

Andrew Anagnost
President and CEO, Autodesk

Kathy Gong
Cofounder and CEO, WafaGames

Daphne Koller
Founder and CEO, Insitro

Sergey Levine
Assistant Professor, EECS, UC Berkeley

Harry Shum
Executive Vice President, Microsoft Artificial Intelligence and Research Group

EmTechDigital.com/2019

Gideon Lichfield is editor in chief of MIT Technology Review.

It's been nearly two decades since the first human genome was sequenced. That achievement opened up the promise of drugs and treatments perfectly tailored to each person's DNA. But personalized or "precision" medicine has always felt a bit like flying cars, sexbots, or lab-grown meat—one of those things we're perpetually being promised but never quite getting. This issue of *MIT Technology Review* makes the case that, in fact, the age of precision medicine has been slowly dawning on us all this time—and we're unprepared.

What's changing fastest now is the sheer volume of medical data available, and the tools for analyzing it. As Antonio Regalado points out in his opening essay (page 8), the number of people getting their DNA tested is now in the tens of millions and doubling each year.

By pairing DNA data with people's medical records, algorithms can predict your risk of certain common diseases and suggest drugs and diets to ward them off, as Ali Torkamani and Erik Topol explain (page 20). Cancer drugs are now being customized to individual patients, as Adam Piore reports (page 46). Epigenetic data can forecast how long you'll live, writes Karen Weintraub (page 80), and new "senolytic" drugs might keep age-related ailments at bay for more of that time, reports Stephen S. Hall (page 84).

Not just DNA sequences but data of all kinds is being scooped up and crunched in vastly greater quantities than before. As Rachel Metz explains (page 56), it's becoming possible to track mental illness just by monitoring how you tap, type, and swipe on your phone.

Better treatments and healthier living aren't the only benefits. Doctors like Rahul Parikh (page 28) hope to be able to spend more time getting to know their patients as algorithms take on the more routine tasks. In a UK trial, AI systems are already replacing physicians for simple consultations, as Douglas Heaven reports (page 22). That could help meet the ballooning health-care needs of an aging population.

The problem? As medicine gets more personalized, it risks getting more unequal. Our cover story is Regalado's gripping and troubling account (page 38) of the parents who raise millions of dollars to finance gene-therapy cures for their children's ultra-rare diseases. Are they trailblazers for a technology that will one day provide cheap, customized care to everyone, or harbingers of a future in which only the super-wealthy and crowdfunding whizzes are saved? IVF combined with genetic screening can weed fatal diseases out of a family for good, but, Laura Hercher argues (page 68), it could also lead to two genetically distinct human castes—one rich and disease-free, the other poor and disease-ridden. The rich and well educated won't only be better able to afford boutique treatments; they'll be more likely to have the technology, and hence the data, that helps them avoid falling ill in the first place.

All this has more than merely medical consequences. Nathaniel Comfort warns (page 16) that our growing ability to find genetic correlations with things like intelligence is threatening to revive the ugly dogma of eugenics. And what, asks Mary Madden (page 34), can any of us do to keep tabs on how the oceans of data about us are being used, or misused?

We'll face this question even in death. As Courtney Humphries reports (page 72), people now in their 30s will have generated enough data by the time they die to power quite convincing digital avatars of themselves. So who will own you when you're gone? At least Simson Garfinkel (page 76) has some advice on how to prepare.

This issue of the magazine, therefore, spans the entire arc of human existence, from before you're born until after you die. Through it all runs a simple question. We know that in health care, human beings are unequal. But just how unequal are we willing to be?

NEPHI NIVEN

WORLD-CLASS EXECUTIVE EDUCATION AT BROWN UNIVERSITY

Transforming Mid-Career Professionals

An Executive Master's degree from Brown University will prepare you to lead your organization, transform your field, and build a powerful lifelong professional network.

You will join a vibrant learning community and apply your new knowledge and skills through a critical challenge project.

The results can be seen in our alumni who are proven leaders; impacting their organizations and the world.

BROWN
School of Professional Studies

brown.edu/professional

IE Brown Executive MBA

Executive Master in Cybersecurity

Executive Master in Science and Technology Leadership

Executive Master of Healthcare Leadership

Editorial

Editor in chief
Gideon Lichfield

Executive editor
Megan McCarthy

Editor at large
David Rotman

Deputy editor
Michael Reilly

News editor
Niall Firth

Managing editor
Timothy Maher

San Francisco bureau chief
Martin Giles

Senior editor, MIT News
Alice Dragoon

Senior editor, AI and robotics
Will Knight

Senior editor, cyborgs
Rachel Metz

Senior editor, biomedicine
Antonio Regalado

Senior editor, energy
James Temple

Senior editor, infrastructure
Elizabeth Woyke

Reporters
Karen Hao (AI), Charlotte Jee (news), Mike Orcutt (blockchain), Erin Winick (future of work)

Copy chief
Linda Lowenthal

Senior production director
James LaBelle

Design

Chief creative officer
Eric Mongeon

Lead designer
Emily Luong

Marketing and events designer
Kyle Thomas Hemingway

Art assistant
Emily Caulfield

Product development

Director of product
Vanessa DeCollibus

Project manager
Allison Chase

User interface/
user experience designer
Jon Akland

Engineers
Shaun Calhoun, Molly Frey, Jason Lewicki

Corporate

Chief executive officer and publisher
Elizabeth Bramson-Boudreau

Director, human resources
Hilary Siegel

Assistant to the CEO
Katie McLean

Manager of information technology
Colby Wheeler

Office manager
Linda Cardinal

Licensing and communities

Vice president, licensing
and communities
Antoinette Matthews

Client services manager
Ted Hu

Events

Senior vice president,
events and strategic partnerships
Amy Lammers

Director of events programming
Laura Janes Wilson

Senior events manager
Nicole Silva

Content and program developer, events
Kelsie Pallanck

Events associate
Bo Richardson

Finance

Finance director
Enejda Xheblati

General ledger manager
Olivia Male

Accountant
Letitia Trecartin

Administrative assistant
Andrea Siegel

Consumer marketing

Senior vice president,
marketing and consumer revenue
Doreen Adger

Director of analytics
Tom Russell

Director of audience development
Rosemary Kelly

Director of digital marketing
and communications
Josh Getman

Assistant consumer marketing manager
Sanjeet Chowdhury

Assistant product marketing manager
Amanda Saeli

Advertising sales

Vice president, sales and brand partnerships
Marii Sebahar
marii@technologyreview.com
415-416-9140

Senior director, brand partnerships
Kristin Ingram
kristin.ingram@technologyreview.com
415-509-1910

Director, brand partnerships
Debbie Hanley
debbie.hanley@technologyreview.com
214-282-2727

Director, brand partnerships
Ian Keller
ian.keller@technologyreview.com
203-858-3396

Digital sales strategy manager
Ken Collina
ken.collina@technologyreview.com
617-475-8004

Advertising services
webcreative@technologyreview.com
617-475-8004

Media kit
www.technologyreview.com/media

MIT Technology Review Insights

Vice president of international business development, head of MIT Technology Review Insights
Nicola Crepaldi

Senior editor
Mindy Blodgett

Senior project manager
Anna Raborn

Content manager
Jason Sparapani

Director of consulting, Asia
Claire Beatty

MIT Enterprise Forum, Inc.

Chairman and president
Elizabeth Bramson-Boudreau

Executive director and clerk
Antoinette Matthews

Treasurer
Enejda Xheblati

Director of chapter leadership and process
Gaylee Duncan

Board of directors

Martin A. Schmidt
Whitney Espich
Jerome I. Friedman
Joichi Ito
Israel Ruiz
David Schmittlein
Alan Spoon

Customer service and subscription inquiries

National
800-877-5230

International
903-636-1115

E-mail
customer_service@
mittechnologyreview.info

Web
www.technologyreview.com/
customerservice

MIT Records (alums only)
617-253-8270

Reprints
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions
licensing@technologyreview.com

MIT Technology Review

One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to bring about better-informed and more conscious decisions about technology through authoritative, influential, and trustworthy journalism.

MIT Technology Review
Insights

Be a part of MIT Technology Review. Join our Global Panel.

You have opinions on today's tech trends. So share them.

The MIT Technology Review Global Panel is a growing group of executives and innovators. As a panelist, you'll be asked to take part in surveys, weighing in on a range of technology topics. And you'll be among the first to see the results.

Apply at
technologyreview.com/globalpanel

Skeptics say drugs based on genetic insights have underdelivered.

But look carefully and they're everywhere.

by Antonio Regalado

Look how far precision medicine has come

Sometime this fall, the number of people who have spit in a tube and sent their DNA to the largest consumer DNA testing companies, like Ancestry and 23andMe, is likely to top 20 million. The list by now is certain to include some of your classmates and neighbors. If you are just tuning in, this figure will seem huge. And you might wonder: how did we get here?

The answer is little by little. The number of people getting DNA reports has been doubling, roughly, every year since 2010. The figures are now growing by a million each month, and the DNA repositories are so big that they're enabling surprising new applications. Consumers are receiving scientific predictions about whether they'll go bald or get cancer. Investigators this year started using consumer DNA data to capture criminals. Vast gene hunts are under way into the causes of insomnia and intelligence. And 23andMe made a $300 million deal this summer with drug company GlaxoSmithKline to develop personalized drugs, starting with treatments for Parkinson's disease. The notion is that targeted medicines could help the small subset of Parkinson's patients with a particular gene error, which 23andMe can easily find in its database.

Ever since the Human Genome Project—the 13-year, $3 billion effort to decipher the human genetic code—researchers and doctors have been predicting the arrival of "precision medicine."

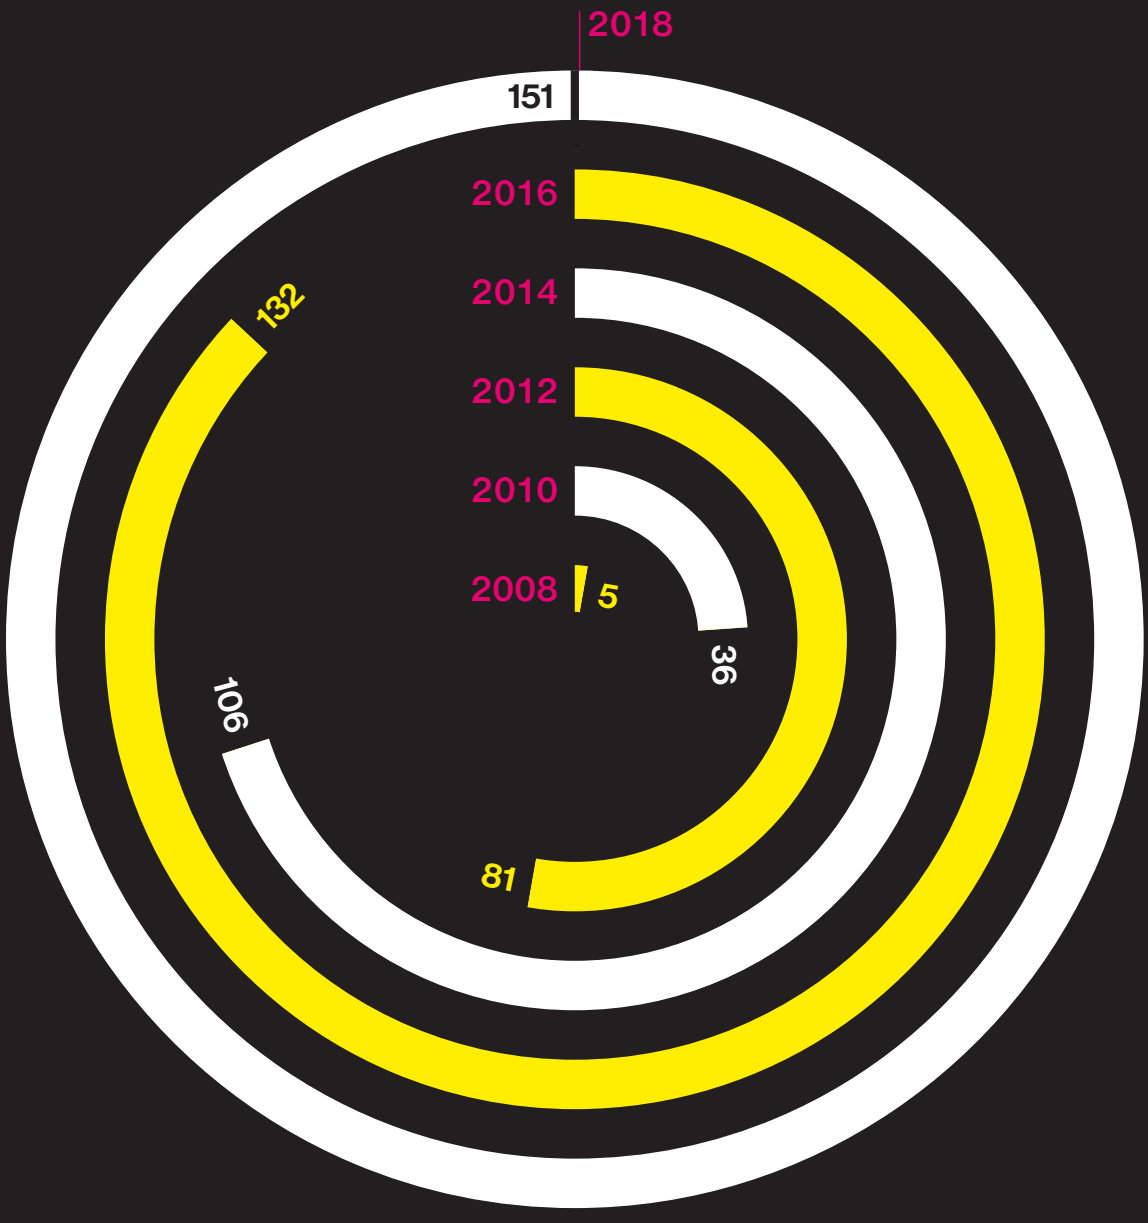

Genetic Rx

The number of "personalized" drugs on the US market has multiplied during the past decade

Personalized Medicine Coalition

The growing tally of personalized or targeted medicines consists of those drugs whose label includes information about how genetic makeup can affect a person's response to a drug.

A time line of precision medicine

Your genes and mine differ. Can science create drugs to match?

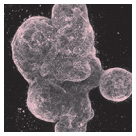

1998

TARGETED DRUGS

The breast cancer treatment Herceptin is approved for sale in the US. It is the first cancer drug to target an underlying genetic defect responsible for producing tumors.


THE WALL STREET JOURNAL.
MARKETPLACE
New Era of Personalized Me
Targeting Drugs For Each Unique Genetic Profile

1999

PERSONALIZED MEDICINE

The Wall Street Journal declares a "New Era of Personalized Medicine" based on genetic mapping of one-letter DNA differences between humans. Drug makers call it the start of a "grand experiment."

2000

HUMAN GENOME PROJECT

The Human Genome Project and Celera Genomics, a company started by entrepreneur J. Craig Venter, both announce that a working draft of the genome sequence is complete.

CONTINUED

It's a term with no agreed-upon definition, although it suggests most strongly just the kinds of medicines that Glaxo and 23andMe are pursuing: more targeted and more effective because they take into account a person's particular genetic makeup. President Bill Clinton, at the unveiling of the genome's first draft back in June 2000, said the data would "revolutionize the diagnosis, prevention, and treatment of most, if not all, human diseases."

Almost two decades after those big promises, it is in vogue to question why precision medicine has not delivered more. A report in the *New York Times* this summer, noting that deaths from cancer still outnumber cures by a wide margin, asked: "Are We Being Misled About Precision Medicine?" One reason for this seemingly slow progress is that not all precision medicine involves drugs. As gene hunts gain in scope—the latest involve comparisons of more than a million people's DNA and health records—an inconvenient fact about many common diseases has emerged: they don't, by and large, have singular causes. Instead, many hundreds of genes play small roles, and there is no obvious point at which to intervene with a pill.

So instead of drugs, we are seeing a new predictive science in which genetic risk profiles may say which people should lower their blood pressure, which should steel themselves for Alzheimer's, and which cancer patients aren't going to benefit from chemotherapy and can skip the ordeal. To be sure, these sorts of prognostics aren't widely accepted, and it's hard to get people to change their behavior. Yet for many people, these predictions may begin to offer a concrete route to precision health and increased knowledge of their own biology.

Look beyond cancer, and some definitive cures *have* arrived. As with those growing millions sending in their DNA, it's easy to miss the change before it's everywhere. Here are just two medications of note: a drug that mops up hepatitis C in 90% of those who take it and an experimental gene therapy that is curing a rare, fatal, and previously untreatable childhood disease, spinal muscular atrophy. Though these treatments come from different corners of biology, it's what they have in common that's important: each benefits from detailed understanding of genetic information *and* tools to control it.

To our thinking, these drugs display real precision. The hep C pill, called Solvadi, consists of a chemical that is irresistible to the replicating virus, but when the drug comes in contact with the virus's genome, replication quickly grinds to a halt. The treatment for spinal muscular atrophy, meanwhile, is a genetic replacement part. With gene therapy, doctors can add fresh DNA instructions to the child's nerve cells. The dozen or so kids who've gotten the therapy at a young age don't develop the disease.

All this traces back to even before the Human Genome Project. Think instead of the foundational act of the biotechnology industry, 40 years ago. On September 6, 1978, Genentech announced "the successful laboratory production of human insulin." Before then, diabetics had injected insulin from pigs. It took around two tons of pig parts to extract eight ounces (227 grams) of pure insulin. But Genentech had found a way to splice the human version of the insulin-producing gene into *E. coli* bacteria, which then manufactured the hormone. Genentech still keeps the 40-year-old press release online.

To the pharmaceutical houses of the 20th century, with their roots in commercial dye making and synthetic chemistry, these new biotech drugs looked at first like a sideshow. They were hard to make and inconvenient to take (by injection, mostly). The pharma giants could easily believe their way of doing things would always dominate. Until well into the 1990s, a single drug company, Merck, was more valuable than all biotech companies combined. It probably seemed as if biotech would never arrive—until it did. Of the 10 best-selling drugs in the US during 2017, seven (including the top seller, the arthritis drug Humira) are biotech drugs based on antibodies. Antibodies embody biological precision too. These tiny blood proteins, normally part of our immune

COURTESY OF NCI CENTER FOR CANCER RESEARCH (TUMOR); COURTESY OF CELERA GENOMICS CORPORATION (VENTER)

Seeking better drugs

The proportion of patients who actually benefit from a best-selling drug in each category

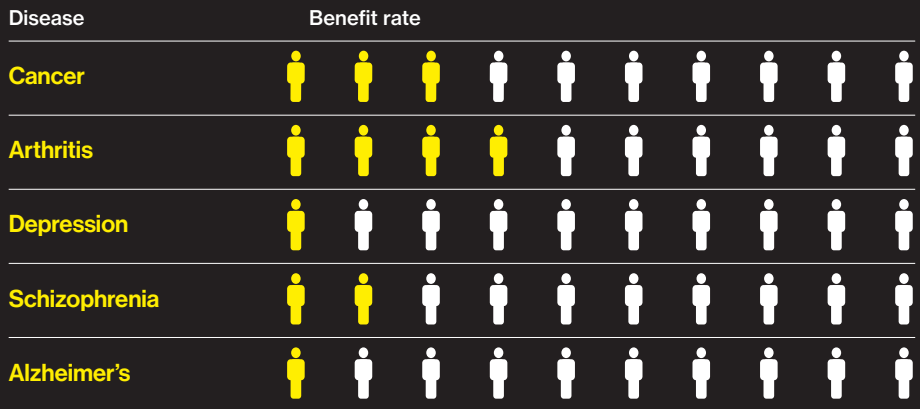

Schork, NJ, Nature; PubMed

Genetic information explodes

Cost of sequencing a genome

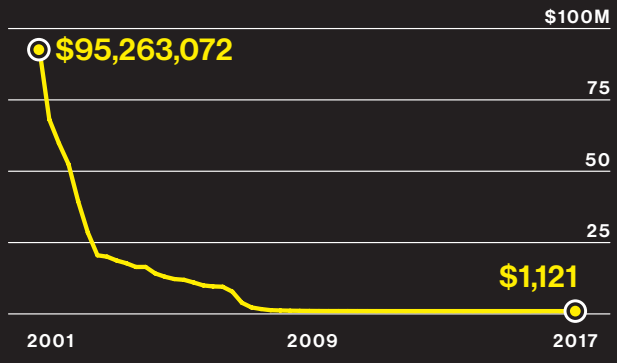

NHGRI

Number of people who have bought consumer DNA tests

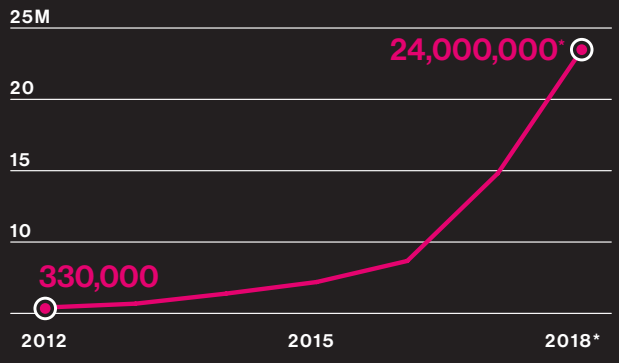

Company reports, Leah Larkin, ISOGG

*2018 DATA ESTIMATED

Drugs based on DNA

Percentage of drugs in development that may be tailored to a person's genetic profile

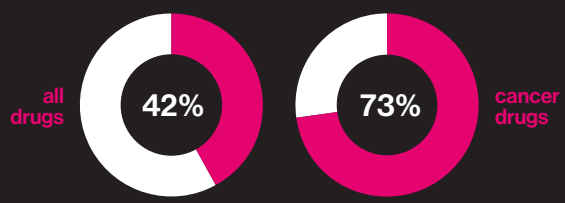

Tufts; Personalized Medicine Coalition

Number of the 10 best-selling drugs in the US that are biological molecules

MIT Technology Review

2013

A STEP BACK

The US bars consumer DNA testing company 23andMe from calculating people's chances of common diseases and cancer, calling results inaccurate. It's a major blow for consumer genetics.

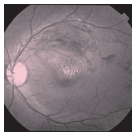

2016

AI DOCTOR

A team at Google uses an artificial-intelligence method called deep learning to diagnose symptoms of blindness by reading retina scans. It performs as well as an ophthalmologist.

2017

GENE THERAPY

The US approves three gene therapies in a four-month span. Two of them prime immune cells to kill blood cancers. The other is a one-time treatment for an inherited form of blindness. It costs $850,000.

2018

RISK PREDICTION

Giant gene studies lead to "risk score" technology based on measuring millions of sites in a person's genome. Doctors say it can predict heart disease and other conditions.

response, fit—like a key in a lock—onto other molecules, like those dotting the surface of a cancer cell. And just like insulin, they're often constructed using DNA code retrieved from our bodies.

Insulin and antibodies are meant to work the same way on everyone. But no two people's genomes are exactly the same—about 1% of the DNA letters differ between any two of us. Those differences can explain why one person is ill and another isn't, or why one person's version of diabetes is different from another's. Drugs that take into account these differences in genetic information are called "targeted" drugs.

The cancer drug Herceptin, an antibody that reached the market in 1998, was among the first. It was effective, but mostly in people whose newly diagnosed breast cancer was growing because of specific genetic damage—about 20% of cases. It depended on the genome of the tumor itself. Herceptin came to market with the admonition that, to get it, you should first have a test to see if you would benefit. According to the US National Cancer Institute, there are now more than 80 such targeted medicines for cancer on the market.

Critics argue rightly enough that such medications still do too little for too few people at too great a cost (often $10,000 a month). In fact, on the whole, those who survive cancer still owe little to targeted drugs. "The single biggest determinant of who survives cancer is who has insurance," Greg Simon, who leads the Biden Cancer Initiative, has said—not whether there's a drug to match their mutation. Some think we are spending too much time searching under the lamplight shed by genetic tools. "Perhaps we had been seduced by the technology of gene sequencing—by the sheer wizardry of being able to look at a cancer's genetic core," a Pulitzer Prize-winning cancer doctor, Siddhartha Mukherjee, wrote this summer.

He's right that the impulse toward precision medicine, cost be damned, springs from new technology. It's what it *can do*. And so you can be sure even more personalization is on the horizon. Genentech (which created Herceptin) now imagines what it calls "cancer vaccines," tailored not just to broad subtypes of people but to the unique signature of a person's tumor. The new approach involves collecting information about the peculiarities of a person's cancer through high-speed genome sequencing; using software to analyze and predict what a custom biological drug would look like (they will be reverse images of antibodies, known as antigens, that stimulate the immune system); and then quickly manufacturing it. No two of these vaccines would be alike. Also, note this: if and when the US Food and Drug Administration approves these vaccines, it won't be greenlighting a particular compound. Instead, it will approve a computerized *process* for turning DNA information into drugs.

Medicine as programmatic and predictable as a computer? The idea has begun to exert a potent appeal in Silicon Valley, where some of tech's biggest names now see biology as "just a code" they can crack. Marc Andreessen (best known for inventing the web browser) is one of them. The venture fund he cofounded, Andreessen Horowitz or a16z, has set aside a total of $650 million since 2015 to put into biotech investments. As the firm's blog states with awe, "You don't just read the code of biology but you can also write, or design, with it."

Welcome to biotech, a16z. Yet they're on to something. Even 40 years after Genentech's insulin press release, genetic engineering is a marvel worth rediscovering. The ability to see, understand, and manipulate human genes and the proteins they make is the great advance that is still unfolding in all its immense complexity four decades later. Biology isn't anywhere as neat as a computer program, but little by little, we're learning how to control it. To enzymes and antibodies we've added gene therapy and gene editing. We haven't sequenced one genome—we've sequenced a million. An astute observer might realize we've already come a long way.

Antonio Regalado is MIT Technology Review's senior editor covering biomedicine.

STUART ISETT/FLICKR (WOJCICKI); COURTESY OF SPARK THERAPEUTICS; BMC OPHTHALMOLOGY (RETINA); COURTESY OF UK BIOBANK

Big questions need big data

Studies are using DNA data from more people than ever

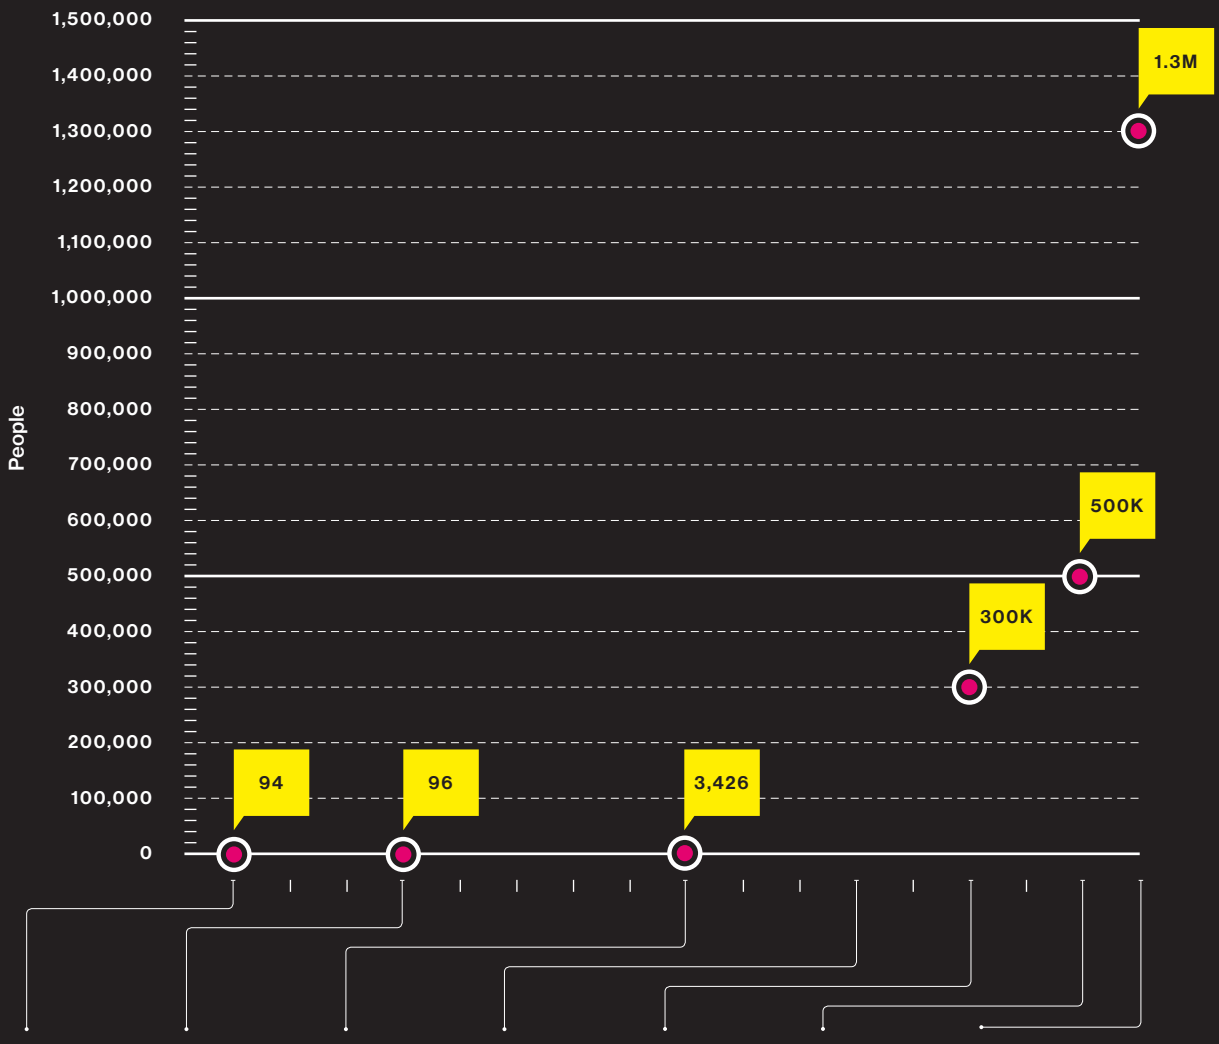

2002
Japanese scientists use a new approach—the genome-wide association study—to hunt for the causes of heart attack.

2005
A gene hunt reveals critical mutations that increase the risk of macular degeneration, a common cause of blindness.

2010
Consumer test company 23andMe contributes user data to a search for Parkinson's genes.

2013
The FDA cracks down on consumer test companies offering genetic health predictions from DNA, calling the results unreliable.

2015
Why are some people fatter than others? Clues from a genetic study are quickly offered to consumers in the form of "DNA diet" tests.

2017
A massive trove of gene data from the UK Biobank permits simultaneous analysis of 2,000 human traits and diseases.

2018
Researchers identify genes linked to educational success. They warn against using the results as a "DNA IQ test."

A search for the genes behind insomnia is the largest genetic study ever. It relies heavily on the consumer DNA database of 23andMe.

Shape the future of MIT Technology Review

Join our Beta Panel

Did you know we're redesigning the MIT Technology Review website?

Be a part of the process by joining our beta tester panel. As one of the first people to test-drive the new website design, your feedback will be a vital part of making our new site a better user experience for all.

Apply at
technologyreview.com/beta-test

Let's talk about health

It's all about one thing: data. Your data can be used to predict disease, improve your lifestyle, and make your doctor smarter and less overworked. But data can be misused if it's in the wrong hands. It can be flawed if it doesn't come from a diverse range of people. And it can be twisted to bolster theories of racial superiority.

1

OPENING A DOOR TO EUGENICS

New ways of using your genetic data could bolster scientific racism and encourage discrimination.

By **NATHANIEL COMFORT**

Want to predict aggression? Neuroticism? Risk aversion? Authoritarianism? Academic achievement? This is the latest promise from the burgeoning field of sociogenomics.

There have been many "DNA revolutions" since the discovery of the double helix, and now we're in the midst of another. A marriage of the social and natural sciences, it aims to use the big data of genome science—data that's increasingly abundant thanks to genetic testing companies like 23andMe—to describe

the genetic underpinnings of the sorts of complex behaviors that interest sociologists, economists, political scientists, and psychologists. The field is led by a group of mostly young, often charismatic scientists who are willing to write popular books and op-eds, and to give interviews and high-profile lectures. This work shows that the nature-nurture debate never dies—it is just cloned and raised afresh in a new world.

Advocates of sociogenomics envision a prospect that not everyone will find entirely benevolent: health "report cards," based on your genome and handed out at birth, that predict your risk of various diseases and propensity for different behaviors. In the new social sciences, sociologists will examine the genetic component of educational attainment and wealth, while economists will envision genetic "risk scores" for spending, saving, and investment behavior.

Without strong regulation, these scores could be used in school and job applications and in calculating health insurance premiums. Your genome is the ultimate preexisting condition.

Such a world could be exciting or scary (or both). But sociogenomicists generally focus on the sunny side. And anyway, they say with a shrug, there's nothing we can do about it. "The genie is out of the bottle," writes the educational psychologist Robert Plomin, "and cannot be stuffed back in again."

Is this what the science says, in fact? And if it is, is it a valid basis for social policy? Answering these questions demands setting this new form of hereditarian social science in context—considering not merely the science itself but the social and historical perspective. Doing so can help us understand what's at stake and what the real risks and benefits are likely to be.

WEIRD SCIENCE

If this is "the science," the science is weird. We're used to thinking of science as incrementally seeking causal explanations for natural phenomena by testing a series of hypotheses. Just as important, good science tries as hard as it can to disprove the working hypotheses.

Sociogenomics has no experiments, no null hypotheses to accept or reject, no deductions from the data to general principles. Nor is it a historical science, like geology or evolutionary biology, that draws on a long-running record for evidence.

Sociogenomics is inductive rather than deductive. Data is collected first, without a prior hypothesis, from longitudinal studies like the Framingham Heart Study, twin studies, and other sources of information—such as direct-to-consumer DNA companies like 23andMe that collect biographical and biometric as well as genetic data on all their clients.

Algorithms then chew up the data and spit out correlations between the trait of interest and tiny variations in the DNA, called SNPs (for single-nucleotide polymorphisms). Finally, sociogenomicists do the thing most scientists do at the outset: they draw inferences and make predictions, primarily about an individual's future behavior.

Sociogenomics is not concerned with causation in the sense that most of us think of it, but with correlation. The DNA data often comes in the form of genome-wide association studies (GWASs), a means of comparing genomes and linking variations of SNPs. Sociogenomics algorithms ask: are there patterns of SNPs that correlate with a trait, be it high intelligence or homosexuality or a love of gambling?

Yes—almost always. The number of possible combinations of SNPs is so large that finding associations with any given trait is practically inevitable.

The evolutionary biologist Graham Coop shows that big data can lull us into a false sense of objectivity. The success of GWASs, he writes, "seems to suggest that we'll soon be able to settle debates about whether behavioral differences among populations are driven in part by genetics." However, he adds, "answering this question is a lot more complicated than it seems."

Coop offers what he calls a "toy" example of a misleading polygenic study—a thought experiment. The hypothetical research question: Why do the English drink more tea than the French?

Coop's imaginary researcher, Bob, uses data from existing databases like the UK Biobank. He counts up the average number of alleles (different forms of a gene) associated with a preference for tea in English people and French people. "If the British, overall," Coop writes, "are more likely to have alleles that increase tea consumption than French people, then Bob might say that we have demonstrated that the difference between French and UK people's preference for tea is in part genetic."

Being a conscientious scientist, of course, Bob would offer the usual assurances about the quality of his data. He would piously insist that his results do not show that all Brits who drink lots of tea do so because of their genes—only that the overall difference between the populations is partly genetic.

Coop then walks us through the problems with this thinking. It ignores the crucial fact that alleles may behave differently in different genomes and in

different environments: "The issue is that GWAS studies do not point to specific alleles for tea preferences, only to alleles that happen to be associated with tea preference in the current set of environments experienced by people in the UK Biobank." In other words, we can't be sure that a different group of people with the same genetic variations would be equally avid tea drinkers. And even if they were, we still wouldn't know it was those genes that made them love tea.

Bob, then, commits two fallacies. First, he confuses correlation and causation. The study does not show that the putative tea-drinking alleles affect tea drinking—merely that they are associated with it. They are predictive but not explanatory. The second fallacy is one I learned on the first day of class in college biostatistics: statistical significance does not equal biological significance. The number of people buying ice cream at the beach is correlated with the number of people who drown or get eaten by sharks at the beach. Sales figures from beachside ice cream stands could indeed be highly predictive of shark attacks. But only a fool would bat that waffle cone from your hand and claim that he had saved you from a Great White.

"Complex traits are just that—complex," Coop concludes. "Most traits are incredibly polygenic, likely involving tens of thousands of loci [i.e., SNPs or genes]. These loci will act via a vast number of pathways, mediated by interactions with many environmental and cultural factors."

A LONG TRADITION

Sociogenomics is the latest chapter in a tradition of hereditarian social science dating back more than 150 years. Each iteration has used new advances in science and unique cultural moments to press for a specific social agenda. It has rarely gone well.

The originator of the statistical approach that sociogenomicists use was Francis Galton, a cousin of Charles Darwin. Galton developed the concept and method of linear regression—fitting the best line through a curve—in a study of human height. Like all the traits he studied, height varies continuously, following a bell-curve distribution. Galton soon turned his attention to personality traits, such as "genius," "talent," and "character." As he did so, he became increasingly hereditarian. It was Galton who gave us the idea of nature versus nurture. In his mind, despite the "sterling value of nurture," nature was "by far the more important."

Galton and his acolytes went on to invent modern biostatistics—all with human improvement in mind. Karl Pearson, Galton's main protégé (who invented the correlation coefficient, a workhorse statistic of GWASs and hence of sociogenomics), was a socialist who believed in separating sex from love. The latter should be spread around liberally, the former tightly regulated to control who bred with whom—that is, for eugenic ends.

The point is that eugenics was not, as some claim, merely an unfortunate bit of specious science. It was central to the development of biological statistics. This entanglement runs down the history of hereditarian social science, and today's sociogenomicists, like it or not, are heir to it.

Early in the 20th century, a vicious new strain of eugenics emerged in America, based on the new science of Mendelian genetics. In the context of Progressive-era reformist zeal, belief in a strong government, and faith in science to solve social problems, eugenics became the basis of coercive social policy and even law. After prominent eugenicists canvassed, lobbied, and testified on their behalf, laws were passed in dozens of states banning "miscegenation" or other "dysgenic" marriage, calling for sexual sterilization of the unfit, and throttling the stream of immigrants from what certain politicians today might refer to as "shithole countries."

At the end of the 1960s, the educational psychologist Arthur Jensen published an enormous article in the *Harvard Educational Review* arguing that Negro children (the term of the day) were innately less intelligent than white children. His policy action item: separate and unequal school tracks, so that African-American children would not become frustrated by being over-challenged with abstract reasoning. What became known as "Jensenism" has resurfaced every few years, in books such as Charles Murray and Richard Herrnstein's *The Bell Curve* (1994) and the journalist Nicholas Wade's *A Troublesome Inheritance* (2014).

Given the social and political climate of 2018, today would seem a particularly inauspicious time to undertake a new and potentially vastly more powerful expression of genetic determinism. True, the research papers, white papers, interviews, books, and news articles I've read on the various branches of sociogenomics suggest that most researchers want to move past the racism and social stratification promoted by earlier hereditarian social scientists. They downplay their results, insist upon avoiding bald genetic determinism, and remain inclusive in their language. But, as in the past, fringe groups have latched onto sociogenomic research as evidence for their hostile claims of white superiority and nationalism.

Given the social and political climate of 2018, today would seem a particularly inauspicious time to undertake a new and potentially vastly more powerful expression of genetic determinism.

SOCIAL RISKS

Social genomics comes with its own large set of social risks—and number one on the list is failing to grapple sufficiently with those risks. In the 2012 paper that has become the de facto manifesto of genoeconomics (the use of genetic data to predict economic behavior), Daniel Benjamin and his coauthors dedicated two full sections to "pitfalls." Every one of them is methodological and statistical—false positives, studies with too few participants, and so forth. Most could be fixed with more data and better statistics.

Some in the field readily acknowledge the skeletons in the closet. "Eugenics is not safely in the past," wrote Kathryn Paige Harden, a developmental behavior geneticist at the University of Texas, in a *New York Times* op-ed earlier this year. Harden lamented the rise of the so-called human biodiversity movement (referring to it as "the eugenics of the alt-right"), with its ties to white supremacy and its specious claims to scientific legitimacy. Members of this movement, she wrote, "enthusiastically tweet and blog about discoveries in molecular genetics that they mistakenly believe support the ideas that inequality is genetically determined; that policies like a more generous welfare state are thus impotent; and that genetics confirms a racialized hierarchy of human worth."

Sociogenomics is the latest chapter in a history of hereditarian social science dating back more than 150 years.

Indeed, the human biodiversity crowd and other so-called "race realists" love sociogenomics. *American Renaissance*, a publication run by the avowed white supremacist Jared Taylor, features articles about the possibilities of sociogenomics, as does the HBD Bibliography, an aggregator of hereditarian materials. Steve Sailer, a well-known and prolific writer in white supremacist and human biodiversity circles, writes extensively about sociogenomics on "race realist" sites such as Unz Review and VDARE.

To be clear: I am not saying that sociogenomicists are racists. I am saying that their work has serious social implications outside the lab, and that too few in the field are taking those problems seriously.

Genetics has an abysmal record for solving social problems. In 1905, the French psychologist Simon Binet invented a quantitative measure of intelligence—the IQ test—to identify children who needed extra help in certain areas. Within 20 years, Binet was horrified to discover that people were being sterilized for scoring too low, out of a misguided fear that people of subnormal intelligence were sowing feeblemindedness genes like so much seed corn.

What steps can we take to prevent sociogenomics from suffering the same fate? How do we ensure that polygenic scores for educational attainment are used to offer extra help tailored to those who need it—and ensure that they don't become tools of stratification?

Here's one way: when the evolutionary biologist Coop and his student Jeremy Berg published a GWAS paper on the genetics of human height, they took the extraordinary step of writing a 1,500-word blog post about what could and could not be legitimately inferred from their paper.

Why isn't this more common? The field needs more people like Coop—and fewer cheerleaders. It needs scientists who reckon with the social implications of their work, especially its potential for harm—scientists who take seriously the social critique of science, who understand their work in both its scientific and historical contexts. It is such people who stand the best chance of using this potent knowledge productively. For scientists studying human social genomics, doing so is a moral responsibility.

Nathaniel Comfort is a professor of the history of medicine at Johns Hopkins. Jon Phillips contributed research for this article.

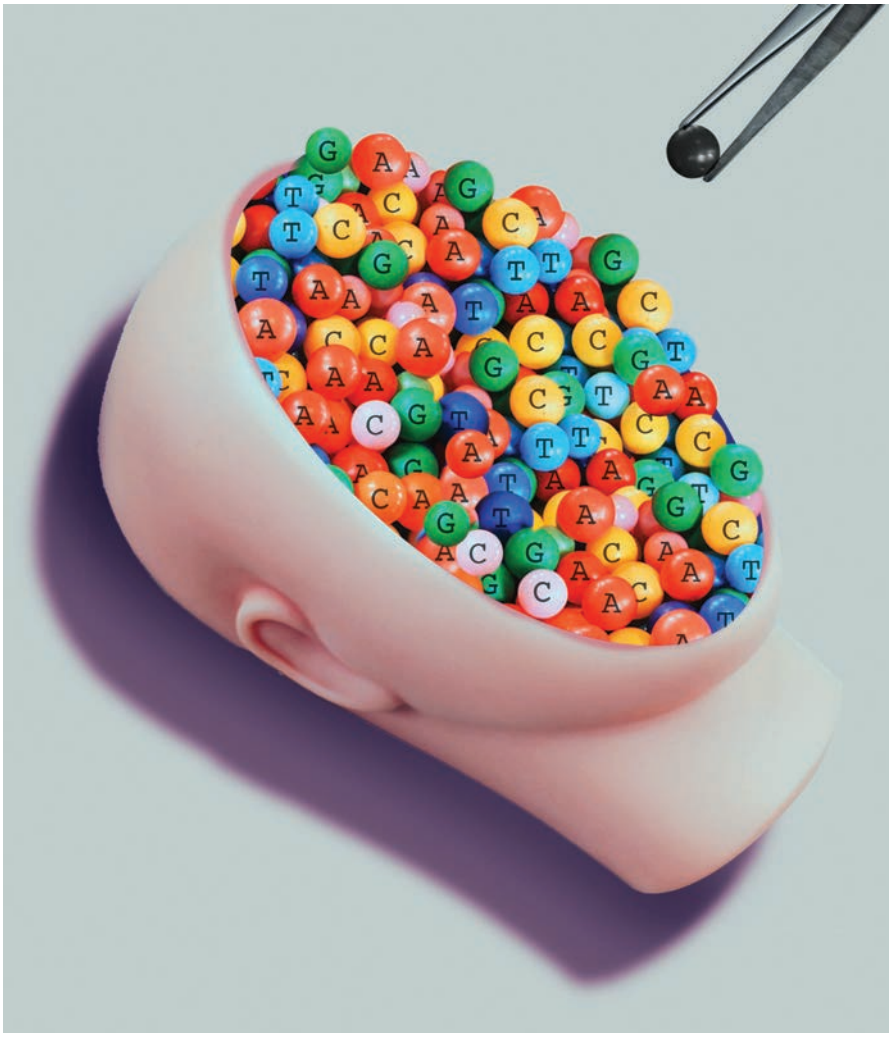

Your genome, on demand

How your detailed genetic profile can predict your risk of diseases and improve your health.

By **Ali Torkamani and Eric Topol** / Illustration by **Nico Ortega**

In early 2018, it was estimated that over 12 million people had had their DNA analyzed by a direct-to-consumer genetic test. A few months later, that number had grown to 17 million. Meanwhile, geneticists and data scientists have been improving our ability to convert genetic data into useful insights—forecasting which people are at triple the average risk for heart attack, or identifying women who are at high risk for breast cancer even if they don't have a family history or a *BRCA* gene mutation. Parallel advances have dramatically changed the way we search for and make sense of volumes of data, while smartphones continue their unrelenting march toward becoming the de facto portal through which we access data and make informed decisions.

Taken together, these things will transform the way we acquire and use personal genetic information. Instead of getting tests reactively, on a doctor's orders, people will use the data proactively to help them make decisions about their own health.

With a few exceptions, the genetic tests used today detect only uncommon forms of disease. The tests identify rare variants in a single gene that causes the disease.

But most diseases aren't caused by variants in a single gene. Often a hundred or more changes in genetic letters collectively indicate the risk of common diseases like heart attack, diabetes, or prostate cancer. Tests for these types of changes have recently become possible, and they produce what is known as your "polygenic" risk score. Polygenic risk scores are derived from the combination of these variants, inherited from your mother and father, and can point to a risk not manifest in either parent's family history. We've learned from studies of many polygenic risk scores for different diseases that they provide insights we can't get from traditional, known risk factors such as smoking or high cholesterol (in the case of heart attack). Your polygenic score doesn't represent an unavoidable fate—many people who live into their 80s and 90s may harbor the risk for a disease without ever actually getting it. Still, these

scores could change how we view certain diseases and help us understand our risk of contracting them.

Genetic tests for rare forms of disease caused by a single gene typically give a simple yes or no result. Polygenic risk scores, in contrast, are on a spectrum of probability from very low risk to very high risk. Since they're derived from combinations of genome letter changes that are common in the general population, they're relevant to everybody. The question is whether we'll find a way to make proper use of the information we get from them. Can they inform us about changes to our lifestyle, or point to medications we should take or a screening test we should get, that might improve our chances of staying healthy?

Statin drugs are a good case study for this. They're widely used, even though 95% of the people taking them who haven't had heart disease or stroke get no benefit aside from a nice cholesterol lab test. We can use a polygenic risk score to reduce unnecessary statin use, which not only is expensive but also carries health risks such as diabetes. We know that if you are in the top 20% of polygenic risk for heart attack, you're more than twice as likely to benefit from statins as people in the bottom 20%; these people can also benefit greatly from improving their lifestyle (stop smoking, exercise more, eat more vegetables). So knowing your polygenic risk might cause you to take statins but also make some lifestyle changes. (And a recent large-scale study in Finland showed that people with high heart-risk scores responded with lifestyle improvements at a much higher rate than those with low risk scores.)

And it's not just about heart disease. A polygenic risk score might tell you that you're at high risk for breast cancer and spur you to get more intensive screening and avoid certain lifestyle risks. It might tell you that you're at high risk for colon cancer, and therefore you should avoid eating red meat. It might tell you that you're at high risk for type 2 diabetes, and therefore you should watch your weight.

Yet despite growing evidence that polygenic risk scores are important, until recently there was no service allowing people to determine their own scores, even if they had invested in their own personal direct-to-consumer genetic profiling. We're attempting to remedy that through the development of MyGeneRank, a free mobile app that estimates users' polygenic risk for heart attack and stroke from their own genetic data. It also allows them to participate in a clinical trial to measure the influence of polygenic risk information on people's behavior, as reported by them, and their heath data, captured by mobile sensors linked to their smartphones.

A polygenic risk score might tell you that you're at high risk for breast cancer and spur you to get more intensive screening.

There are still some issues and controversies we need to deal with. Equal access is one major concern—especially given that the majority of genetic studies have been performed in populations of European ancestry. For now, it appears that the more powerful the predictions become, the less accurate they become with other populations.

In addition, genetic risk information is likely to make some people feel anxious or fatalistic (or might give others a false sense of security). Previous studies suggest that genetic risk information has a minimal influence on these psychological states, but many of those studies were done when the variations in risk you could get via polygenic factors were marginal. As our ability to separate people into increasingly different classes of genetic risk gets better, these issues may become more prominent.

Another challenge will be to convince people to forgo or delay medical interventions if they have a low risk of a certain condition. This will require them to agree that they're better off accepting a very low risk of a catastrophic outcome rather than needlessly exposing themselves to a medical treatment that has its own risks. People tend to overestimate the likelihood of catastrophic events, so if polygenic scores are to achieve their full impact on health outcomes and healthcare spending, we'll need to find a way to effectively communicate those trade-offs.

And finally there are the privacy concerns. We need to maintain our current protections against genetic discrimination so that people can benefit from their own genetic information without having to worry that insurance companies will get access to that information and use it to raise their rates or deny coverage.

You can't change your genetic risk. But you can use lifestyle and medical interventions to offset that risk. We can accelerate breast cancer screening for women with a high risk for the disease, and help people with borderline risk of heart disease to make decisions about whether to take statins or not. If we deliver and track the response to polygenic risk information, we can collect real-world evidence on how to optimize the use of that data to give safe and effective health advice.

In the near future your smartphone might feature technologies that monitor your physiological, genetic, environmental, and behavioral characteristics. And this information could be linked to virtual medical coaches and AI systems that can synthesize all that information and deliver you insights about your own health, on demand.

Ali Torkamani is director of genomic informatics at the Scripps Research Translational Institute. Eric Topol is a cardiologist and the author of books including the upcoming Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again.

Dr. Bot will see you now

An AI chatbot might help you avoid having to make an appointment with your overworked physician.

by Douglas Heaven

illustrated by Nicole Ginelli

"My stomach is killing me!"

"I'm sorry to hear that," says a female voice. "Are you happy to answer a few questions?"

And so the consultation begins. Where's the pain? How bad is it? Does it come and go? There's some deliberation before you get an opinion. "This sounds like dyspepsia to me. Dyspepsia is doctor-speak for indigestion."

Doctor-speak, maybe, but it's not a doctor speaking. The female voice belongs to Babylon, part of a wave of new AI apps designed to relieve your doctor of needless paperwork and office visits—and reduce the time you have to wait for medical advice. If you're feeling unwell, instead of calling a doctor, you use your phone to chat with an AI.

The idea is to make seeking advice about a medical condition as simple as Googling your symptoms, but with many more benefits. Unlike self-diagnosis online, these apps lead you through a clinical-grade triage process—they'll tell you if your symptoms need urgent attention or if you can treat yourself with bed rest and ibuprofen instead. The tech is built on a grab bag of AI techniques: language processing to allow users to describe their symptoms in a casual way, expert systems to mine huge medical databases, machine learning to string together correlations between symptom and condition.

Babylon Health, a London-based digital-first health-care provider, has a mission statement it likes to share in a big, bold font: to put an accessible and affordable health service in the hands of every person on earth. The best way to do this, says the company's founder, Ali Parsa, is to stop people from needing to see a doctor.

40,000 people in London have used the Babylon app.

When in doubt, the apps will always recommend seeking a second, human opinion. But by placing themselves between us and medical professionals, they shift the front line of health care. When the Babylon Health app started giving advice on ways to self-treat, half the company's patients stopped asking for an appointment, realizing they didn't need one.

Babylon is not the only app of its kind—others include Ada, Your.MD, and Dr. AI. But Babylon is the front-runner because it's been integrated with the UK's National Health Service (NHS), showing how such tech could change the way health services are run and paid for. Last year Babylon started a trial with a hospital trust in London in which calls to the NHS's non-emergency 111 advice line are handled partly by Babylon's AI. Callers are asked if they want to wait for a human to pick up or download the Babylon-powered "NHS Online: 111" app instead.

Around 40,000 people have already opted for the app. Between late January and early October 2017, 40% of those who used the app were directed to self-treatment options rather than a doctor—around three times the proportion of people who spoke to a human operator. But both the AI and the humans staffing

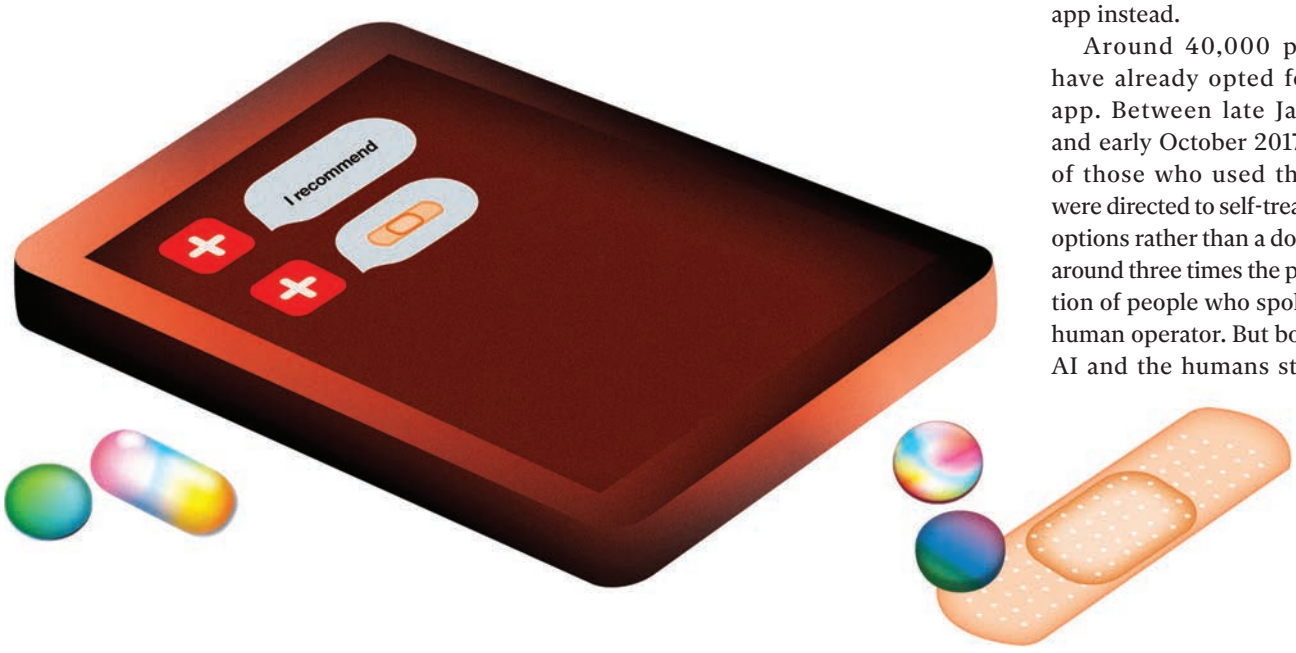

the phone line told the same proportion of people to seek emergency care (21%).

Now Babylon has also co-launched the UK's first digital doctor's practice, called GP at Hand. People in London can register with the service as they would with their local doctor. But instead of waiting for an appointment slot and taking time off work to see a physician in person, patients can either chat with the app or talk to a GP at Hand doctor on a video link. And in many cases the call isn't needed. The human doctor becomes your last resort rather than your first.

GP at Hand has proved popular; some 50,000 people registered in the first few months, among them Matt Hancock, the UK health minister. Babylon now wants to expand across the UK. The service is also available in Rwanda, where 20% of the adult population has already signed up, according to Mobasher Butt, a doctor and a member of Babylon's founding team. And it's setting up services in Canada, with plans to do the same in the US, the Middle East, and China.

Your doctor is overloaded

For 70 years, the NHS has provided free medical care to anyone who needs it, paid for by UK taxpayers. But it is showing signs of strain. Two generations ago there were 50 million Britons, and their average life expectancy was not much over 60 years. There are now 66 million, and most can expect to live into their 80s. That stretches the resources of a system that has never been flush with cash.

On average, people in the UK see a doctor six times a year, twice as often as a decade ago. From 2011 to 2015, the average GP clinic's patient list grew by 10% and its number of contacts with patients (by phone or in person) grew by 15.4%, according to a survey by the King's Fund. In a survey by the British Medical Association in 2016, 84% of general practitioners said they found their workload either "unmanageable" or "excessive," with "a direct impact on the quality" of care they gave their patients.

In turn, people often have to wait days to get a non-urgent consultation. Many show up at hospital emergency departments instead, adding even more strain to the system. "We have the perception that it's older people who turn up [at the emergency room]," says Lee Dentith, CEO and founder of the Now Healthcare Group, a health-tech company based in Manchester, UK. "But it's not. It's the 18- to 35-year-olds who are unwilling to wait a week for an appointment."

Population and life expectancy will continue to grow. By 2040, it is estimated, the UK will have more than 70 million people, one in four of whom will be over 65. Most other rich countries are also getting older.

At the same time, the next few decades will see more people living with long-term illnesses such as diabetes and heart disease. And better treatment for diseases like cancer means millions more people will be living with or recovering from them.

Of course, the UK is not alone. Whether because of prohibitive costs in the US or the lack of medical professionals in Rwanda, "all health systems around the world are stretched," says Butt. "There's not enough clinical resources. There's not enough money."

When the app started giving advice on ways to self-treat, half of patients stopped asking for an appointment, realizing they didn't need one.

Which is where companies like Babylon come in. A chatbot can act as a gatekeeper to overworked doctors. Freeing up even more of the doctor's time, the AI can also handle paperwork and prescriptions, and even monitor care at home.

A chatbot can also direct people to the right provider. "A GP is not always the best person to see," says Naureen Bhatti, a general practitioner in East London. "A nurse might be better at dressing a wound, and a pharmacist might be better for advice about a repeat prescription. Anything that helps unload a very overloaded system, allowing doctors to do what they are best at, is always welcome."

Sometimes AI is just better

Bhatti remembers how upset lots of doctors were when patients first started bringing in printouts from their own web searches. "How dare they try and diagnose themselves!

"How do we make this a job that people want to do? I don't think ... consulting from their kitchen is why people get into medicine. They come to meet patients."

Don't think you can negate my six years at medical school with your one hour on the internet." But she likes to see it from the patients' perspective: "Well, don't think you can negate my six years of living with this illness with your one-hour lecture at medical school."

When a patient does meet a doctor face to face, the AI can still help by suggesting diagnoses and possible treatments. This is useful even when a doctor is highly skilled, says Butt, and it's "really critical" in poorer countries with a shortage of competent doctors.

AI can also help spot serious conditions early. "By the time most diseases are diagnosed, a £10 problem has become a £1,000 one," says Parsa. "We wait until we break down before going to a doctor." Catching a disease early slashes the cost of treating it.

These apps first hit the market as private health services. Now they are starting to integrate with national healthcare providers and insurers. For example, Ada users can share their chatbot sessions with their NHS doctor, and the company is now working with a handful of GP practices to enable the chatbot to refer them to the doctor. Another app, Now Patient, provides video consultations with your existing doctor, and it also acts as an AI pharmacist. Users can buy their drugs from the Now Healthcare Group's drug-delivery service. It's a kind of Amazon for medicines.

"This is a service that patients really want, that they didn't previously have, and that is now being provided to them through the NHS 365 days a year, 24 hours a day, for free," Butt says of Babylon. "And the brilliant thing is it doesn't cost the NHS a single penny more to deliver that."

Not only will the AI in these apps get smarter; it will get to know its users better. "We're building in the ability for patients to manage their health not only when they're sick, but also when they're not sick," says Butt. The apps will become constant companions for millions of us, advising us and coaxing us through everyday health choices.

Death by chatbot?

Not everyone is happy about all this. For a start, there are safety concerns. Parsa compares what Babylon does with your medical data to what Facebook does with your social activities—amassing information, building links, drawing on what it knows about you to prompt some action. Suggesting you make a new friend won't kill you if it's a bad recommendation, but the stakes are a lot higher for a medical app.

According to Babylon, its chatbot can identify medical conditions as well as human doctors do, and give treatment advice that's safer. In a study posted online in June and coauthored with researchers at Imperial College London, Stanford University, and the Northeastern Medical Group, Babylon put its AI through a version of the final exam of the Royal College of General Practitioners (RCGP), which British GPs must pass in order to practice unsupervised. Babylon's AI scored 81%, 9% higher than the average grade achieved by UK medical students.

The RCGP was quick to distance itself from Babylon's hype, however. "The potential of technology to support doctors to deliver the best possible patient care is fantastic, but at the end of the day, computers are computers, and GPs are highly trained medical professionals: the two can't be compared and the former may support but will never replace the latter," said RCGP vice chair Martin Marshall in a statement. "No app or algorithm will be able to do what a GP does."

Others level far more serious charges, suggesting that Babylon has focused on making its service accessible and affordable at the expense of patients' safety. One Twitter user with the handle DrMurphy11 (he's an NHS consultant who told me he needs to remain anonymous because of the corporate culture there) has coined the hashtag #DeathByChatbot. In videos showing interactions with the app, DrMurphy11 suggests that Babylon's AI misses obvious diagnoses and fails to ask the right questions. "I have no concerns about health tech or AI in general,"

he says. "No doctor wants to make mistakes, and any system that helps minimize the risk of harm from human error will be welcomed." But he's worried that companies are misleading doctors and the public with marketing claims that vastly oversell their current tech.

Babylon has also met with criticism in Rwanda, where it runs the Babyl service, for not taking local epidemiology into account. In an interview with the BBC, Rwanda's minister of health claimed that the Babyl app included no questions about malaria, for example (although Babylon disputes this).

Still, while Babylon may not be as good as a real doctor (and such apps are always careful to recommend you see a real doctor when in doubt), playing it too safe would defeat the purpose. "We wanted to re-create the same pragmatic approach that a clinician takes," says Butt. "If we just had a group of non-clinical people building the service, they might have gone for something that was 100 percent safe, but that could mean you send everyone to hospital, which is not what a real doctor or nurse would do."

Another fear is that digital-first services will create a two-tiered health-care system. For example, GP at Hand advises people with serious medical issues to think twice about signing up to a practice that offers mostly remote access to doctors. That might seem prudent, but it has led to accusations that GP at Hand is effectively cherry-picking younger patients with less complex—and less expensive—health-care needs. Since British GP practices get per-patient funding from the NHS, cherry-picking would mean the rest of the health-care system is left to do more with less.

For some GPs, this isn't acceptable. "We take everybody," says Bhatti. But Oliver Michelson, a spokesperson for the NHS, accepts that GP at Hand has to issue some form of caveat—it can't realistically welcome everyone. "They are not denying people access but saying that if you're going to need to come into your GP regularly, a digital-first service may not be the best place to be," he says.

And Butt insists that they exclude nobody. "The service is available to everyone," he says; it just may not suit some people, such as those with severe learning difficulties or visual impairments, who would struggle with the app.

People still come in handy

For Bhatti, having a local doctor who knows you is a crucial part of the health system. "Knowing your doctor saves lives," she says. "Doctors will pick up things because there's continuity." She thinks this is just as much an issue for doctors as for patients. "How do we make this a job people want to do?" she says. "I don't think people working flexibly, consulting from their kitchen, is why people come to medicine. They come to meet patients."

Not even Butt envisions chatbots replacing human doctors entirely. "Care is not just about diagnosing or prescribing medicine," he says. "It's about knowing your patient is going to be able to cope with the chemotherapy you're proposing for them, knowing that their family will be able to offer them the support that they're going to need for the next few months. Currently there is no software that's going to be able to replace that."

Digital-first services may scare off sicker patients.

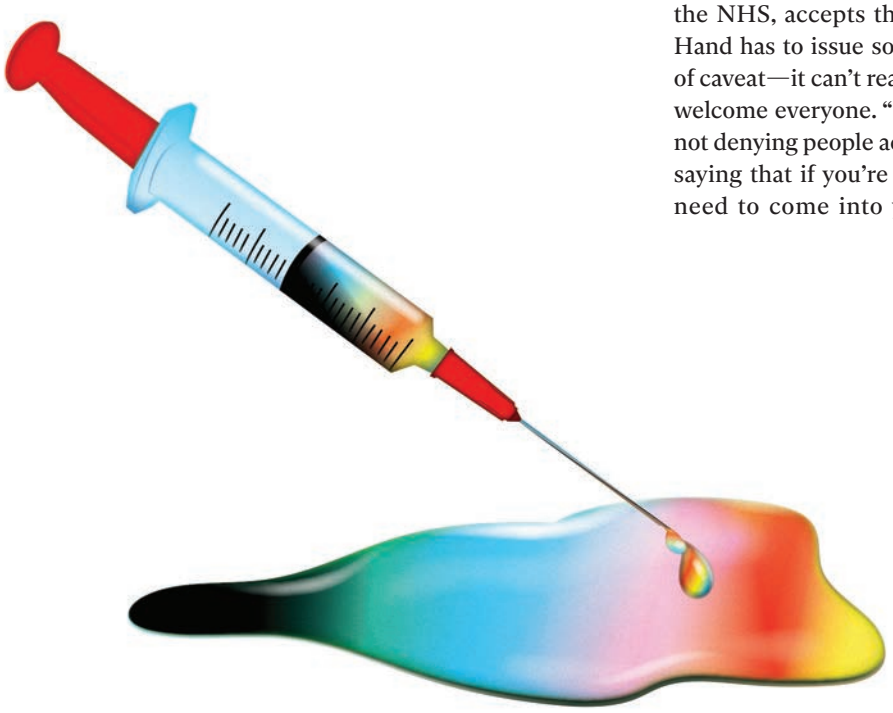

Douglas Heaven is a freelance writer based in London. His most recent story for MIT Technology Review was "Can you spot the cryptocrime in this picture?" in our May/June issue.

AI can't replace doctors. But it can make them better.

A machine can collate environmental data, genetic data, and patient history way better than I can.

By **Rahul Parikh**

Several years ago Vinod Khosla, the Silicon Valley investor, wrote a provocative article titled "Do We Need Doctors or Algorithms?" Khosla argued that doctors were no match for artificial intelligence. Doctors banter with patients, gather a few symptoms, hunt around the body for clues, and send the patient off with a prescription. This sometimes (accidentally, maybe) leads to the correct treatment, but doctors are acting on only a fraction of the available information. An algorithm, he wrote, could do better.

I'm a pediatric and adolescent physician in the San Francisco Bay Area, where entrepreneurs like Khosla have been knocking on the doors of doctors for years with their pilot technologies and software and hardware. I can say with some authority that Khosla's is the voice of a savvy outsider who knows what he knows—which isn't health care.

Yes, AI could help us diagnose and treat disease. It can collate and serve up broad swaths of data in a clear and concise way, cutting down on the imprecise judgments that doctors make because of the pressures and complexity of our practices. There's no doubt that for certain doctors, whose work is highly focused on diagnosis (radiologists or pathologists, for example), that breakthrough may prove an existential threat. A decade ago, for example, researchers showed that AI was as good as radiologists at detecting breast cancer.

But for physicians like me in primary care, managing 1,500 to 2,000 patients, AI presents an opportunity. I went to medical school to connect with people and make a difference. Today I often feel like an overpaid bookkeeper instead, taking in information and spitting it back to patients, prescribing drugs and adjusting doses, ordering tests. But AI in the exam room opens up the chance to recapture the art of medicine. It could let me get to know my patients better, learn how a disease uniquely affects them, and give me time to coach them toward a better outcome.

Consider what AI could do for asthma, the most common chronic medical disease in childhood. Six million American kids suffer from it. In 2013, they collectively missed 14 million days of school. The cost of medications, visits to the doctor and emergency room, and hospitalizations nears $60 billion a year.

I diagnose asthma via a rule of thumb that's been handed down over time: if you've had three or more wheezing episodes and the medicines for asthma help, you have the disease. Once it's diagnosed, I ask the parents to remember—as best they can—how often they administer medicines to their child. I ask: What seems to trigger episodes? Is the child exposed to anyone who smokes at home? I can also review their records to count how many visits to the emergency room they've had, or the number of times they've refilled their prescriptions.

It's not that we don't have the data; it's just that it's messy. We spend a great deal of our time trying to make sense of it.

But even with the most accurate recall by parents and patients, and the most accurate electronic records, it's still just retrospective knowledge. There's no proactive, predictive strategy.

It's not that we don't have the data; it's just that it's messy. Reams of data clog the physician's in-box. It comes in many forms and from disparate directions: objective information such as lab results and vital signs, subjective concerns that come in the form of phone messages or e-mails from patients. It's all fragmented, and we spend a great deal of our time as physicians trying to make sense of it. Technology companies and fledging startups want to open the data spigot even further by letting their direct-to-consumer devices—phone, watch, blood-pressure cuff, blood-sugar meter—send continuous streams of numbers directly to us. We struggle to keep up with it, and the rates of burnout among doctors continue to rise.

How can AI fix this? Let's start with diagnosis. While the clinical manifestations of asthma are easy to spot, the disease is much more complex at a molecular and cellular level. The genes, proteins, enzymes, and other drivers of asthma are highly diverse, even if their environmental triggers overlap. A number of experts now think of asthma in the same way they think of cancer—an umbrella term for a disease that varies according to the tumor's location and cellular characteristics. Ian Adock of the National Heart & Lung Institute at Imperial College, London, studies the link between asthma and the environment. He and his team have been collecting biological samples from asthma patients' blood, urine, and lung tissue and organizing the genetic and molecular markers he finds into subtypes of asthma. The hypothesis is that with that kind of knowledge, patients can be given the drug that works best for them.

AI might also help to manage asthma flares. For many patients, asthma gets worse as air pollution levels rise, as happened this past summer when brush fires swept through Northern California. AI could let us take environmental information and respond proactively. In 2015, researchers published a study showing they could predict the number of asthma-related emergency room visits to a Dallas–Fort Worth hospital. They pulled data from patient records, along with air pollution data from EPA sensors, Google searches, and tweets that used terms like "wheezing," or "asthma." The Google and Twitter data were tied to the user's location data.

If I had this kind of data I could say, "Alexa, tell me which asthma patients I need to worry about today." I could give a heads-up to the affected families. And if I also had some genetic data like Adock's, I could diagnose asthma before the patient suffered three bouts of wheezing, by ordering blood tests and comparing the results against those molecular markers.

This kind of time-saving intelligence frees me to spend more time with my patients. One study showed that asthmatic children only took or received their inhaled medications about half of the time. AI might allow me more time to personally interact with those kids, and get better results.

The author recently came across this drawing by a seven-year-old depicting her idea of a visit to the doctor—she's on the exam table, and the doctor's facing the other way.

DRAWING BY AG, AGE 7, COPYRIGHT THOMAS G. MURPHY MD 2011

Lots of questions lie ahead. Are patients willing to share more of their personal data with us? If the AI shows your care is better one way, but you or your doctor feel differently, will an insurance company accept it? What if the algorithm misses something or is applied incorrectly? Who is liable, the doctor or the machine's maker?

Not long ago, in the *Journal of the American Medical Association,* I saw a colorful picture drawn by a child in crayon. It portrayed her pediatrician, eyes glued to the computer, while she sat on the exam table, looking wide-eyed. I hope that AI will soon allow me to turn my attention back to that little girl.

Rahul Parikh is a pediatrician in the San Francisco Bay area.

Making genomic medicine relevant for everyone

Carlos D. Bustamante's hunt for genetic variations between populations could correct health disparities and drive drug discovery.

By David Rotman
Photographs by Christie Hemm Klok

In the 15 years since the Human Genome Project first exposed our DNA blueprint, vast amounts of genetic data have been collected from millions of people in many different parts of the world. Carlos D. Bustamante's job is to search that genetic data for clues to everything from ancient history and human migration patterns to the reasons people with different ancestries are so varied in their response to common diseases.

Bustamante's career has roughly spanned the period since the Human Genome Project was completed. A professor of genetics and biomedical data science at Stanford and 2010 winner of a MacArthur genius award, he has helped to tease out the complex genetic variation across different populations. These variants mean that the causes of diseases can vary greatly between groups. Part of the motivation for Bustamante, who was born in Venezuela and moved to the US when he was seven, is to use those insights to lessen the medical disparities that still plague us.

But while it's an area ripe with potential for improving medicine, it's also fraught with controversies over how to interpret genetic differences between human populations. In an era still obsessed with race and ethnicity—and marred by the frequent misuse of science in defining the characteristics of different groups—Bustamante remains undaunted in searching for the nuanced genetic differences that these groups display.

Perhaps his optimism is due to his personality—few sentences go by without a "fantastic" or "extraordinarily exciting." But it is also his recognition as a population geneticist of the incredible opportunity that understanding differences in human genomes presents for improving health and fighting disease.

David Rotman, *MIT Technology Review*'s editor at large, discussed with Bustamante why it's so important to include more people in genetic studies and understand the genetics of different populations.

How good are we at making sure that the genomic data we're collecting is inclusive?

I'm optimistic, but it's not there yet.

In our 2011 paper, the statistic we had was that more than 96% of participants in genome-wide association studies were of European descent. In the follow-up in 2016, the number went from 96% to around 80%. So that's getting better. Unfortunately, or perhaps fortunately, a lot of that is due to the entry of China into genetics. A lot of that was due to large-scale studies in Chinese and East Asian populations. Hispanics, for example, make up less than 1% of genome-wide association studies. So we need to do better. Ultimately, we want precision medicine to benefit everybody.

Aside from a fairness issue, why is diversity in genomic data important? What do we miss without it?

First of all, it has nothing to do with political correctness. It has everything to do with human biology and the fact that human populations and the great diaspora of human migrations have left their mark on the human genome. The genetic underpinnings of health and disease have shared components across human populations and things that are unique to different populations.

How does that play out?

Diabetes is a great example. If we look at the genetics of diabetes, they are different in different parts of the world. In the early 2010s, the Broad [Institute of MIT and Harvard] did a study with the National Institute of Genomic Medicine in Mexico to study the genetics of diabetes. Sure enough, they found a genetic variant that has a 25% frequency in Mexico that you don't see in European, East Asian, or African populations. It is largely seen only in the Americas, and it underscores a large part of ethnic disparity in diabetes.

We've done research on seemingly innocuous traits like blond hair. There is no more striking phenotype. Some people have blond hair and some people don't. And the cause of blond hair in Melanesia is completely different from the cause in Europe—and that's blond hair. So why do you think diabetes, heart disease, all these other complex traits will have identical causes in all humans? It doesn't make sense.

It turns out the highest prevalence of asthma [in the US] is in individuals

of Puerto Rican ancestry, followed by individuals of African-American ancestry, followed by European ancestry. The people with the lowest rate of asthma are those of Mexican ancestry. You have two of the Hispanic populations at the opposite ends of the spectrum.

Why is detailing these genetic differences helpful for medicine?
If the genetic etiology of disease is different, it gives us an opportunity to discover new drug targets. It gives us new biology that then can be used even for those that don't necessarily suffer from the disease in that way. It's important for drug discovery. If you think of it like looking for oil, we've only been looking for oil in the North Sea. There are plenty of other places to search, and that benefits everyone.

Secondly, we're finding that polygenic risk scores [disease-risk predictions based on genetic tests] for European ancestry don't translate easily into other populations. If we don't have broad representation in medical and population genetics, then we run the risk of widening health disparities, which will be a terrible outcome for precision medicine and precision health.

So aren't you disappointed by the lack of progress in including more populations in genomic data?
I'm actually super-excited. We've done a great job of mining for drug targets in Europe. Iceland led the way, Britain led the way, and now Finland. So we're tapping all those resources—awesome. But what about Latin America? What about Africa? What about South Asia? All of those places have tons to contribute to our understanding of health and disease.

It is both a moral obligation and a missed scientific opportunity if we don't go to work in those populations.

Many genetic researchers have long argued that race has no basis in science. But the debate doesn't seem to go away.

"We can't use genetics for the purpose of trying to define the stories we tell about ourselves."

In a global context there is no model of three, or five, or even 10 human races. There is a broad continuum of genetic variation that is structured, and there are pockets of isolated populations. Three, five, or 10 human races is just not an accurate model; it is far more of a continuum model.

Humans are a beautifully diverse species both phenotypically and genetically. This is very classic population genetics. If I walk from Cape Horn all the way to the top of Finland, every village looks like the village next to it, but at the extremes people are different.

But as a population geneticist?
I don't find race a meaningful way to characterize people.

You walk a tricky line, though, don't you? You're pointing out the importance of variance between different populations, but you don't want to reinforce old categories of race.

We can't use genetics for the purpose of trying to define the stories we tell about ourselves. Social determinants of health are often far more important than genetic determinants of health, but that doesn't mean genetic determinants aren't important. So you've got to embrace the complexity and figure out how we translate this to a broad general public.

I'm actually an optimist. I think the world is becoming a less racist place. If you talk to the next generation of people, millennials on down, those abhorrent ideologies are thrown away. That means it gives us a space to now think about what role does genetics play in health and diseases and human evolution in ways that we can soberly understand and bring to bear on important problems.

We can't allow genetics to get hijacked by identity politics. If you begin to allow politics and other interests to come in, you just muddy the waters. You need to let the data lead. You need to let outcomes lead. And the rest will follow.

DATA BIAS IN DNA STUDIES

Precision medicine is getting more precise for some but leaving many others behind. And those left behind are often people with Latin American, African, Native American, and other ancestries that are underrepresented in genomic databases.

By far, most of the data in genome-wide association studies, which have been critical in spotting genetic variants tied to common diseases, comes from people with European ancestry. In 2011, Carlos D. Bustamante and his colleagues called out the disparities and the resulting threat that genomic medicine "will largely benefit a privileged few." In subsequent years, the collection of genomic data has exploded, but the disparities remain. In 2016, Alice Popejoy, who was a PhD student at the University of Washington and is now a postdoc in Bustamante's lab, updated the results in the journal *Nature*, finding little progress for most population groups.

One result of this lack of data is that genetic tests may be less relevant and accurate for people from underrepresented groups. Increasingly popular consumer genetic tests can be misleading or just plain wrong, and medical genetic tests for some common diseases are often inconclusive. Likewise, Popejoy says, false positives and false negatives in genetic diagnoses are more common in people with non-European ancestry, because the results are interpreted using databases that are incomplete or biased toward European ancestry.

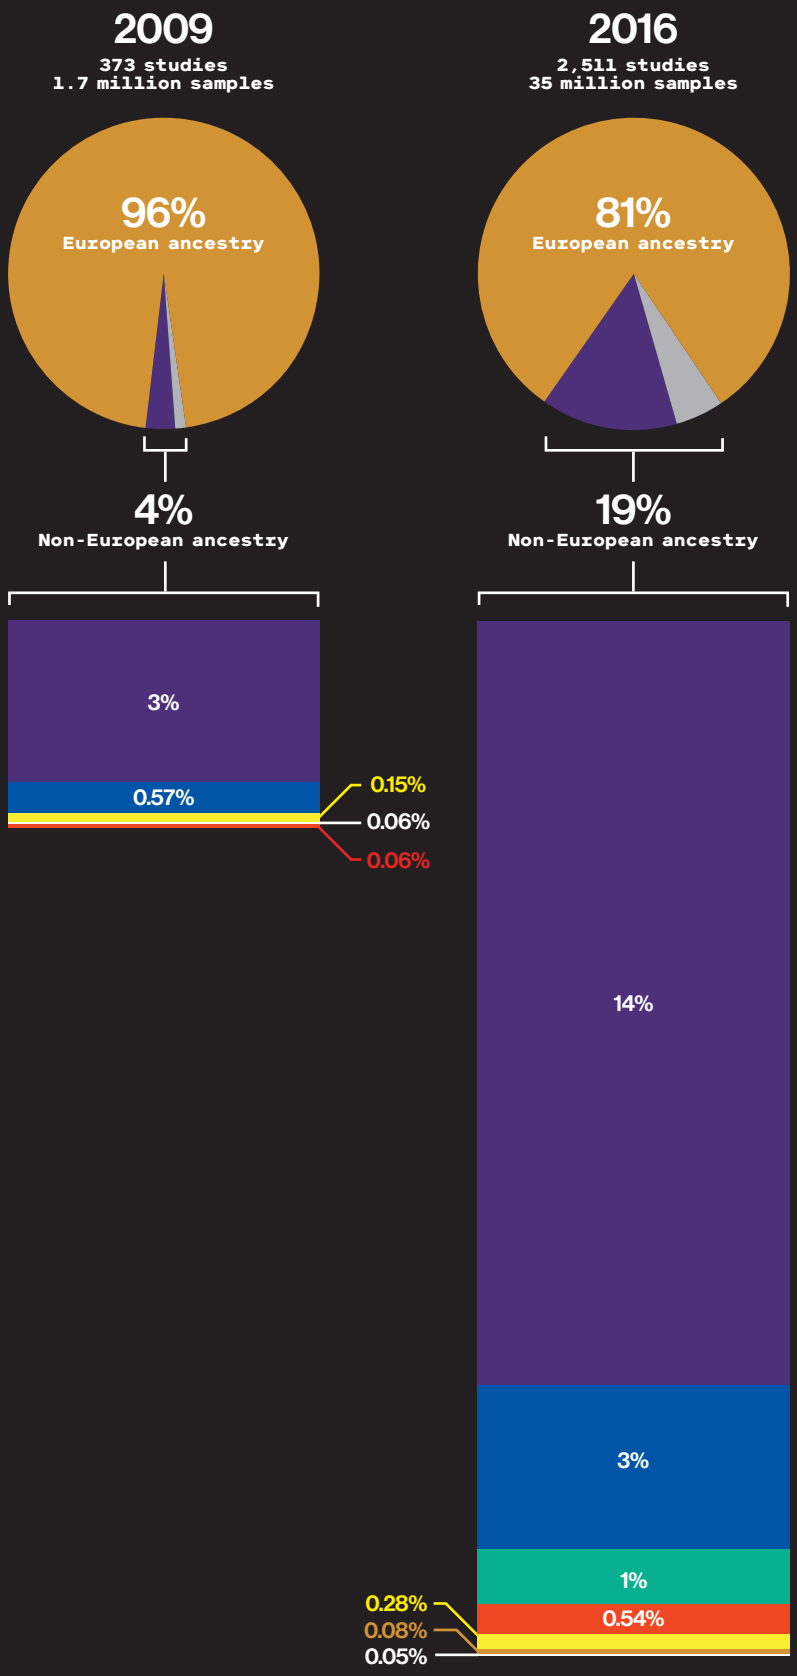

POPEJOY AND FULLERTON, *NATURE* 538 (2016); BUSTAMANTE, DE LA VEGA, AND BURCHARD, *NATURE* 475 (2011)

When you're sick, you're vulnerable—and that's when your doctor pressures you into participating in a data-gathering experiment.

By **Mary Madden**

Need medical help? Sorry, not until you sign away your privacy

Last summer I found myself running late for a doctor's appointment I'd waited months to get. Even though the back injury I had sustained three months earlier was finally starting to improve, I was eager to get an expert opinion from an orthopedic surgeon. When I arrived, breathless and apologetic, the doctor's office was filled with patients—many with much more serious injuries than mine—who had also waited months to see the renowned specialist. As I was about to take my seat, I was called back to the front desk: Could I also please answer some questions about my personal health history using the office's new tablet-based system?

As a social science researcher who has studied digital privacy and security issues for much of my career, I was less than thrilled to be a guinea pig for their new data-management system. But … I had waited so long for this appointment, and I had already kept the doctor waiting, and maybe this would save me time at future appointments with other doctors? At that moment, as if in response to my frustrated realization that there was no clear way to opt out and still receive the care I needed, my back muscles tightened up.

Mary Madden is a technology researcher and writer. She leads a project with the Data and Society Research Institute to understand the social and cultural impacts of data-driven technologies on health equity and well-being.

I nodded politely and brought the tablet back to my chair. From the institutional perspective, this was a totally reasonable request for verification. But it was also a clear instance of surveillance, and the power dynamics between me and the administrative authority were not at all equal. I was in pain and in no mood to argue.

By agreeing to use the tablet, I'd already consented to a form of data collection I wasn't entirely comfortable with. I had never heard of the branded tablet the office was using, and the logo assuring me that it was "antibacterial" didn't ease my concerns about letting scores of other patients handle a device into which I'd put my private data. The awkward software interface did little to suggest that my data would be dealt with carefully; worse than the clunky visual design, there was no indication of whether or not the tablet was internet-connected, and there was no explanation of how my data would be stored or protected once it entered their system.

So what did I do? I dutifully entered my info anyway—immediate physical needs have a way of leapfrogging over data privacy concerns, even for people like me who feel strongly about maintaining control over how their information is collected and used.

Not the first time this happened

As I scrambled to consult my phone for records of my grandparents' cause of death and the appropriate medical term to describe the blood condition that runs in my family, I realized that this was probably the fourth time over the past year that I'd been asked to enter some version of this data digitally in other systems—in addition to various paper versions of the same information. Instead of making the patient experience more efficient and less stressful, it made me feel as though doctor's offices were crowdsourcing their work to stressed-out patients with little explanation of why.

When I'd finished digitally detailing my health history, the final screen seemed to mock me with one last request: Could

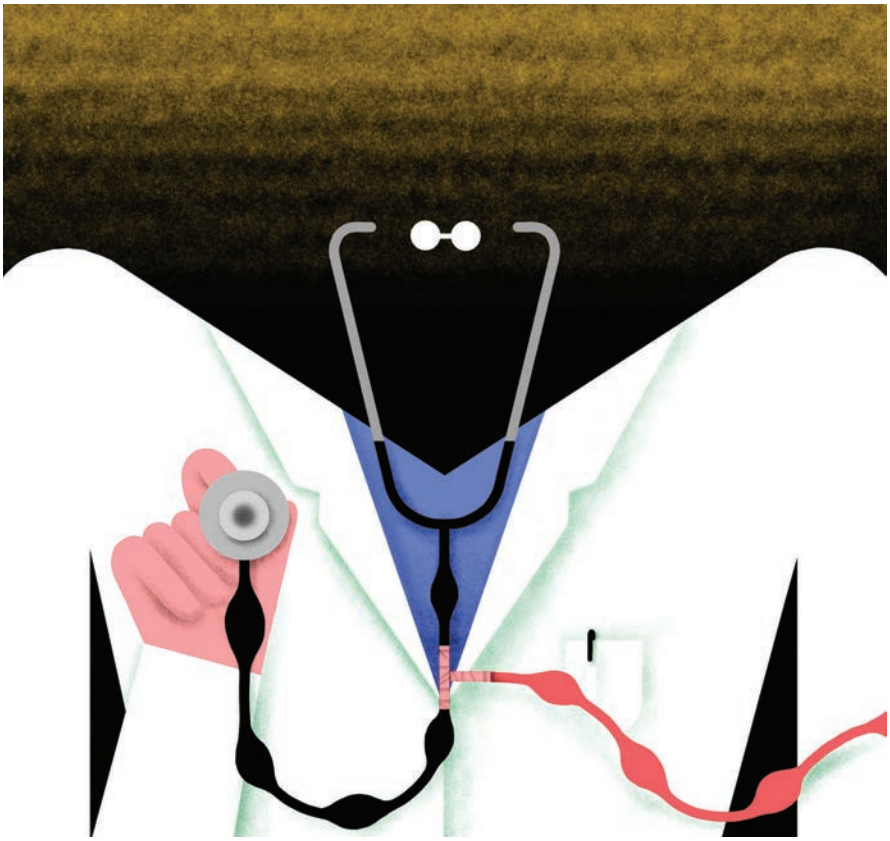

Opt in and you risk losing control over how your health data is used. Opt out and you risk losing access to the medical care you need.

I please acknowledge that I'd received a copy of the office's privacy practices? (I hadn't.) But what were the consequences of opting out at this point? And what about people who were much less comfortable with technology than I was? How were they dealing with questions or concerns about this process?

The banality of Big Brother

In the internet age, it's become repetitive and banal to simply agree to terms of service that we don't fully understand. And while it would be nice to think that my doctors and their third-party software vendors will forever treat my health data with the utmost care, the reality is that digital health data systems have been vulnerable to numerous ransomware attacks, genetic testing companies have opened up their customers' data to use by pharmaceutical companies, and the market for health data is massive and growing.

I've spent more than a decade studying Americans' attitudes to different kinds of digital information, and I have seen repeatedly that health data is one of the most sensitive categories. In a study I contributed to at the Pew Research Center, respondents were asked whether they would participate in a web-based system that their doctor's office used to manage patient records. Even in this scenario (which notably involved a much more transparent system than the one I'd used at the orthopedic surgeon's office), only a little more than half of American adults definitively said they'd be comfortable sharing their data.

Health data is one of the few categories of information that enjoy a robust (if outdated) set of privacy protections by law in the US, but the definition of what even counts as health data is rapidly evolving. More and more companies are looking to use diagnostic insights from social-media data and other nonregulated categories that currently exist in the lucrative marketplace of predictive analytics. The current Wild West environment allows health data brokers to create risk scores that are sold to insurance companies that in turn use these metrics to charge higher rates to the most vulnerable among us. Not only is this bad for patient privacy, but it further exacerbates inequalities in our society.

Care shouldn't require data consent

Americans' concerns about the sanctity of their health data have been cited as one reason that Google and Apple have recently partnered with the likes of the American Heart Association and doctors from Massachusetts General Hospital. Such household names can help allay patients' fears about entrusting their data to Big Tech. But we're now at the point where the stakes are growing much higher when we make decisions to share our data with a platform or participate in a study. When we opt in, we risk losing control over how our health data is used and who can profit from it. When we opt out, we risk losing access to the care we need.

In the era of data-driven medicine, systems for handling data need to avoid anything that feels like manipulation—whether it's subtle or overt. At a minimum, the process of obtaining consent should be separated from the process of obtaining care.

If you don't want to hand over your information right away, or if you have concerns about the security of your doctor's data-gathering efforts—you should be able to see the doctor anyway.

BENEDIKT LUFT

Get MIT Technology Review delivered to your inbox.

THE DOWNLOAD
Your daily dose of what's up in emerging technology

CLOCKING IN
A daily look at the workplace of the future

CHAIN LETTER
Blockchains, cryptocurrencies, and why they matter

THE ALGORITHM
News and views on the latest in artificial intelligence

Stay in the know:
technologyreview.com/inbox

MIT Technology Review

Then there's illness

Now it's about money. Whatever ails you, there's a growing chance we can cure it—for a price. Who can pay? And should they? If some people can excise genetic disease from their bloodlines, will we end up creating two human races, one sick and one healthy? Or are both the promise and the perils of genomic medicine being wildly overhyped?

2

A cure for one

By
Antonio Regalado

Illustrations by
Sébastien Thibault

One day, gene therapy may help with the rarest of diseases. Some parents aren't waiting.

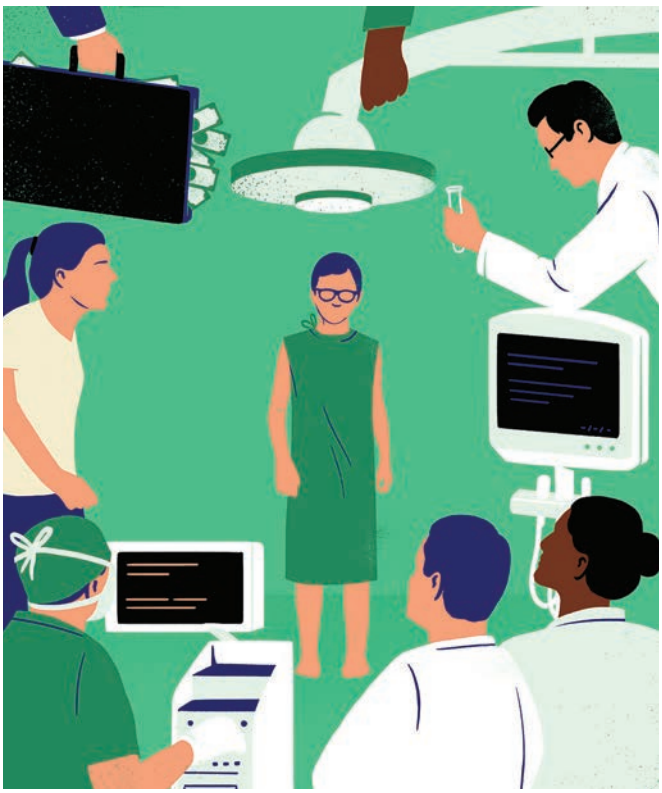

Jennie and Gary Landsman launched an appeal to save their sons on Thanksgiving of 2017. By the end of the weekend the family, who live in Marine Park, Brooklyn, had raised $200,000.

In a moving three-minute video posted online, they sit on an overstuffed leather couch. Jennie glances away from the camera, betraying little emotion, as Gary talks. "We need your help, we really do," Gary says, his voice breaking. The Landsmans' two sons—Benny, then 18 months, and Josh, four months—both have a fatal genetic brain disorder called Canavan disease. Benny, limp on his mother's lap, is already affected by nerve loss. Josh isn't yet. But he will be if nothing is done.

Canavan is an "ultra-rare" disease. So rare, in fact, that there is no reliable understanding of how many children are born with it. Relatively few researchers study Canavan, and no drugs are approved to treat it. There isn't even a single clinical trial open for some last-ditch remedy, the kind people battling cancer can turn to. Doctors told Jennie there was not much to be done. She should go home and make her boys comfortable until they died.

The Landsmans refused to accept that advice. Instead, Jennie hit Google and started e-mailing scientists. Here's what she learned: there may be a way to fix the genetic error in the boys' brains. But the family would have to pay for it themselves. And it would be expensive.

"We need one and a half million dollars, and our goal is to get it in the next six months," Jennie says in the video.

The Landsmans had discovered gene therapy, technology that uses viruses to add healthy genes to cells with defective ones. After several decades caught in scientific backwaters, gene therapy has entered a golden age. During a span of four months, from August to December 2017, the US Food and Drug Administration (FDA) approved three such therapies, two for blood cancer and one for an inherited cause of blindness. Companies are investigating treatments for hemophilia and muscular dystrophy. "Once just a theory," said the FDA's chief, Scott Gottlieb, in July, gene therapy "may have the potential to treat and cure some of our most intractable and vexing diseases."

The technology's medical logic is especially irresistible for parents of children with the rarest diseases on earth. These are the 7,000 or so conditions that typically, like Canavan, are caused by errors in a single gene. Gene therapy suggests the ultimate bug fix—just give people working DNA instructions. The problem for the Canavan kids is that there have been too few patients for anyone to bring the research out of the lab and put dollars behind it. The same is true for countless other diseases you've never heard of, some of which are known to affect fewer than 50 people on the globe.

"The simple math is that there are a very limited number of patients. That is what created this crazy, crazy paradigm," says Eric David, an executive with BridgeBio, a biotech in Palo Alto, California, that specializes in treatments for rare diseases. "Families are saying, 'Oh my God, no one is going to pay for this. I have to fund it myself.'"

Gene therapy already has a reputation as medicine's gnarliest economic case. The problem is who will pay. Even those few treatments approved for sale carry price tags as high as $1 million. Underlying the unheard-of prices is the cost of years of research, human tests, and paperwork to win the FDA's sign-off, all in tiny markets with small pools of patients. Costly, too, is the still unwieldy process of manufacturing trillions of viruses, into each of which a gene is placed so it can be conveyed into people's cells. The result? A growing gap between the list of diseases that could be treated with gene therapy and those that actually are.

I learned of six cases where parents financed clinical trials for gene therapy in which their own children were treated. These include Karen Aiach, who started a

"The simple math is that there are a very limited number of patients. That is what created this crazy, crazy paradigm."

biotechnology company, Lysogene, based outside Paris; it funded the trial in which her daughter was treated for Sanfilippo syndrome. A connected Hollywood couple, the Grays, raised $7 million to pay for a trial that infused gene-carrying viruses into two of their daughters and several other children with a rare form of Batten disease. More than 20 other parent-financed trials are in the planning phase.

Other families are avoiding the rigors of formal studies and trying to secure untested gene therapies as emergency treatment. In Florida, a single boy was treated with a Canavan gene therapy in 2017 after his parents paid for the experiment. They did it under an exemption in federal rules called "expanded access," which can allow unapproved drugs to be offered to specific patients "whose life is immediately threatened."

That experiment fell into a gray zone, not quite research and not quite medicine. It is the same pathway the Landsman family is trying to follow, with the help of Paola Leone, a gene therapist in New Jersey, and Christopher Janson, a neurologist in Chicago. Leone and Janson asked the FDA last June to permit emergency use of their own Canavan gene therapy in up to five children they have designated in advance. The first two names on the waiting list: Benny and Josh Landsman.

According to the FDA, the strategy is not as unusual as it sounds. Of the approximately 700 gene-therapy trials it oversees, 77 fall into the desperate-case category, according to an agency spokesperson. It is not known in how many of these cases the families are covering the costs, but that is entirely legal, too. "We would love to do it in a broader, systematic fashion that would lead to a drug treatment, but we don't have the money," says Janson, a physician at the University of Illinois College of Medicine. "Until then, we are stuck on our own trying to help a couple of kids."

Some scientists who know of the Landsmans' plan fear it represents a new form of boutique medicine—a way to give those with fat checkbooks or a knack for viral fund-raising campaigns special access to cutting-edge breakthroughs. A different perspective is that it's just a preview of the personalized genetic medicine that will increasingly be available more generally.

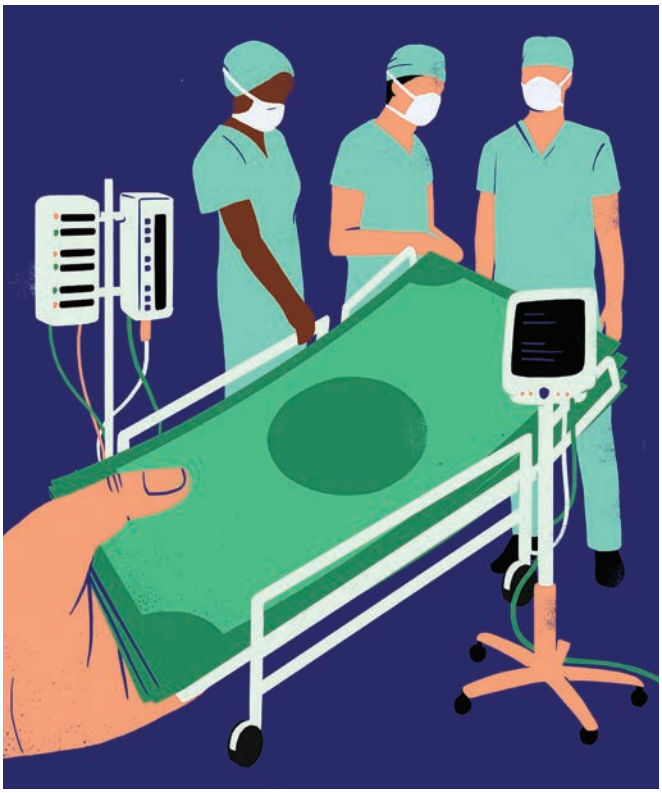

In the future, health officials believe, it could become commonplace for scientists to detect a genetic mutation and whip up a custom DNA antidote for one person. "Those 7,000 diseases are ones where we know the molecular defect for most of them. We know exactly what the initial glitch was that has led to this outcome," Francis Collins, the director of the National Institutes of Health (NIH), said in a speech this year. "Shouldn't we think about ways to do that in a fashion that scales to hundreds or maybe thousands of diseases? So what would that take?"

Nobody really knows. And the Landsmans can't wait for Washington, DC, or drug companies to figure it out. At today's rate of new drug approvals for rare diseases—about 15 a year—it could take 1,000 years for companies to get around to all of them. With two sons slipping away at home, Jennie and Gary are measuring time in months instead. Josh has a big smile but never learned to crawl. He'll soon become like Benny, who moves his arms only weakly and communicates by glancing at pictures Velcroed to a felt pad. "He's never said 'mommy,'" Jennie told me. But he can still ask for her—one of the pictures pinned up there is hers.

Jennie says she hopes that all Canavan kids will someday benefit and that the researchers helping her "will become famous." But she did not raise all that money to fund an experiment or to become a philanthropist. "This is not a clinical trial," Jennie told me. "This is private. This is for Benny and Josh."

Perfect timing

Canavan disease is rare, but it's significantly more common among people of Ashkenazi Jewish descent, like the Landsmans. Like Tay-Sachs, it's enough of a threat that prospective parents in this population are tested to see if they are silent carriers of the gene error. About one in 40 are. A series of medical miscommunications, Jennie says, led her to mistakenly believe she had tested negative. Since it takes two mutated gene copies, one from each

parent, to cause the illness, they saw no reason to test Gary.

Jennie and her pediatrician were slow to pick up on Benny's symptoms. Her sister said the toddler seemed "mushy," but the Landsmans' doctor said not to worry. By then she was pregnant with Josh. The disastrous diagnoses unfolded over a few days last summer. In late July, a blood test finally showed Benny had Canavan. Two weeks later, on Gary's birthday, they learned their newborn had it, too.

As Jennie remembers it, she spent weeks in depression, staring at sunbeams coming under the awning and trying to "live in the moment." But when I visited Leone, the gene therapist, at the Rowan University School of Osteopathic Medicine in New Jersey, she showed me e-mails Jennie had sent her between the two boys' diagnoses. "Can you help?" she had asked.

The idea of gene therapy traces back to 1970, but only recently have scientists mastered its components. In 2017, doctors at Nationwide Children's Hospital, in Ohio, described in *The Lancet* how they had prevented a group of infants from developing spinal muscular atrophy, a nerve disorder that, like Canavan, is fatal. The key elements: a virus that infects the right cells (nerves, in this case), immense doses, and timing. Give a one-month-old infant the missing gene, and the nerve damage doesn't begin. It now appears to scientists—and parents—that similar strategies must be capable of saving kids with other inherited nervous-system diseases.

Leone was a logical person for Jennie to approach. Between 2001 and 2005, Leone and Janson had, with government funding, treated 13 children with Canavan in one of the first attempts to change the genetic code inside a person's brain. At that time, scientists were unsure of the concept's potential, and their treatment, though it had some effect, was no cure.

Leone had been working toward a new Canavan gene therapy. But her last federal grant had run out in January 2018. In her lab, I saw a scientist curse at an old-model Mac that was slow to load an image. The shoestring budget is nothing new. "When I started this work," she says, "people looked at me and said, 'You must be out of your mind to work on a rare disease—you are never going to find any money and no one is going to be interested.'"

Leone keeps pictures and memorials to Canavan kids she has known in her office. Over the years, she had told many of their parents there was no chance at a cure. But the Landsmans' timing was perfect. By the fall of 2017, Leone had given the new gene therapy to enough mice to see what she calls dramatic effects. The disease seemed to have greatly slowed, even reversed. "Then I was prepared to say 'Yes' to the family that came along," she says.

"There is a lot we can do"

Leone met the Landsmans in New York, near the 9/11 memorial, in September of 2017. Gary confessed that if he had a choice between fighting and fleeing, he wanted to flee. Many parents institutionalize children with Canavan. Gary wished he could take Jennie far away and never come back. "Excruciating pain," recalls Leone. "Eyes that had cried so much they were hard to see."

Jennie wanted to know what they could do. Leone told her: "There is a lot we can do, but the first thing is how much it will cost. I can tell you it's approximately $1.5 million."

"We can do it," Jennie said without blinking.

Leone tallied the costs. They would need to hire a company to chemically synthesize healthy copies of the gene that's broken in Canavan, set aside payments for neurosurgeons, and hire consultants to prepare a request to the FDA. The biggest single expenditure would be manufacturing. Making the viral particles—they're grown in thin sheets bathed in components of cow blood—remains a delicate craft, and there are long waiting lists at production centers. Leone believed it would cost at least $1 million just to make enough virus to treat Benny, Josh, and perhaps a few others.

The Landsmans didn't have the money. The family is squarely middle class. "But there's money everywhere, isn't there?" Jennie reasoned. She was right. Their video, posted to Facebook and later GoFundMe, a crowdfunding site, went viral. By now, they've been on TV and in *People* magazine. Eight thousand donors have already

given more than $1.5 million. "This was all local people in a small Jewish community in Brooklyn," says Ilyce Randell, a Canavan patient advocate who has been in contact with the Landsmans and has funded Leone's work in the past. "It was a perfect storm—everyone rallied."

But if the Landsman children end up benefiting, she says, it will be because of research under way long before they were born. "To make it seem like they bought a cure for a million bucks—that is misleading," she says. "What is true is they came at the right time. Ten years ago you couldn't say, 'I'll raise money and get my kid treated.' Three years ago you couldn't do it. The science was not ready."

Unfair system

In August, many of the families and key researchers in the rare-disease world arrived in the security line at the NIH in Bethesda, Maryland. During a two-day meeting, cosponsored by the FDA, scientists gave talks whose topics teetered between remarkable results of tailored therapies and what the organizers called "unanswered questions" about how these could ever reach patients at affordable prices.

At today's rate of new drug approvals for rare diseases—about 15 a year—it could take 1,000 years for companies to get around to all of them.

The event attracted parents hoping to find gene-therapy specialists who would treat their children. One, Amber Freed, wore a name tag reading "*SLC6A1*," the scientific designation of a little-studied gene. Freed told a story that was by now becoming familiar. After months criss-crossing the country trying to diagnose her son's unexplained symptoms, she finally had his genome sequenced. In May, she learned he had a pathological *SLC6A1* mutation. Freed had been working as an equity analyst in Denver, Colorado, but quit the day of the diagnosis. "I stood up from my desk and never looked back," she says.

Until recently, many children with clusters of unusual symptoms would remain undiagnosed. Starting in 2010, genome sequencing became inexpensive enough to employ as a routine diagnostic tool. Now, more often than not, even mysterious inherited disorders can be linked to a genetic misspelling. "Now you can walk out of a hospital with a genetic cause," Freed told me. "I think pretty soon kids will walk out the door with a solution."

Without treatment, Freed's son will come to experience a violent form of seizure called a "drop attack." The victim remains conscious but frozen and can topple to the ground, unable to break the impact. "It's coming, but we are going to get the cure before it gets to that point," said Freed, who came to the meeting in a power suit and positioned herself near the stage. "We are going to find the cure for him. Our secondary mission is to help those who come after us."

That evening I spotted Freed perched on a stool at a Bethesda eatery, speaking to a researcher named Steven Gray. A soft-spoken southerner and gene-therapy specialist at the University of Texas Southwestern Medical Center, Gray has become the go-to scientist for parents like Freed. During the conference, he showed a slide listing 23 rare diseases for which he is trying to develop genetic treatments. Gray says he finds the kids' stories tragic and a powerful motivator.

Part of Gray's job is to reset parents' expectations. Gene therapy is not as simple as packaging a gene into a virus. Many diseases can be poor candidates—for instance, those in which a gene is overactive rather than broken. Often scientists have groundwork to do, such as engineering a mouse to mimic the condition. Bypassing these steps can be perilous. If a child's body has been missing a vital molecule since birth, for example, adding it may provoke a violent immune response. "We have gotten this wrong in the lab and we have killed mice," says Gray. "Gene therapy is not a pill you can stop taking."

Gray's best-known client is Lori Sames, whose daughter suffers from giant axonal neuropathy. The disease affects only about 80 people in the world. Sames managed to raise $6 million, which she funneled to Gray and into animal tests. In 2016, her daughter became the fifth child treated in a study of Gray's gene therapy at the NIH.

Gray told me that if a gene looks like a good candidate, and a family has money to support laboratory work, he will agree to take on their cause, no matter how rare the disease. "This is the most unfair system imaginable," he admitted. "If you don't have money, it won't happen."

To some bioethicists, when parents fund treatments it has the potential to create sharp ethical dilemmas. "There is a fairness issue if only the people who have the money get to be first in line," says Mildred Cho, a bioethicist at Stanford University, who has consulted on similar cases. "And there is a scientific integrity issue, because those with the money may not be the most appropriate or the best candidates. These decisions should be objective."

I asked Sames if she had created a conflict of interest by paying for research. The question makes "the hairs on my arms stand up," she said. "Anyone who suggests it's corrupt that parents privately fund development of a treatment for a child, in an attempt to save the child—well, I think it's irrational and rather insane. If the parents don't drive it, it's never going to happen. Wake up." Sames adds that the fund-raising she did never guaranteed her

daughter, Hannah, a spot in the NIH trial. Hannah had to pass a lung function test like others to get in. "We were no different than any other candidate," she says. "I wept the day she passed the test." Since then the trial has been moving forward at a "glacial" government pace, according to Sames, and other parents are mad at her. "They are hurt—their child is failing before their eyes—and they are angry, angry their kid is not injected," she says. "But there is nothing I can do."

Some families are managing to move even faster to a treatment than Sames did. The Hollywood couple, film producer Gordon Gray and his wife, Kristen, were able to get two daughters treated at Nationwide Children's Hospital about one year after the girls were found to have Batten Cln6, an inherited nervous-system disease believed to affect just a few hundred kids. Kristen Gray says the couple paid for the trial in its entirety. They also formed a company to take commercial rights to the treatment.

Few parents, though, are able to raise millions or start a company. On GoFundMe, hundreds of appeals mention gene therapy, but most raise only a few thousand dollars. One woman from Texas appealed for funds because she has muscular dystrophy; she has gathered just $35. The medical possibilities are out there, "but I don't think there is the regulatory infrastructure or the funding infrastructure to really make it happen," says Steven Gray, the gene therapist from Texas.

Another obstacle is that most of the key components of gene therapy are patented—including the viruses, the production tricks, and engineered genes. That means parents, and the scientists who help them, are often working in a cloud of legal uncertainty. Leone says to treat the Landsman boys she will have to buy $250,000 worth of trial insurance. "I could have been stopped with a phone call, but I wasn't. People have been very kind," she says. "But I will tell you, there are so many pieces in the patent puzzle … it's like a contemporary symphony, one that is atonal. It makes you want to scream."

Calling Bill Gates

Of the 7,000 rare diseases, around 90% currently have no treatment whatsoever. Gene therapy could potentially help with many, and in the future, new technologies like gene editing could, in theory, make it possible to fix nearly any genetic mutation. Christopher Austin, chief of an NIH branch responsible for new therapies, says eventually there may be as many different treatments as there are unique DNA flaws. To Austin, that means made-to-order, hyper-personalized medicine isn't some ethical mistake to avoid; it is the next step forward. "All of us need to think deeply that this is possible now," he says. It's something "that people have thought about for decades—and now it seems to be coming true."

Exactly who will pay to discover, develop, and deploy this Noah's ark of medicines is not clear. Lori Sames told me she sometimes fantasizes about approaching Bill Gates, whose foundation is trying to eradicate malaria and polio. Leone envisions a different solution: a global institute of cures, with access to manufacturing, hospital beds, and agreements in place to streamline the "biblical" work of dealing with insurers and regulators. "So any new disease, any new genetic mutation, we'd have everything set up," she says. "We would bring patients from all over the world for treatment."

Biotech companies have raced into gene therapy, but so far, much of their effort has been aimed at more common genetic conditions like hemophilia. The US government's clinical-trial website lists more than 20 open studies exploring gene treatments for that disease, which could be the technology's first blockbuster. The markets for ultra-rare diseases haven't drawn as much commercial interest. "Imagine a company with 75 employees that exists to treat 75 people. You can see the problem," says Eric David, the executive with BridgeBio.

In April 2018, however, something happened to make biotechnology executives take a fresh look. The Swiss drug company Novartis announced that it would buy the gene-therapy company AveXis for $8.7 billion. AveXis had just one drug in the clinic—it owned rights to the treatment for spinal muscular atrophy that had been tested at Nationwide Children's Hospital. The acquisition price was immense for a treatment used, at that time, on only about 15 kids, and for a disease that affects one in 10,000 births.

"My jaw hit the floor. I don't even know what $8 billion is," says Jerry Mendell, the doctor who led the trial. Mendell didn't hold shares in AveXis, but one of his center's former employees, Brian Kaspar, did. Kaspar, who joined the company, is now $400 million richer. "In my mind, the AveXis deal—there is a before and there is an after," says David. "After that, people who would not have looked at gene therapy for a disease quite that rare said, 'Wow—if I can get a trial going, maybe I can be worth a billion dollars too.'"

"All of us need to think deeply that this is possible now. It's something that people have thought about for decades—and now it seems to be coming true."

One reason AveXis was worth so much is that the treatment seemed to be an outright cure. That could let Novartis charge $2 million per patient, and perhaps more. To Walter Kowtoniuk, a principal at the investment company Third Rock Ventures,

in Boston, such medical successes mean diseases previously thought to affect too few people to attract companies are suddenly drawing intense interest. He says he has been "shocked" by the "massive competition" to gain control of gene-therapy programs.

That's created a situation in which desperate bids to treat children can rapidly turn profitable. In October, the Gray family—which had helped form a virtual company around the Batten Cln6 treatment—sold the rights to another biotechnology company, Amicus, for $100 million. Some investors are starting to think Canavan looks pretty interesting too. It's widely known among Canavan parents that a couple in Florida spent more than $1 million to get their child treated in a one-person study organized by the University of Massachusetts and the University of Florida. I reached the boy's father, who asked to remain anonymous but did say his son seems to have benefited.

The Florida experiment helped launch Canavan out of the too-rare-to-care category. Early in 2018, David's company, BridgeBio, entered an agreement to license the treatment from the University of Massachusetts and created a subsidiary, Aspa Therapeutics, which he now leads. But Kowtoniuk says other investors have been angling to take over the project because the risk seems much lower now that one boy has been treated. "There is a battle, literally a battle, to license the technology," he says. "I think there is such a tidal-wave shift in what is going on right now."

The growing biotech interest could mean the end of the parent-led scrambles. David told me he doesn't think the epoch of parent-financed gene therapy will last very long. "It's transitional," he says. "I think it's going to be for a limited time."

A ticking clock

David says the formal clinical trial of his company's Canavan treatment, different in design from Leone's, won't begin for another year, maybe two. For his company it's important to plan carefully and not rush, since that would jeopardize its investments and its aim of getting a treatment approved.

Jennie Landsman's children, though, can't wait that long. A self-financed experiment is probably the only way her two kids can get gene therapy in time. When I visited the Landsman home, I stood behind Josh, who was belted into a high chair, as his mother showed him pictures of lunch foods: chicken, macaroni and cheese, corn. Jennie followed his gaze.

She had hoped the boys would be treated by now. The team submitted its proposal to the FDA in June 2018, but the agency responded with a notice, called a "clinical hold," delaying the experiment. At the time of writing, in October 2018, Jennie was counting on December at the latest. Janson, the doctor running the trial, thought sometime in 2019 was more likely. He and Leone have plans to submit a new request following a meeting with FDA officials. He has also started testing the treatment on monkeys, a costly safety step he predicts regulators may insist on.

Nerves are fraying. At Benny's age, Canavan patients often have a steep decline in brain function. Even gene therapy might not reverse it. "My blood pressure is really going up," Janson says. "We probably lost at least three to four months."

When I visited her, Jennie had the idea of going to the FDA meeting and bringing her kids. She wanted to know what I thought. If the regulators saw them, how could they say no? Janson doesn't think it's a good idea. "I think we have to go within the system," he says. "We aren't a drug company. We don't have unlimited resources to lobby the FDA."

I asked Janson if he thought it was fair that the Landsmans' kids could end up getting treated while some other family without a surprise GoFundMe success would not be. "Unfortunately, there are a lot of things in society that are not fair," he said. "There are parents who want to see me in my neurology clinic and can't because they don't have insurance. We have a problem in society."

Precision medicine, it seems, is just another example: "There is no easy answer to your question, because the system is not set up to deal with this."

Antonio Regalado is MIT Technology Review's senior editor covering biomedicine.

TUMOR SHAPE

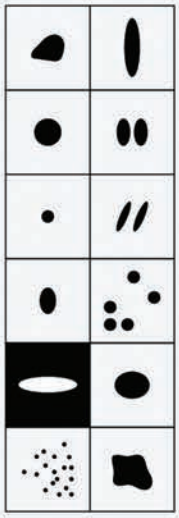

TYPE OF CANCER

LUNG CANCER

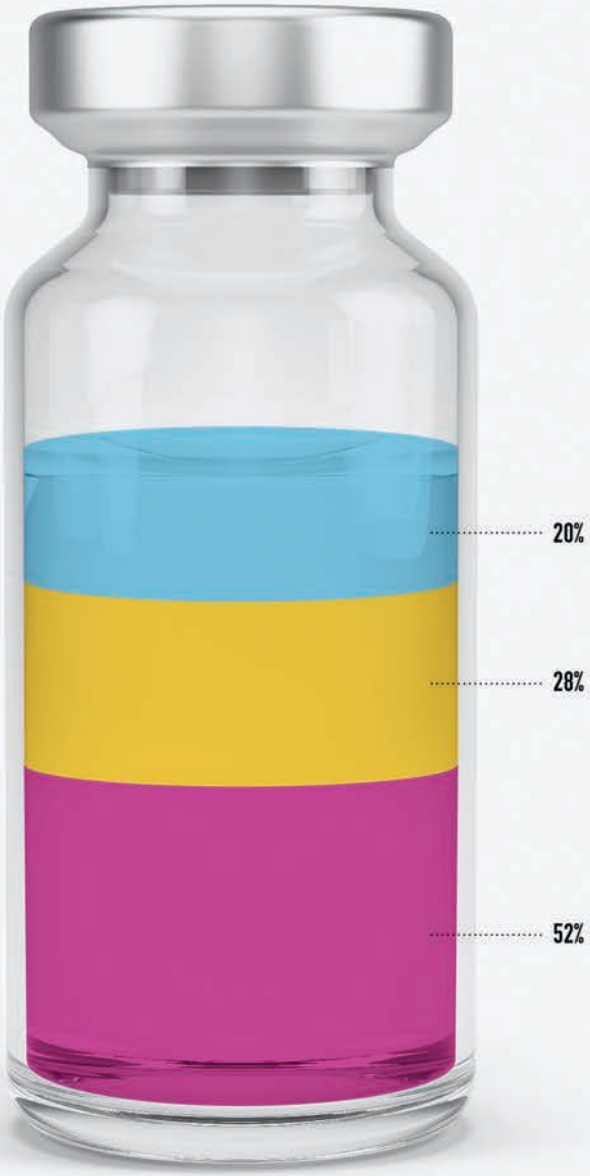

METHOTREXATE

VINCRISTINE

DOXORUBICIN

CYCLOPHOSPHAMIDE

CISPLATIN

5-FLUOROURACIL

BLEOMYCIN

DOSAGE

3

BLOOD TYPE

B+

AGE

54

GENDER

M

TREATMENT TIME

4–6 WEEKS

CANCER STAGE

1 2 3 4

One tumor at a time

Personalized cancer vaccines are a scientific breakthrough, but can they be a sustainable business?

by

Adam Piore

illustrations

Selman Design

The first time someone pitched Genentech's senior leadership on a personalized cancer vaccine, it did not go well. "I thought there was going to be a riot," Ira Mellman, then Genentech's head of research oncology, recalls.

From across the table, he watched the scientific review committee grimly shaking their heads as his team member and longtime collaborator Lélia Delamarre made her case. Then he overheard the head of clinical development turn to the person sitting next to him and mutter, "Over my dead body. A vaccine will never work."

That was in 2012. Cancer immunotherapy, which uses a person's own immune system to attack tumors, is now one of medicine's most promising fields, and one of the greatest breakthroughs in oncology in decades. But it took a long time to get there. Until the recent advent of a new class of blockbuster immunology drugs, the field was notorious for questionable science, hype, and spectacular disappointments.

And what Mellman and his team were proposing that day went further than turbocharging immune cells to make them better able to attack cancers. They were talking about a vaccine precisely tailored to stimulate the immune system to react to specific tumors. If it worked, the approach could, in some cases, be even more potent than other types of immunotherapy. But it faced a series of daunting hurdles. If Genentech, a San Francisco–based biotech company owned by the Swiss pharma giant Roche, were to attempt to develop a vaccine that could attack individual tumors, it wouldn't just have to accept new scientific advances; it would also have to embrace an entirely new and untested business model. That's because the vaccine Mellman and Delamarre envisioned could not be manufactured the traditional way, in large batches that could be packaged in bulk, warehoused, and dispensed off the shelf at your local pharmacy.

When Mellman and Delamarre said "personalized," they really meant it. The composition of each vaccine would be based on the characteristics of each patient's tumor DNA. The company would have to, in essence, make a separate treatment for every single patient.

Nor would this be the kind of drug you could order up with a prescription in hand and get in a few days, like Genentech's highly successful cancer drugs Herceptin and Avastin. To create this drug, the company would have to orchestrate a multistep process for each patient, performed at multiple sites. Each patient would need a biopsy, the tumor tissue would have to undergo full genome sequencing, the results would require complex computational analysis, and the individual vaccines would then need to be designed and queued up for manufacture. Theoretically, if the vaccines were to be produced on a large scale, this would have to happen hundreds of times a week. And it would have to happen fast.

If any single step in the process went awry, if a shipping mistake occurred or a batch was contaminated, it could prove deadly—because cancer doesn't wait.

No wonder the Genentech leadership was so skeptical.

After that calamitous first pitch meeting, Mellman and Delamarre retreated to their laboratories. They returned a few months later with more exciting data: they had identified specific targets on cancer cells, targets that would readily be attacked by immune cells. They also had fresh, convincing research from a growing number of other academic groups on the feasibility of their approach. And, critically, they had a preliminary plan for how Genentech itself might take the first tentative steps toward making tailor-made treatments an economically viable product.

This time the reception was different. The committee signed off on an exploration that would culminate in 2016 with a $310 million deal with BioNTech, a German company that has a technique for producing personalized vaccines to target tumors. Last December, the partners launched a massive round of human testing, targeting at least 10 cancers and enrolling upwards of 560 patients at sites around the globe.

At Genentech headquarters, Mellman and Delamarre's small team has grown by now into an army of hundreds, consisting not just of lonely lab workers but supply-chain specialists, regulatory experts, diagnosticians, and a whole host of consultants, all focused on the laborious task of figuring out how the production of their promising new product—should it continue to demonstrate the powerful effects seen so far—might be scaled up in a way that won't bankrupt the company.

"It's never been done, so we are learning as we go," says Sean Kelley, the project team leader overseeing the effort.

Nor are Genentech and BioNTech the only companies now pushing into this new territory. In late 2017, Moderna, a biotech based in Cambridge, Massachusetts, announced that, in partnership with pharmaceutical giant Merck, it intended to start human trials with a vaccine targeting solid tumors. Another company, Neon Therapeutics, founded by researchers at Dana Farber Cancer Institute and

The company would have to, in essence, make a separate treatment for every single patient.

Washington University, treated its first patient in phase 1 trials in May with a similar vaccine derived using a different method. It raised $100 million in an IPO this summer, driven largely by optimism over its approach.

The technology for the first truly personalized cancer vaccine is not yet proven. And these therapies are all likely to be expensive, Mellman acknowledged recently, sitting in a spacious conference room outside his office at Genentech's headquarters in South San Francisco. But he insists that if it's all done right, the extra costs and thinner margins will be more than offset by the sheer number of people who would use the treatment.

"You can imagine a scenario where every single cancer patient would benefit from this vaccine," he says. "That's unheard of."

Fighting against yourself

Scientists have been intrigued for decades by the possibility that cancer's greatest strength—its ability to mutate and evolve—might also be one of its greatest vulnerabilities.

Mutations in cellular DNA are, after all, what cause cancer in the first place, by prompting the cells carrying them to grow and proliferate uncontrollably. As far back as the 1940s, some researchers were arguing that it might be possible to put the immune system's cellular bloodhounds onto the scent of a specific tumor by somehow priming them with a vaccine that helped it recognize the tumor's mutations. A number of researchers have experimented and continue to experiment with techniques that involve removing immune cells from the body, genetically engineering them, and then reinfusing them in the hopes of triggering a robust response. Other cancer immunologists have focused on developing drugs to turn off molecular switches on the immune system's T cells that can interfere with their ability to attack.

But until recently, the scientific tools simply didn't exist to take the sophisticated personalized approach Genentech is now pursuing—an approach that requires scientists to fully characterize an individual cancer tumor, identify the most attackable mutations, and then design a personalized vaccine that would provoke the immune system to target them.

The problem was identifying the right target molecules on the tumor cell, or—as researchers thought of them—the antigens that would catch the attention of the immune cells. "It was so much work to identify antigens in the past," says Robert D. Schreiber, director of immunotherapy at Washington University. "You could do all this work, and then you end up with one antigen from one individual that is not necessarily ever seen again in any other individual."

That all changed with the advent of cheap genetic sequencing. In 2008, five years after the Human Genome Project published the sequence of the first human genome, scientists published the first genome sequence of a cancerous cell. Soon after, scientists began to compare the DNA in tumor cells and healthy cells to characterize the myriad ways that they differed. These studies confirmed that all cancer cells contain hundreds—if not thousands—of mutations, most of which are unique to each tumor.

In 2012, a team of German researchers, led by scientists at BioNTech, sequenced a widely used mouse tumor cell line designed to mimic human melanoma cells. They identified 962 mutations and used RNA sequencing to identify 563 that were expressed in genes. The group then created vaccines made of protein fragments that contained 50 of the mutations and injected them into mice to see if this would prime the immune system to respond. About one third—16 of the mutations—were detected

by the immune system, and five of those generated an immune response designed specifically to attack any cell found to harbor such mutations.

It was concrete evidence suggesting that genome sequencing could be used to design an effective cancer vaccine capable of putting the immune system on the trail of multiple mutations at the same time—and that such a vaccine might indeed provoke the immune system to attack a tumor. The race was on to answer the next logical questions: Why is it that the human immune system can be stimulated to attack some mutations and not others? And how can we figure out which mutations are most likely to be vulnerable?

At the urging of Mellman, Delamarre took Genentech's own lab mice and sequenced their tumor cells, identifying 1,200 individual mutations not present in normal tissue. Then she measured how T cells naturally responded to them. Of those 1,200 mutations, she found, the mice's immune system had begun to mount attacks against only two.

To answer why only those two mutations appeared to attract an immune response, Delamarre took a closer look at the interaction between the cancer DNA and a key component of the mouse immune system known as major histocompatibility complex, which in humans is called the human leukocyte antigen system (HLA). The HLA complex comprises 200 different proteins that protrude from cellular surfaces like microscopic thumbtacks on a poster board. When passing immune cells detect the presence of a protein fragment that doesn't belong—a piece of an unwanted virus or bacterium, or a mutation—they sound the alarm and cause the body to attack it.

Delamarre had determined that roughly seven of the 1,200 tumor mutations she'd identified were displayed on the cellular surface by HLA. When she examined the structure of these seven protein fragments, something got her attention: in the two that the immune system had recognized, the mutations were prominent on the cellular surface, facing up toward passing immune cells. Those the immune system had ignored faced down and were hidden in grooves in the cellular surface or obscured on the edges of the HLA. The immune system attacked those two mutations because they were the easiest to detect. By injecting mice with a vaccine designed to target those two mutations, she could enhance their bodies' ability to fight the tumors.

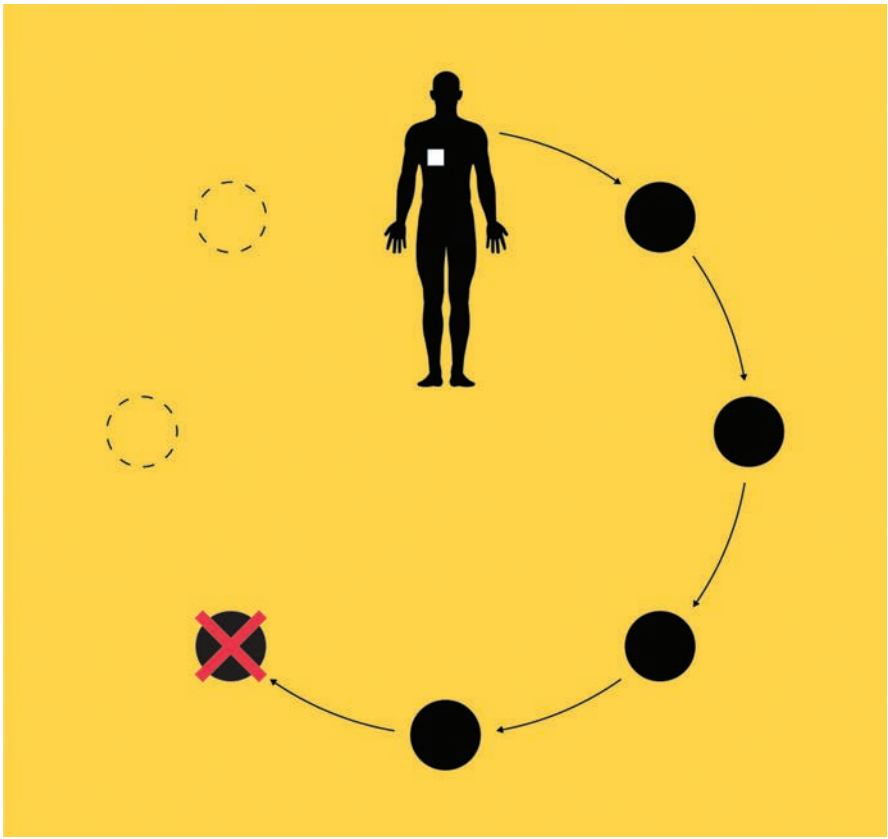

Together, these findings were what helped her and Mellman convince Genentech's review committee that a cancer vaccine was worth pursuing.

Facing the music

Genentech's headquarters, in an industrial park just off California's Highway 101, is a sprawling campus of glass buildings, hulking warehouses, and grassy courtyards. On a sunny morning this past August, cheerful groups of men and women in shirtsleeves and T-shirts strolled casually through a courtyard outside the company cafeteria. A band was setting up, getting ready to regale the lunchtime crowd with some blues, while nearby some kitchen workers prepared outdoor grills to cook food for employees.

Much of this is paid for by cancer drugs. Genentech won approval for its first cancer treatments in 1997, and since then the company has fielded no fewer than 15 of them.

But a cancer *vaccine* is unknown territory. The initial human trials that Genentech and BioNTech launched last year are shaping up as a test not just of the vaccine's efficacy but of the two

If any single step in the process went awry, if a shipping mistake occurred or a batch was contaminated, it could prove deadly.

partners' ability to scale up the new technology. By design, the geographic scope and the number of conditions targeted in the trial are broad—so far Genentech and BioNTech have opened sites in the US, the UK, Belgium, Canada, and Germany, and they are likely to expand to other nations around the globe.

Producing the vaccines even for the small number of patients in early trials "was an extremely challenging process," says BioNTech CEO Ugur Sahin, a veteran cancer researcher who cofounded the company in 2008. "Everything was driven by pipetting and by people on the bench producing the vaccines," he says. "So we had a very small capacity."

BioNTech has been able to automate some functions and reduce the time it takes to manufacture each vaccine from three months to about six weeks. It is shooting to get that down to four weeks by the end of the year.

The company can now produce hundreds of vaccines in a year—it aims to reach 1,500 over the next year. But if Genentech and BioNTech are ever to bring the product to market, they will need to be able to produce between 10,000 and 20,000 a year, Sahin says.

In San Francisco, teams from Genentech and BioNTech track progress in a designated space, consisting of a suite of rooms. On the walls, there are huge charts spelling out the patient status, the manufacturing and supply chain, the duration and schedule for each activity. "The key thing is that on paper it can look like a very coordinated process, but if any of those steps break down, then you can be in a situation where you have to start over," Genentech's Sean Kelley notes.

A number of unanticipated challenges have arisen. Early on, the team was surprised to discover that workers at BioNTech were contractually prohibited from working on weekends—so there was no one to receive patient tissue samples arriving then.

Gregg Fine, a senior medical director who is overseeing the trials, says he has been surprised by how variable the turnaround time has been at clinics and labs where patient biopsies themselves are collected and analyzed—a problem, since individual vaccines can't be manufactured until the samples are received.

The issue, Fine believes, is that patients with metastatic cancer may have problems getting to the doctor in a timely manner because they are too sick. And many collection sites don't yet have a procedure for flagging their samples as urgent, which means they can get lost in the stack with other biopsies.

Getting the vaccines back to the patients themselves has also proved problematic. At least one vaccine has been held up at customs in New York City.

For now, the problems are manageable and informative because the number of patients is relatively small. But all these problems will have to be solved if the vaccines are ever to go mainstream. "You're not going to be able to wait six months for a vaccine if you have a patient with fast-progressing pancreatic cancer," says Kelley.

Genentech officials declined to speculate about the eventual price of the vaccine, insisting it was too early to know. "It's going to be more expensive," says Kelley. "This will cost us much more to make per person."

The cost of sequencing might come down, building out a manufacturing network would increase efficiencies, and new assays might be developed, or new technologies that allow the cheaper manufacture of the vaccines themselves. "We've done estimates, and we feel that right now it is viable, but we would like it to become, obviously, more and more viable," he says.

For now, though, one of the most promising advances in cancer research remains an experimental treatment. It might be a medical breakthrough, but it is facing a familiar logistical challenge: how to get the product cheaply and quickly where it needs to go.

Adam Piore is the author of The Body Builders: Inside the Science of the Engineered Human, about how bioengineering is changing modern medicine.

The skeptic: What precision medicine revolution?

The benefits of genomic drugs are exaggerated, hurting patients and the practice of medicine, says one high-profile oncologist.

By Stephen S. Hall
Portraits by John Clark

Vinay Prasad is relatively young (35) and still climbing the academic ladder (he's an associate professor of medicine at Oregon Health & Sciences University in Portland), but he has already established an outsize reputation as a "professional scold" for his sharp critiques of contemporary biomedical research, including personalized medicine. In commentaries in high-profile medical and scientific journals, and in a Twitter account with some 25,000 followers, Prasad has questioned the evidence (or lack thereof) to support the use of precision oncology, the practice of selecting drugs for patients on the basis of specific mutations in their tumors. He has also criticized the inflated cost of cancer drugs and the financial conflicts of interests bedeviling contemporary research.

Prasad brings several unique perspectives to the role of medical scold. Born in Euclid, Ohio, outside Cleveland, to an immigrant couple from India, he developed an interest in philosophy in college before attending medical school at the University of Chicago. As a practicing oncologist, the prolific Prasad has generated a boatload of peer-reviewed papers, gathering evidence to suggest, among other things, that genomic-based evidence hasn't made much of an impact on cancer patients. As a sometimes prickly online persona, he has been faulted for unleashing expletive-laden putdowns but has also attracted a robust audience for what he calls "tweetorials," which dissect the design of high-profile studies and the data they generate. In the following conversation with veteran medical writer Stephen S. Hall, he takes aim at "precision oncology," the gaps in direct-to-consumer genetic testing, and what it really costs to bring a new drug to market.

Proponents have been promising a revolution in personalized medicine for decades. What's the reality?
I would say, and I think many people will agree, that the promises that were made around the time of the Human Genome Project have largely not materialized, and that the impact of personalized medicine has probably been exaggerated.

What's the danger of exaggerating the promises?
I think we have a schizophrenia in science and medicine. On the one hand, people who are good scientists understand that science is difficult. You should not be, nor will you be, having breakthroughs all the time. Breakthroughs are rare. Science is hard. It takes years of slogging to understand very fundamental pathways.

On the other hand, we often are tempted to—and I see experts continue to—make grandiose promises, and have a lofty, unrealistic vision for what might be achieved in the next few years.

That harms the public understanding of science, because the public comes to believe that unless you guys and gals are producing breakthroughs all the time, we shouldn't be funding this. That's wrong, because science needs more funding—needs a lot more funding than what we're currently investing.

Does it hurt the patient?
I would say inflated rhetoric about the value of medical practices, technologies, or science harms patients because it distorts their understanding of what a therapy or intervention might do. And by distorting the understanding, it robs them of autonomy. I'll give you just one example.

Sometimes cancer patients are on medications that add real side effects to their life, but they believe that there's going to be some survival benefit by taking this medicine. Every person is making kind of a daily decision: Do I stick with this medicine or not? Are the side effects worth it to me or not? And if that decision is made in a very impartial way, with a good understanding of what the drug does, that's the right way. But if that decision is made under the cloud of hype, when it's surrounded and marinated in hype and misinformation, then I think what we're really doing is that we're preventing the person from making the decision compatible with their wishes. We're kind of taking away that choice. And I do fear that that happens quite often.

You recently published a study indicating that most cancer patients don't benefit from personalized genomic

medicine, even though it's been in practice since at least 2006. Why do you think that's the case?
Some people have said that study is pessimistic. It's neither pessimistic nor optimistic; it is simply the most realistic estimate of how many people have benefited from genome-driven therapies. There clearly are some situations in cancer where drugging a single cancer-causing gene is important, and that should not be taken away. Those clearly do exist.

The problem is that they simply don't exist for the majority of patients who will be diagnosed with metastatic cancer. The purpose of our paper was to document what that number is, and what has been the change over time. I've heard the rhetoric that we're reaching exponential growth, or that [precision oncology] is taking off, or there's an inflection point. We simply don't see that evidence if you look objectively at the data.

Does that mean you're reluctant to use them in your own practice?
Of course I use genome therapies. I love [them]. Where they work, they work well. In fact, I would increase the funding to research them. But at the same time, I think we should be realistic about their prospects. We're also doing that same kind of analysis right now for immunotherapy drugs and cytotoxic drugs and different kinds of drugs. Can we more accurately compare what has been the impact of these different types of therapies?

In a recent article, you suggested that if adopted prematurely, the use of precision medicine might actually increase the risk of inappropriate medical care. How so?
Every day there are new potential treatments or therapies or strategies to treat any disease, and they all have some degree of bio-plausibility. When it comes to a new cancer drug, bio-plausibility is just not enough. You should also test it and prove that it does what you think it does. Precision medicine should be held to the same standard.

One of the differences is that precision medicine is very, very seductive. Some of its bio-plausibility is just such a compelling story that I think we do see this temptation by proponents that it shouldn't be assessed in the same way. It's so plausible, it should just be adopted—that kind of attitude. That kind of attitude might paradoxically lead us to adopt potentially more things that ultimately turn out not to do what you think they should do.

Do you think direct-to-consumer marketing by companies like 23andMe has made it seem as though personalized medicine has arrived already?
Yes, I think the constant rhetoric that this is wonderful has shifted the public perception. In terms of the direct-to-consumer advertising, we actually have a paper on the *BRCA* breast cancer gene test that appeared in [the *Journal of the American Medical Association*] about a month or two ago. It points out that there are some limitations to that direct-to-consumer *BRCA* testing. The test is actually only for three mutations that are very common in the Ashkenazi Jewish population, but not perhaps the most common *BRCA* mutations among all people with deleterious mutations. And thus there are some unintended consequences. A woman with a family history who may be worried will send off that test, get a negative result, and feel reassured. But that person may have a deleterious *BRCA* mutation. It may actually be counterproductive.

If genomic testing and these other aspects of personalized medicine are not currently predictive of outcomes for individual patients, are the drug companies and medical institutions taking advantage of consumers by pushing these methods?
It's a big category, and there are some things that are very well validated. But

I think there are some things that are not. And the consumer doesn't always know which ones are which, and that's the challenge. Even some of the people in the field apparently seem to forget which ones are which, and that's what I try to remind them of.

When you remind them, it sounds like you get pretty strong pushback.
I appreciate pushback when it's about the technical merits of any of these arguments. Where I think pushback is counterproductive is when pushback becomes personal or when pushback is about the intention.

There are a number of people who have voiced concern that one or more precision therapies don't have the data. And sometimes I feel as if the argument devolves into the people who want that therapy saying, "Well, we want what's best for patients. And you people who are saying that we don't have data, you apparently don't want what's best for patients." I think we have to recognize we all want what's best for patients. This is an argument about the evidence. And I get personally frustrated when I see people try to pervert the argument in that way.

You've also criticized the high cost of drugs, and you recently argued that industry estimates of the cost of bringing a new drug to market are wildly exaggerated. What does it really cost?
I think that the cleanest estimate that I've seen—and I'm a little bit personally biased—is the estimate that Sham Mailankody and I put out in *JAMA Internal Medicine*, where we estimate that it costs something like $800 million in R&D to bring a cancer drug to market. The industry estimate is $2.6 billion. There's a big difference there. But at the end of the day, this is one of those few things in life where you don't have to settle for estimates. Since the industry repeatedly uses the cost of R&D as a justification for the high price—and unsustainable price—of drugs, I think it's probably fair game for governments to ask them to show the data. Let's just put all the data on the table and let's see what it really costs.

One of the other things you've suggested is that the expert panels that advise the FDA have financial conflicts of interest. Is that compromising the quality of medicines that consumers are getting?
I just want to clarify my view here, which is that I wholeheartedly support collaboration between academic investigators and for-profit companies. The additional complexity and challenge is when you have payments made to physicians personally. I think those payments—and they've been shown to—do affect our perception of products. If you're receiving a lot of money from a manufacturer, you may not view their product as impartially as you would if you were not receiving that money. That's the concern. I think we should try to curb the financial conflicts of for-profit companies in the healthcare space.

There are some legitimate questions here about the role of financial conflicts in this space. Does it distort the impartiality around adjudicating medical practices? I fear it does.

Given the implications of the kind of critiques that you have been publishing pretty prolifically, why aren't more people saying the same thing?
I ask myself that all the time. These questions feel very obvious to me. There are a lot of people who do care. A lot of them are general internal-medicine folks. I think we see it a little less in the specialties. And I think we see it much more in the younger crop of physicians than the older crop, in the sense that people who have done this, practiced for many years in this environment and who have found their niche in the environment, they're comfortable where they are, and they don't really feel the urge to comment about these more problematic areas. But people who are younger, and approach this field with fresh eyes, feel as if these things are problematic.

You don't always sound like a scold.
I'm very optimistic about science, that we will improve outcomes. I just think that we would benefit from a lot more empiricism and impartiality in the process. That's what I feel is missing—empiricism, impartiality, and more modest rhetoric. I think those three things would go like 90% of the way.

Is it true, as reported by *The Cancer Letter*, that you've closed your Twitter account?
No, it's not true at all! I'm on Twitter, @VPplenarysesh. I believe that there are a number of inaccuracies in the *Cancer Letter* stories about me. I'll save that for another day.

Precision medicine is very, very seductive. The temptation is that it shouldn't be assessed in the same way as other treatments.

A psychiatrist in every pocket

Our obsession with smartphones may actually be a boon for treating disorders like depression and schizophrenia.

BY
RACHEL METZ

PHOTOGRAPHY BY
JESSICA CHOU

There are about 45 million people in the US alone with a mental illness, and those illnesses and their courses of treatment can vary tremendously. But there is something most of those people have in common: a smartphone.

A startup founded in Palo Alto, California, by a trio of doctors, including the former director of the US National Institute of Mental Health, is trying to prove that our obsession with the technology in our pockets can help treat some of today's most intractable medical problems: depression, schizophrenia, bipolar disorder, post-traumatic stress disorder, and substance abuse.

Mindstrong Health is using a smartphone app to collect measures of people's

Mindstrong founder
and CEO Paul Dagum

Cofounder Tom Insel, a psychiatrist and former director of the National Institute of Mental Health

cognition and emotional health as indicated by how they use their phones. Once a patient installs Mindstrong's app, it monitors things like the way the person types, taps, and scrolls while using other apps. This data is encrypted and analyzed remotely using machine learning, and the results are shared with the patient and the patient's medical provider.

The seemingly mundane minutiae of how you interact with your phone offers surprisingly important clues to your mental health, according to Mindstrong's research—revealing, for example, a relapse of depression. With details gleaned from the app, Mindstrong says, a patient's doctor or other care manager gets an alert when something may be amiss and can then check in with the patient by sending a message through the app (patients, too, can use it to message their care provider).

For years now, countless companies have offered everything from app-based therapy to games that help with mood and anxiety to efforts to track smartphone activities or voice and speech for signs of depression. But Mindstrong is different, because it's considering how users' physical interactions with the phones—not what they do, but how they do it—can point to signs of mental illness. That may lead to far more accurate ways to track these problems over time. If Mindstrong's method works, it could be the first that manages to turn the technology in your pocket into the key to helping patients with a wide range of chronic brain disorders—and may even lead to ways to diagnose them before they start.

Digital fingerprints

Before starting Mindstrong, Paul Dagum, its founder and CEO, paid for two Bay Area–based studies to figure out whether there might be a systemic measure of cognitive ability—or disability—hidden in how we use our phones. One hundred and fifty research subjects came into a clinic and underwent a standardized neurocognitive assessment that tested things like episodic memory (how you remember events) and executive function (mental skills that include the ability to control impulses, manage time, and focus on a task)—the kinds of high-order brain functions that are weakened in people with mental illnesses.

The assessment included neuropsychological tests that have been used for decades, like a so-called timed trail-tracing test, where you have to connect scattered letters and numbers in the proper order—a way to measure how well people can shift between tasks. People who have a brain disorder that weakens their attention may have a harder time with this.

Subjects went home with an app that measured the ways they touched their phone's display (swipes, taps, and keyboard typing), which Dagum hoped would be an unobtrusive way to log these same kinds of behavior on a smartphone. For the next year, it ran in the background, gathering data and sending it to a remote server. Then the subjects came back for another round of neurocognitive tests.

As it turns out, the behaviors the researchers measured can tell you a lot.

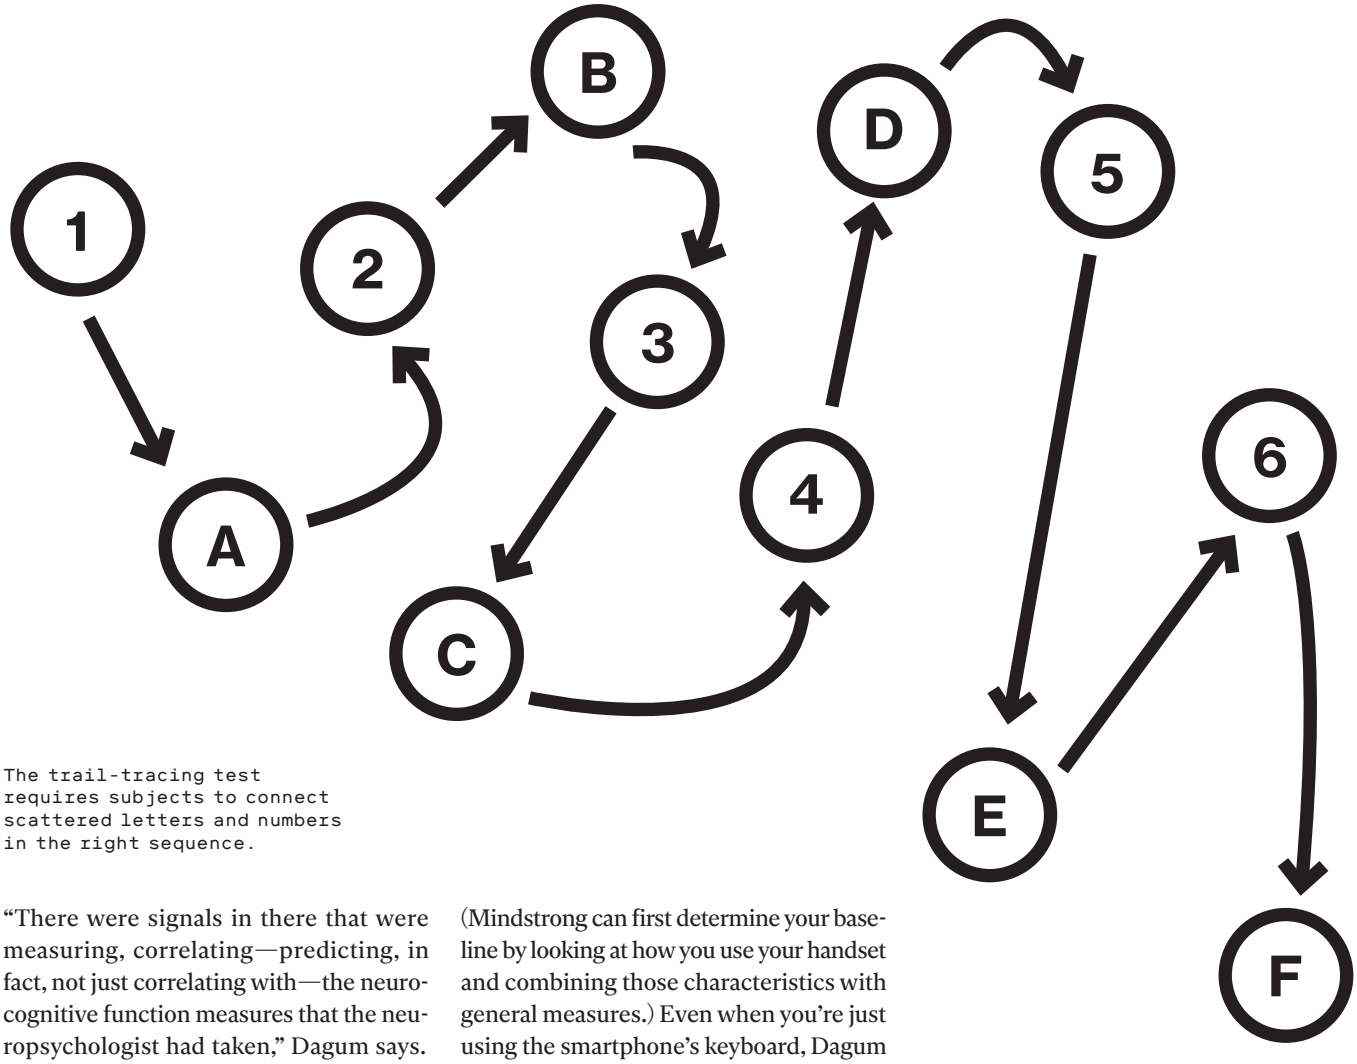

The trail-tracing test requires subjects to connect scattered letters and numbers in the right sequence.

"There were signals in there that were measuring, correlating—predicting, in fact, not just correlating with—the neurocognitive function measures that the neuropsychologist had taken," Dagum says.

For instance, memory problems, which are common hallmarks of brain disorders, can be spotted by looking at things including how rapidly you type and what errors you make (such as how frequently you delete characters), as well as by how fast you scroll down a list of contacts.

The assessment included classic neuropsychological tests that have been used for decades, like a so-called timed trail-tracing test.

(Mindstrong can first determine your baseline by looking at how you use your handset and combining those characteristics with general measures.) Even when you're just using the smartphone's keyboard, Dagum says, you're switching your attention from one task to another all the time—for example, when you're inserting punctuation into a sentence.

He became convinced the connections presented a new way to investigate human cognition and behavior over time, in a way that simply isn't possible with typical treatment like regularly visiting a therapist or getting a new medication, taking it for a month, and then checking back in with a doctor. Brain-disorder treatment has stalled in part because doctors simply don't know that someone's having trouble until it's well advanced; Dagum believes Mindstrong can figure it out much sooner and keep an eye on it 24 hours a day.

In 2016, Dagum visited Verily, Alphabet's life sciences company, where he pitched his work to a group including Tom Insel, a psychiatrist who had spent 13 years as director of the National Institute of Mental Health before he joined Verily in 2015.

Verily was trying to figure out how to use phones to learn about depression or other mental health conditions. But Insel says that at first, what Dagum presented—more a concept than a show of actual data—didn't seem like a big deal. "The bells didn't go off about what he had done," he says.

Over several meetings, however, Insel realized that Dagum could do something he believed nobody in the field of mental health had yet been able to accomplish. He had figured out smartphone signals that correlated strongly with a person's cognitive performance—the kind of thing usually possible only through those lengthy

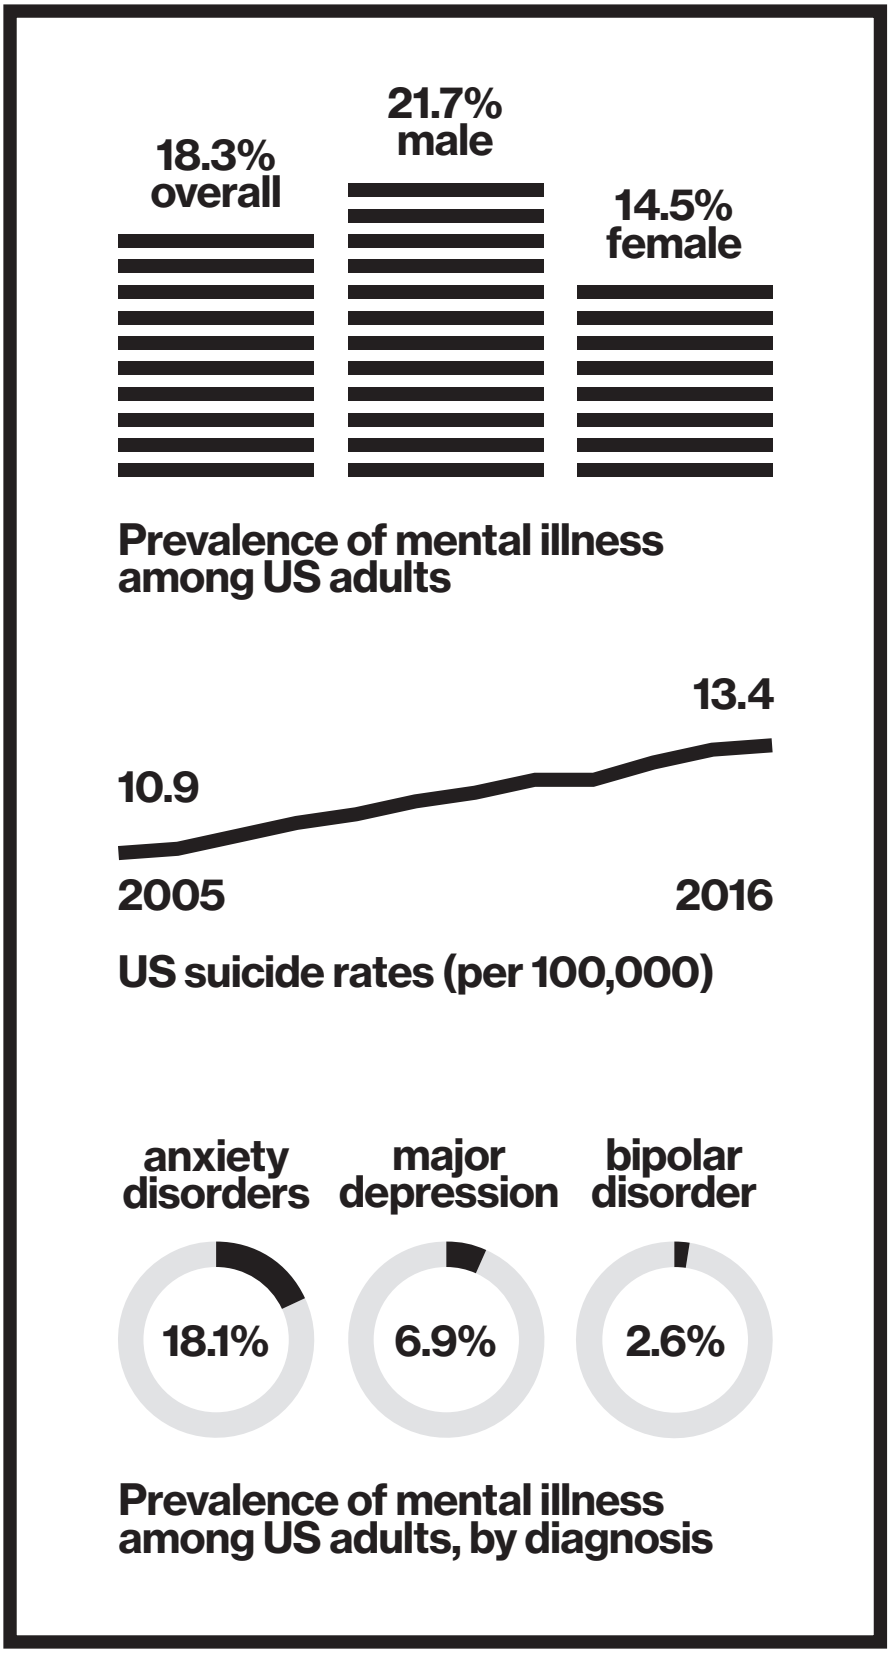

lab tests. What's more, he was collecting these signals for days, weeks, and months on end, making it possible, in essence, to look at a person's brain function continuously and objectively. "It's like having a continuous glucose monitor in the world of diabetes," Insel says.

Why should anyone believe that what Mindstrong is doing can actually work? Dagum says that thousands of people are using the app, and the company now has five years of clinical study data to confirm its science and technology. It is continuing to perform numerous studies, and this past March it began working with patients and doctors in clinics.

In its current form, the Mindstrong app that patients see is fairly sparse. There's a graph that updates daily with five different signals collected from your smartphone swipes and taps. Four of these signals are measures of cognition that are tightly tied to mood disorders (such as the ability to make goal-based decisions), and the other measures emotions. There's also an option to chat with a clinician.

For now, Insel says, the company is working mainly with seriously ill people who are at risk of relapse for problems like depression, schizophrenia, and substance abuse. "This is meant for the most severely disabled people, who are really needing some innovation," he says. "There are people who are high utilizers of health care and they're not getting the benefits, so we've got to figure out some way to get them something that works better." Actually predicting that a patient is headed

Sources: National Institute of Mental Health; National Alliance on Mental Illness

toward a downward spiral is a harder task, but Dagum believes that having more people using the app over time will help cement patterns in the data.

There are thorny issues to consider, of course. Privacy, for one: while Mindstrong says it protects users' data, collecting such data at all could be a scary prospect for many of the people it aims to help. Companies may be interested in, say, including it as part of an employee wellness plan, but most of us wouldn't want our employers anywhere near our mental health data, no matter how well protected it may be.

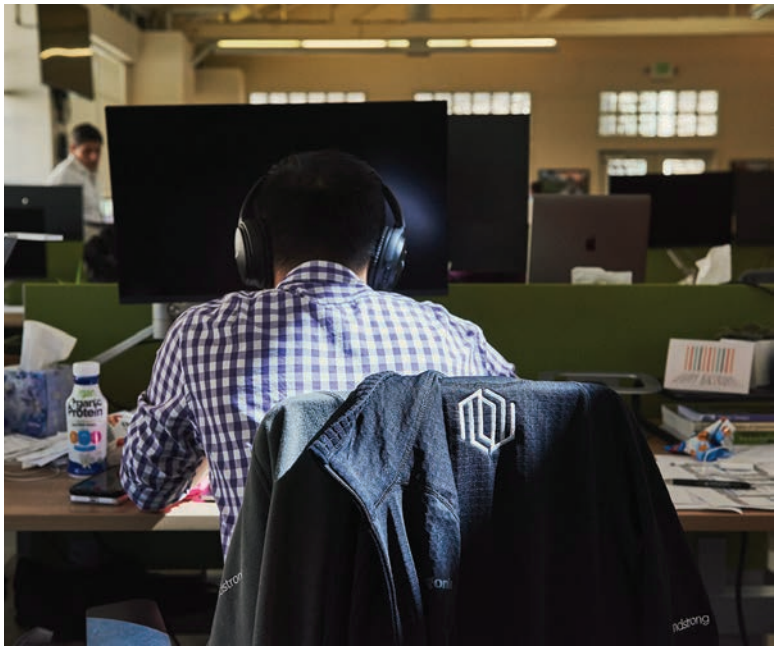

The Palo Alto startup wants assessing your mental health to fit into your regular life.

Spotting problems before they start

A study in the works at the University of Michigan is looking at whether Mindstrong may be beneficial for people who do not have a mental illness but do have a high risk for depression and suicide. Led by Srijan Sen, a professor of psychiatry and neuroscience, the study tracks the moods of first-year doctors across the country—a group that is known to experience intense stress, frequent sleep deprivation, and very high rates of depression.

We don't know how many different illnesses are in the category of depression. Insel hopes Mindstrong can use patient data to find out.

Participants log their mood each day and wear a Fitbit activity tracker to log sleep, activity, and heart-rate data. About 1,500 of the 2,000 participants also let a Mindstrong keyboard app run on their smartphones to collect data about the ways they type and figure out how their cognition changes throughout the year.

Sen hypothesizes that people's memory patterns and thinking speed change in subtle ways before they realize they're depressed. But he says he doesn't know how long that lag will be, or what cognitive patterns will be predictive of depression.

Insel also believes Mindstrong may lead to more precise diagnoses than today's often broadly defined mental health disorders. Right now, for instance, two people with a diagnosis of major depressive disorder might share just one of numerous symptoms: they could both feel depressed, but one might feel like sleeping all the time, while the other is hardly sleeping at all. We don't know how many different illnesses are in the category of depression, Insel says. But over time Mindstrong may be able to use patient data to find out. The company is exploring how learning more about these distinctions might make it possible to tailor drug prescriptions for more effective treatment.

Insel says it's not yet known if there are specific digital markers of, say, auditory hallucinations that someone with schizophrenia might experience, and the company is still working on how to predict future problems like post-traumatic stress disorder. But he is confident that the phone will be the key to figuring it out discreetly. "We want to be able to do this in a way that just fits into somebody's regular life," he says.

Rachel Metz is MIT Technology Review's senior editor covering the cyborg beat.

Profiles in precision medicine

Advances in DNA testing and gene editing have given people choices that would have been impossible a few decades ago. Here, in their own words, are the stories of four people confronted with these dilemmas.

When I was younger I watched my mother's decline, and eventual death, from Huntington's disease. It was a terrifying experience. So naturally when my wife and I learned that I had inherited the mutation that causes the disease, we were devastated.

We both understood the odds for any children we might have. As the child of a parent with Huntington's, you either have the mutation (and will surely develop the disease) or you don't, in which case your family is free of it forever. Since Huntington's is a dominant disease, meaning you need only one copy of the mutated gene to fall ill, 50% of the children of people with Huntington's face the same fate as their sick parent.

After receiving my test results, my wife and I decided we'd never have biological children.

Years later we heard about a new procedure called pre-implantation genetic diagnosis, whereby embryos generated for in vitro fertilization can be screened for Huntington's. By implanting only healthy embryos, we could reduce our child's risk of inheriting the disease from 50% to essentially zero. So we changed our minds.

Our city had an IVF clinic that had recently established a pre-implantation diagnosis program, but it hadn't yet had a successful pregnancy as a result of genetic screening. We would be its guinea pigs.

We got lucky. In 2006 my wife gave birth to healthy twins who did not inherit the mutation. Now that more than 12 years have passed, I can say that using this technology to have healthy kids was one of the most powerful ways I've had to fight back against my diagnosis.

We still don't have any disease-modifying treatments for Huntington's, although I'm hopeful there will be in time for me. But either way, pre-implantation genetic diagnosis gives parents affected by genetic diseases a cure for their kids, and for all future generations.

IAN ALLEN

Jeff Carroll

ASSOCIATE PROFESSOR,
Western Washington University
Bellingham, Washington

Lora Moser **CYSTIC FIBROSIS PATIENT, Austin, Texas**

I was born in 1977. At age two I was diagnosed with cystic fibrosis, a disease caused by a defective gene that changes a protein that regulates how cells process salt. The average life expectancy was 14 years. My parents, who had never heard of cystic fibrosis, were in panic. They devoted their lives to the Cystic Fibrosis Foundation and to date have raised more than $750,000 for drug research and development.

The disease didn't get to me as early as my parents feared. I began working at 16, stayed employed through college, and ultimately had to stop studying only because my need for medical insurance trumped a college degree. I ended up in retail management, and by the time I was 36, I'd been a store manager for two different multimillion-dollar companies.

But the long hours and inconsistent schedule took a toll on my health. When my pulmonologist noticed a significant drop in my lung function, I had to apply for disability. I had lived my life with the philosophy: "I have CF. CF doesn't have me." I was wrong.

The good news was that my parents' hard work paid off. We now have drugs that don't only treat the symptoms of cystic fibrosis but attack the disease itself at the molecular level. These drugs, tezacaftor and ivacaftor, work together to address a missing protein, known as CFTR, caused by a genetic mutation.

The bad news is that the cost of this treatment is beyond my family's means. The retail price of the drugs I need exceeds $38,000 a month. My husband has a private medical insurance policy provided by his employer. But because I'm on Medicare, I can't use his insurance, since Medicare recipients are prohibited from being on private medical policies. It's considered "double dipping."

I know we may be close to being able to edit the CF mutation out of embryos, or select embryos without the mutation. I support these advances, since they can end the suffering and early mortality caused by CF.

But those advances won't help me. I've been off the drugs I need for eight months. This landed me in the hospital on intravenous antibiotics for three weeks in August. I lost 26% of my lung function. Each day is a mental and physical battle with an unknown outcome.

Modern medicine gives us many gifts. But for many of us, those gifts are out of reach.

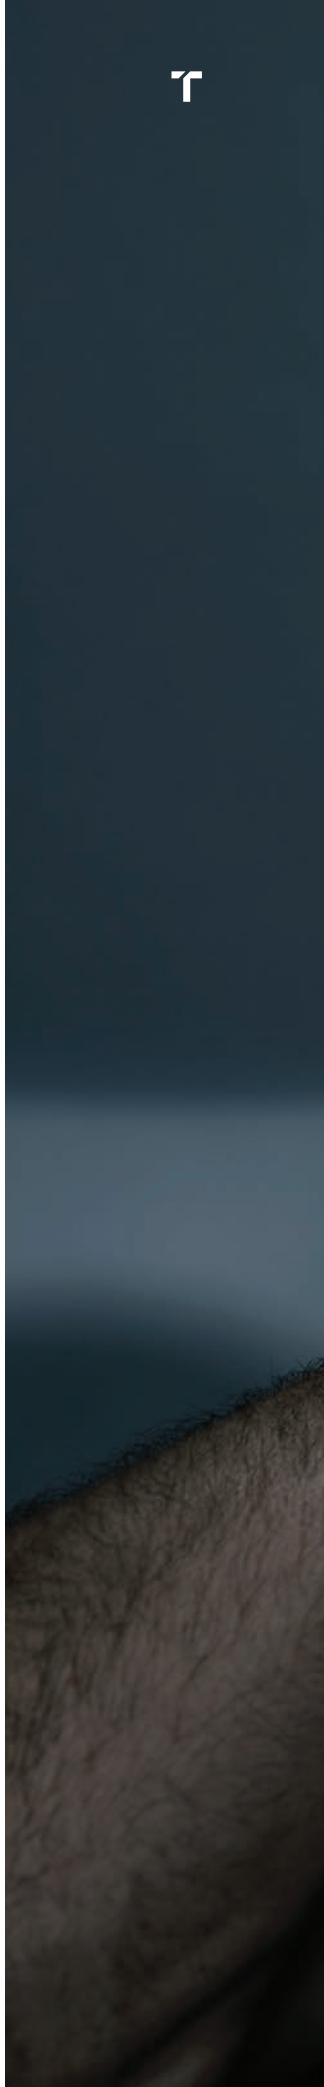

PHIL KLINE

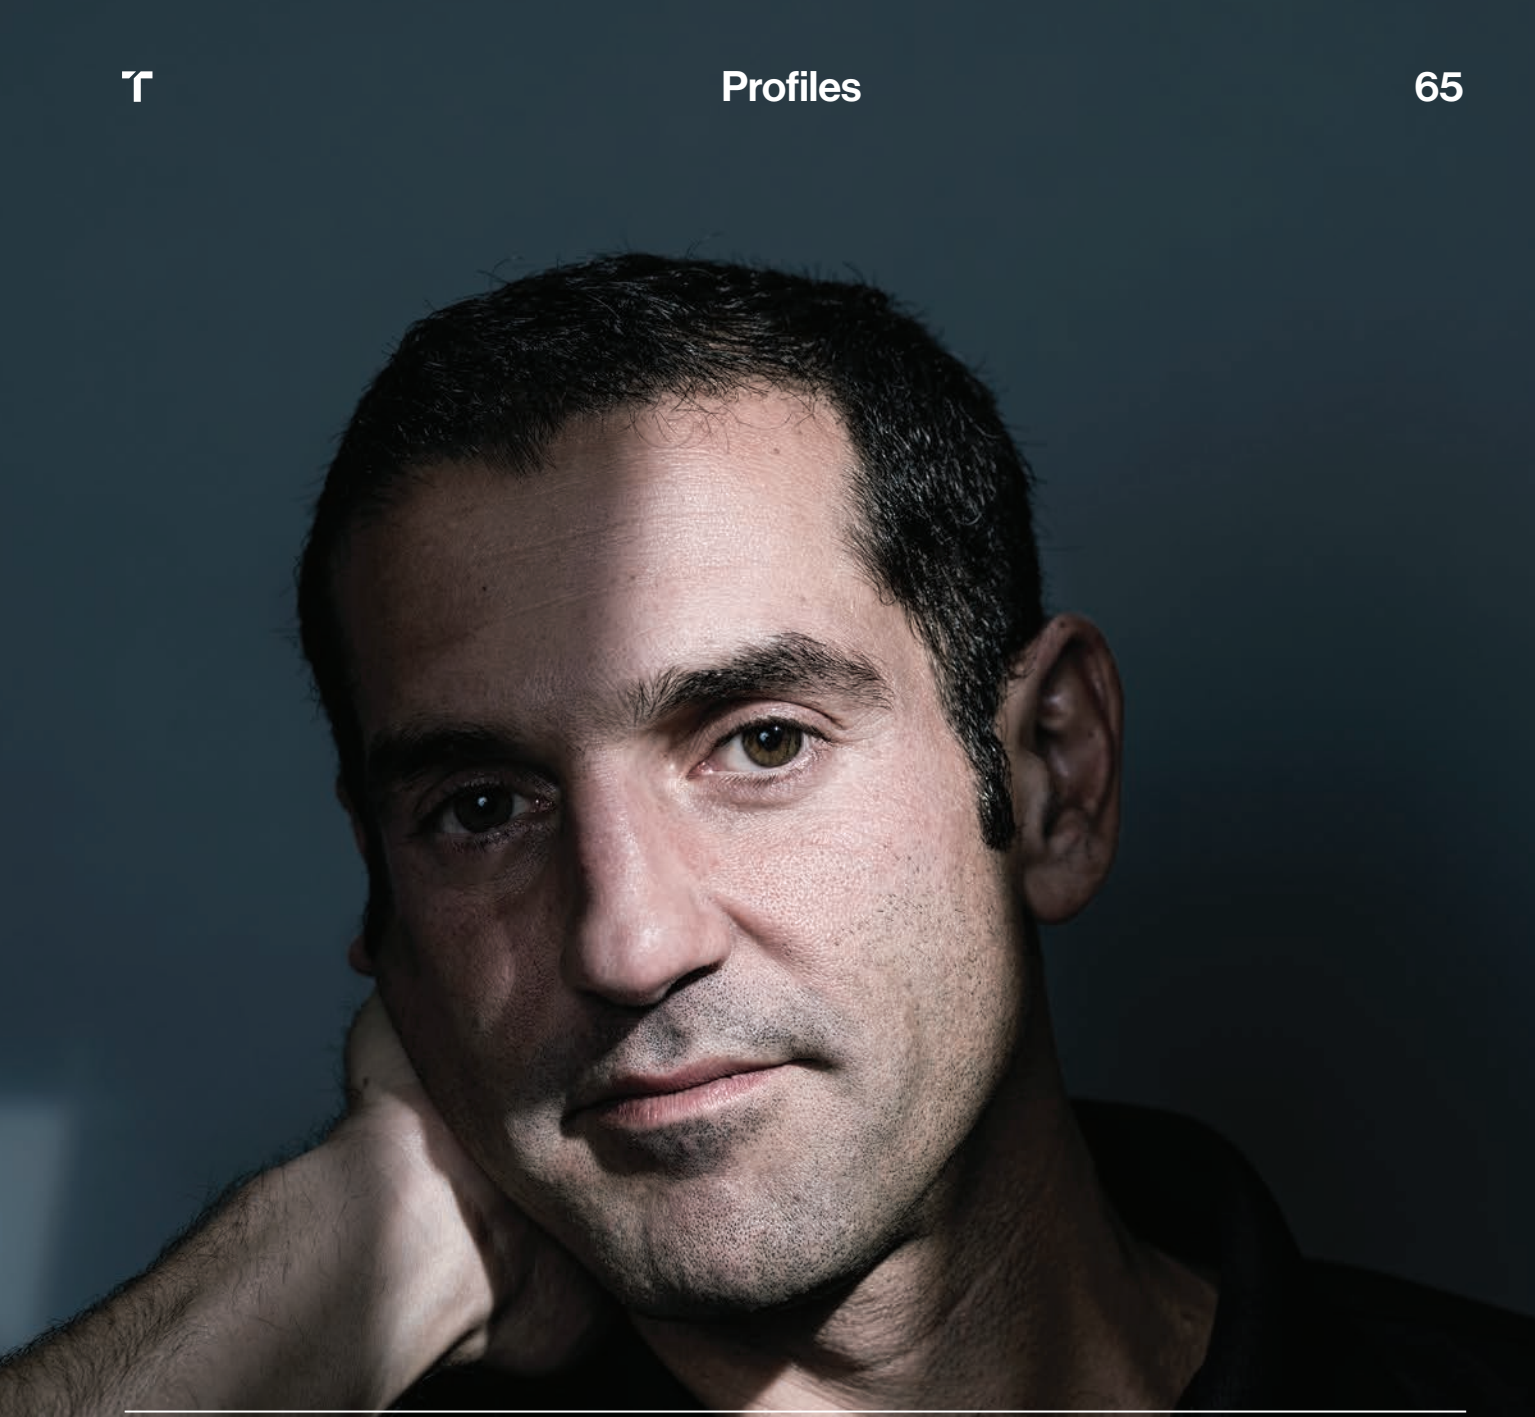

WINNI WINTERMEYER

Othman Laraki

CEO, COLOR GENOMICS
San Francisco, California

My grandmother died from breast cancer. My mother, who survived two breast cancers, got tested and discovered that she is a *BRCA2* mutation carrier, explaining our family's history. I later got tested and discovered that I, too, was a carrier of the same mutation: 1466DelT. One typo with consequences.

I learned of this almost 15 years ago. I'm now the CEO of a major genetics company.

When my wife, Elizabeth, and I started to discuss having children, I raised the issue of my mutation. We decided that we'd do pre-implantation genetic testing: that way we could select embryos based on whether they carried the mutated gene. I wanted to ensure that our children would be free of that cancer risk.

Elizabeth was just coming off birth control, so our doctor told us it was unlikely that she would get pregnant for a while, and suggested we return for the testing in a few months. Thinking nothing of it, we went back to life as usual. Almost immediately, Elizabeth got pregnant.

At first I was distressed that we were expecting a child without knowing whether he or she carried my mutation. But I came to terms with it and put the question out of my mind—after all, there wasn't much I could do about it.

Fast-forward a couple of years, and as we watched our first child, Kamal, take his first steps, Elizabeth said something striking: "I am so grateful for this child. This specific child. If we had gone through embryo selection, it would have been a different child, who would not have been Kamal. This child who we love and adore would not exist."

To me, that changed everything. I believe that this is a deeply personal choice—without a cosmically right or wrong answer. However, the thought that our choice would have caused our beautiful child to not exist convinced me that—at least for my *BRCA2*—we were willing to let fate call the shots.

Today, we have three wonderful boys: Kamal (five), Rami (four), and Zak (one). All of them may or may not have the *BRCA2* mutation, and we would have it no other way.

Paula Amato, MD ASSOCIATE PROFESSOR

Oregon Health & Science University, Portland, Oregon

Last year, I took part in an experiment that a lot of people think was ethically questionable. I was part of a team at Oregon Health & Science University that used CRISPR gene editing to correct a disease-causing gene mutation in human embryos. In other words, we were "editing" humans.

Why would we do this? With our work we were able to correct a mutation in a gene called *MYBPC3*. This mutation causes a deadly heart condition known as hypertrophic cardiomyopathy. Our work was potentially a first step toward eliminating the disease from that family and all its descendants.

Some people argue that we shouldn't pursue our research, and that instead women should simply undergo a pre-implantation genetic diagnosis, which could identify any embryos with the mutation before they're implanted. This sentiment is most likely uttered by people who have never treated an IVF patient.

One recent patient of mine and her husband easily conceived their first baby, who unfortunately was born with a disease called spinal muscular atrophy (SMA), a rare genetic neuromuscular disorder characterized by loss of motor neurons and progressive muscle wasting. The baby passed away at age one. The mother subsequently completed two physically burdensome rounds of IVF at a cost of tens of thousands of dollars. She made a total of four embryos, only one of which was chromosomally normal and unaffected with SMA. We transferred that embryo, but unfortunately it did not take.

Such cases are not at all unusual. The type of gene editing we're researching would complement pre-implantation diagnosis by reducing the number of cycles of IVF required. It would relieve patients of the associated physical burden and costs. And it would rescue the affected embryos.

Another patient of mine, who carried the *BRCA* gene mutation, which increases the risk of breast and ovarian cancer, came to see me for IVF and pre-implantation genetic testing to avoid passing on the gene to her children. She was conflicted—if her parents had made a similar choice, she wouldn't be here today. And she was right. They would have selected a different embryo. But what if instead, she (the very same person) could have been born, just without the *BRCA* mutation?

That's what our research promises. That's why we're doing it.

IAN ALLEN

Designer babies aren't futuristic. They're already here.

By **Laura Hercher**

Are we designing inequality into our genes?

At first, Matthew assumed the weakness in his knee was the sort of orthopedic nuisance that happens when you turn 30. It was weeks before he consulted a doctor, and months before it occurred to him that there could be a connection between his worsening limp and a cousin's shoulder problem when they were kids. DNA testing confirmed it: Matthew, like his cousin, had a genetic form of dystonia, a condition where muscles contract uncontrollably. Their grandfather most likely had dystonia as well.

I'd met Matthew only a few months earlier, when he'd married my friend's daughter, Olivia, in one of those hip old New York hotels with an elegant downtown vibe. Since I was the only genetic counselor of their acquaintance, they brought their questions to me. With their permission, I am sharing their story. I have changed their names to preserve their privacy.

Matthew was lucky. His was a mild version of DYT1 dystonia, and injections of Botox in his knee helped. But the genetic mutation can cause severe symptoms: contractures in joints or deformities in the spine. Many patients are put on psychoactive medications, and some require surgery for deep brain stimulation.

Laura Hercher is director of research at the Sarah Lawrence College Program in Human Genetics.

Their kids, Matthew and Olivia were told, might not be as lucky. They would have a 50–50 chance of inheriting the gene variant that causes dystonia and, if they did, a 30% chance of developing the disease. The risk of a severely affected child was fairly small, but not insignificant.

My friends learned there was an alternative. They could undergo in vitro fertilization and have their embryos genetically tested while still in a laboratory dish. Using a technology called pre-implantation genetic testing, they could pick the embryos that had not inherited the DYT1 mutation.

It would be expensive—costs for IVF in the US average over $20,000 for each try, and testing can add $10,000 or more. And it would require an unpleasant two-week process of ovarian stimulation and egg harvesting. "It wasn't the way I saw myself making a baby," Olivia told me. But they wanted what the procedure could offer them: a guarantee that dystonia was eliminated for the next generation, and beyond.

Matthew and Olivia don't think of themselves as having a "designer baby." That term has negative associations, suggesting something trivial, discretionary, or unethical. They weren't choosing eye color or trying to boost their kid's SAT score. They were looking out for the health and well-being of their future child, as parents should.

Public opinion on the use of assisted reproductive technology consistently draws a distinction between preventing disease and picking traits. The Johns Hopkins Genetics and Public Policy Center, which contacted over 6,000 people through surveys and focus groups from 2002 to 2004, summed up its findings this way: "In general, Americans approve of using reproductive genetic tests to prevent fatal childhood disease, but do not approve of using the same tests to identify or select for traits like intelligence or strength." The dystonia gene is in a gray zone—some people born with it live perfectly healthy lives—yet presumably few parents would criticize Matthew and Olivia's choice to weed it out.

All embryo testing does fit the "designer" label in one important way, however: it is not available to everybody.

BENEDIKT LUFT

We risk creating a society where some groups, because of culture or geography or poverty, bear a greater burden of genetic disease.

Matthew and Olivia opted in to what is a quiet but significant trend. Although the number of couples using this technology remains small, it is growing rapidly. According to the Society for Assisted Reproductive Technology, the number of US IVF attempts with single-gene testing rose from 1,941 in 2014 to 3,271 in 2016, an increase of almost 70%.

This is only the beginning. As the price of genetic testing of all kinds drops, more adults are learning about their genetic makeup as part of routine medical care and discovering specific genetic risks before pregnancy. But these people are still most likely to be affluent and educated, like Olivia and Matthew. While they consulted with IVF clinics, Olivia's own brother and his wife got news of a gene that increased risk for cancer in their kids. "If you could get rid of it, why wouldn't you?" he asked.

Cost was not a concern for these couples, but it is an obstacle for many Americans. The Centers for Disease Control and Prevention (CDC) estimates that 1.7% of babies born in the US today are conceived using IVF. It's much higher in countries that publicly fund assisted reproductive technology: 4% in Belgium, 5.9% in Denmark. A 2009 study found that 76% of the medical need for assisted reproduction in the US is unmet.

Insurance doesn't normally cover IVF in the US, except for a handful of states where coverage is mandated. Even policies that cover fertility treatment are inconsistent in what they reimburse. Coverage for pre-implantation genetic testing is downright Kafkaesque. Under many policies, testing the embryos is covered, but the IVF procedure itself is not, because the couples are not infertile.

"The analogy I like to use," says James Grifo, director of the Division of Reproductive Endocrinology and Infertility at NYU Langone Health, "is if you were having coronary bypass surgery and they didn't pay for cracking the chest."

At least part of the reason the IVF industry is growing is not that more people can afford it but that those who can are paying for new kinds of services. Egg banking, for example, is now aggressively marketed to younger women as an insurance policy against age-related infertility. In 2011, egg banking did not even exist as a category in the CDC's annual report on IVF; by 2016, storing eggs or embryos was the purpose of 25% of all IVF cycles. Elite companies like Facebook offer egg freezing as a perk, but for most people it remains a luxury.

Cost isn't the only barrier. Reproductive technology is less acceptable in racial, ethnic, and religious groups where being seen as infertile carries a stigma. Language barriers can reduce awareness and referrals. Geography also plays a role, since IVF clinics cluster in areas of greatest demand.

Presumably, many people would make the same decision as Matthew and Olivia if given the option, but many don't have that choice. Our discomfort around designer babies has always had to do with the fact that it makes the playing field less level—taking existing inequities and turning them into something inborn. If the use of pre-implantation testing grows and we don't address these disparities, we risk creating a society where some groups, because of culture or geography or poverty, bear a greater burden of genetic disease.

What could change society more profoundly than to take genetic disease—something that has always epitomized our shared humanity—and turn it into something that only happens to some people?

Foundational inventions that change entire industries.

At Qualcomm, inventing comes first. When we connected the phone to the internet, our foundational inventions created the mobile revolution. Now, as we lead the world to 5G, our inventions are going to enable new industries to be created, and the next great product the world can't live without.

qualcomm.com/weinvent

Qualcomm

Inventing the tech the world loves

What about death?

In the age of big data, death isn't an end, but more like a ... reformatting. Why not do some consulting work after you die? Or stick around to nag the grandkids? But beware—they might do things with your digital self that you don't like. Take an epigenetic test and put your date of demise on your calendar. Meanwhile, some new drugs might keep you young until then.

3

Never let me go

YOU'RE DEFINITELY TEMPORARY. BUT A DIGITALLY ENHANCED VERSION OF YOU DOESN'T HAVE TO BE.

Hossein Rahnama knows a CEO of a major financial company who wants to live on after he's dead, and Rahnama thinks he can help him do it.

Rahnama is creating a digital avatar for the CEO that they both hope could serve as a virtual "consultant" when the actual CEO is gone. Some future company executive deciding whether to accept an acquisition bid might pull out her cell phone, open a chat window, and pose the question to the late CEO. The digital avatar, created by an artificial-intelligence platform that analyzes personal data and correspondence, might detect that the CEO had a bad relationship with the acquiring company's execs. "I'm not a fan of that company's leadership," the avatar might say, and the screen would go red to indicate disapproval.

BY COURTNEY HUMPHRIES / PHOTOGRAPH BY TONY LUONG

Creepy? Maybe, but Rahnama believes we'll come to embrace the digital afterlife. An entrepreneur and researcher based at Ryerson University in Toronto, and a visiting faculty member at MIT's Media Lab, he's building an application called Augmented Eternity; it lets you create a digital persona that can interact with people on your behalf after you're dead.

While most older people haven't amassed enough digital detritus to build a working artificial intelligence, Rahnama posits that in the next few decades, as we continue to create our digital footprints, millennials will have generated enough data to make it feasible. Even as we speak, the digital remains of the dead accumulate. Something like 1.7 million Facebook users pass away each year. Some online accounts of the dead are deleted, while others linger in perpetual silence. "We are generating gigabytes of data on a daily basis," Rahnama says. "We now have a lot of data, we have a lot of processing power, we have a lot of storage capability." With enough data about how you communicate and interact with others, machine-learning algorithms can approximate your unique personality—or at least some part of it.

And what would the digital "you" look like? Well, what do you want it to look like? It might be a text-based chatbot like the CEO's or an audio voice like Siri or a digitally edited video or a 3-D animated character in a virtual-reality environment. It might be embedded in a humanoid robot.

TWENTY THOUSAND PERSONALITIES AT ONCE

We're not there quite yet. It's hard enough to create software agents that can carry on a natural-sounding conversation, let alone capture the personality of a specific person. There's no software that can interact, communicate, and make decisions the way you do. Rahnama says the CEO's avatar will be a "decision support tool," but it won't be capable of running the company.

"There is one thing that is missing in AI today, and that is context," he says. Most chatbots simply offer responses based on the content of a conversation, but our communication changes depending on who we're talking to, where we are, and what time of day it is. The need to include this kind of context was the basis for Rahnama's company, Flybits (for which he was named one of this publication's 35 Innovators Under 35 in 2012). Flybits provides a platform that lets companies tailor their communications to customers on the basis of contextual cues. A bank, for example, might offer different messages through its mobile app depending on your purchase history, your calendar schedule, or whether you're walking or taking a train.

The contextual part was something Rahnama found useful when he started Augmented Eternity. If you're going to construct a digital self, it's not enough to know that somebody said something. You have to know the context in which it was said—was the person joking? Annoyed? Reacting to today's news? These same kinds of clues end up being crucial when piecing together a digital personality, which is why the Augmented Eternity platform takes data from multiple sources—Facebook, Twitter, messaging apps, and others—and analyzes it for context, emotional content, and semantics.

A similar concept grabbed headlines a few years ago when Russian software developer Eugenia Kuyda created a chatbot representation of her best friend, Roman Mazurenko, who died in late 2015. Kuyda made the bot by plugging Mazurenko's personal messages with friends and family into a neural network built with Google's open-source machine-learning framework, TensorFlow. The bot was, by Kuyda's own admission, not very precise or polished, but when it answered questions, it often sounded uncannily like her friend.

Kuyda says the main complication with trying to create digital versions of the dead is that people are complicated. "We're extremely different when we talk to different people," she says. "We're basically like twenty thousand personalities at once." For example, Mazurenko had said things to her that he might have left out of a conversation with his parents. She could consult with his family and other friends to figure out which information was too sensitive to share. Could any company realistically do the same?

Rahnama obviously thinks so. He says Augmented Eternity will take a step toward accommodating various personalities by tailoring the conversation according to context and letting users control what data is accessible to whom. So someday his daughter might consult with his digital family persona, while a former student could ask questions of his academic persona. He sees it as one way of leaving a legacy—a way to keep contributing to society instead of fading to black.

IT'S NOT JUST FOR THE DEAD

But a digital avatar might also come in handy even when you're still around. AI could help transform your professional expertise from a scattered written record to a representation of your knowledge that people can interact with. A lawyer who charges hundreds of dollars an hour could let people consult a digital avatar instead,

for a much lower price. Celebrities, politicians, and other public figures could outsource some of their public interaction to digital versions of themselves. AI would allow us to consult experts with whom we'd never be able to meet in real life. The ability to represent and share expertise, Rahnama says, "can actually contribute to new business models on the internet." Rather than speaking with a generic Siri or Alexa, you could ask an eminent scientist, a politician, or a coworker. And why attend a business meeting when you could send your avatar?

Another startup, Eternime, based in Mountain View, California, offers to incorporate your personal information into "an intelligent avatar that looks like you" and that will "live forever and allow other people in the future to access your memories." Its founder, Marius Ursache, has been promoting the idea for years, and more than 40,000 people have signed up to Eternime's waiting list, but the self-funded company has still launched only limited beta versions. Ursache thinks the problem is less technical than behavioral: "People don't invest much time in activities that will pay off in decades," he says.

Whether or not it takes off as a business, Rahnama hopes Augmented Eternity will start conversations about privacy and data ownership. "The reason I like this research project is that it addresses a lot of key ethical questions around data science and AI," he says. "Like, who is going to own my information after I pass away?"

In a paper published in *Nature Human Behavior* earlier this year, ethicists Carl Öhman and Luciano Floridi from the Oxford Internet Institute argue that we need an ethical framework for the burgeoning digital afterlife industry. Should we treat digital remains by the same code that museums use for human remains? Doing so would severely limit the ways in which companies can use (or exploit) our data. If digital remains are like "the informational corpse of the deceased," they write, they "may not be used solely as a means to an end, such as profit, but regarded instead as an entity holding an inherent value."

Creepy? Maybe, but Rahnama believes we'll all come to embrace the digital afterlife.

HOLD A BLACK MIRROR UP TO NATURE

Just about every discussion of the digital afterlife, Öhman points out, mentions "Be Right Back," an episode of the British show *Black Mirror*, in which a bereaved young widow interacts with a digital avatar of her late husband. Over the course of the episode, she progresses from sending a few hesitant texts to a chatbot to purchasing a lifelike robot in her husband's image.

What's often overlooked in discussions about the show is the role of the company that created the avatar. In real life, Öhman says, we should be skeptical of such companies. The power of the digital dead to manipulate the living is enormous; who better to sell us a product than someone we've loved and lost? Thus our digital representations might be more talkative, pushy, and flattering than we are—and if that's what their makers think is best, who's going to stop them?

In the *Black Mirror* episode, the avatar periodically elicits more of the dead husband's data and upsells his widow on more expensive representations of him, until it becomes so lifelike that she can't "kill" it. The rhetoric around immortal digital selves focuses on our desire to be remembered. But wouldn't most of us want our loved ones to be able to let us go?

Courtney Humphries is a freelance writer who covers science and the environment for a variety of publications.

How to make sure your loved ones can get into all your accounts. Or, alternatively—how to cover your tracks.

By **Simson Garfinkel**

Six things to do with your data before you die

What would happen to your digital estate if you died, suddenly, before finishing this paragraph? Would your survivors be able to find what you left behind?

There is nothing hypothetical about this for many people: the problem emerges, wholly formed, when tragedy strikes. What's worse, more than half of Americans don't have a will, let alone one that's up to date, according to a 2016 Gallup Poll. As a result, most survivors lack a road map to the deceased's assets (physical and digital) or even, in some cases, the legal authority to proceed.

Fortunately, there are many things you can do now, without a lawyer, to make things easier for your survivors.

Simson Garfinkel is a science writer living in Arlington, Virginia, and coauthor of The Computer Book: From the Abacus to Artificial Intelligence, 250 Milestones in the History of Computer Science, published this November by Sterling Milestones.

#1 Build a back door

Fifteen years ago, if you died and your next of kin got your laptop, that person was pretty much guaranteed access to your data. Then, in 2003, Apple introduced full disk encryption, designed to protect your data from a thief, but also keeping it out of the reach of your survivors. Cryptocurrencies pose a similar problem: if no one has access to your digital wallet, then any value there is lost—there's no Bitcoin central control to complain to.

Today there's a debate as to whether tech companies should put back doors in their crypto technology so law enforcement can get access to data on devices they seize during an investigation. Short of that, it's easy to back-door your encryption yourself: just write down your hard drive's master password, put the paper in an envelope, and seal it. Do the same with your Bitcoin wallet. Make sure it's well hidden but in a location that's known to your loved ones.

#2 Sign up for Inactive Account Manager

If you have a Gmail account, use Inactive Account Manager to specify an e-mail address that will be automatically notified three months after your Google account goes inactive. Google defines "activity" broadly: if you check Gmail, log in to a Google website, or perform a search with a Chrome browser that's logged into your account, Google will assume you're not dead. But when your digital heartbeat stops, this approach ensures that someone you trust can access your Gmail account, Google Photos, and other data.

#3 Download your medical records

Your doctor is supposed to keep copies of your test results and other records, but it's a good idea to keep your own. Ask for copies and scan them. You might also be able to get your records directly if your health-care provider participates in the US government's Blue Button Connector, which lets you download PDF files for yourself and a special format for other health-care providers (should you wish to give it to them).

My elderly father keeps a copy of his records on a USB stick that he carries with him at all times. It comes in handy when he sees a specialist who might not have access to his primary care provider's computer. Yes, there's a risk the stick could fall into the wrong hands, but he's decided that the risk of medical professionals not having access to his records is greater.

#4 Use a password manager

It used to be straightforward to identify the deceased's accounts by waiting for bank

BENEDIKT LUFT

If you're an avid Facebook or Twitter user, take some time to read their data-after-death policies. You might not like what you find.

statements and tax bills to arrive by snail mail. These days, two thirds of Americans do their banking online (according to a 2017 survey by the American Bankers Association), and many people no longer receive paper statements. This significantly increases the chance that your bank accounts or retirement accounts might be declared "abandoned" in the event that you die.

So use a password manager like 1Password or LastPass. Now make sure that your spouse, or lawyer, or children, or parents, or *somebody* has some way to get to your accounts (so they can, for example, save any cherished photos or easily delete your accounts after you're gone).

One way that couples can simply access each other's accounts is by sharing their passwords. This is getting harder as websites implement two-factor authentication, but it's still possible by registering multiple second factors (like a FIDO Universal 2nd Factor device) and giving one to each partner.

#5 Ponder the complexities of social media

If you are an avid user of Facebook or Twitter, take some time to read their data-after-death policies. You might not like what you find.

When Facebook is notified that one of its users has become medically incapacitated or died, the company allows authorized individuals to request that the user's account be either "memorialized" or removed. Be aware: memorialized accounts can be managed by a legacy contact (who has to be specified in advance), but that person can't log into the Facebook account, remove or change past posts, or read private messages. In one famous case, parents of a 15-year-old German girl who died after being hit by a subway train were unsuccessful in trying to force Facebook to open the girl's account so that they, the parents, could determine if she had experienced cyber-bullying or depression, or if her death really was a tragic accident.

Twitter's policy is similar: after you die, a family member can contact the company and ask that your account be deleted, according to a help page on its website. Twitter will also, if requested, remove specific imagery or messages sent just before or after an individual's death. But Twitter will not give family members access to a deceased user's private messages.

So if you're storing something on Facebook that you'd like people to have access to after you're gone, you should download that data regularly and store it where your loved ones will have access—for example, in Google Drive.

#6 Be careful what you wish for

I gave much of this advice at a cybersecurity training seminar a few months ago, and almost everybody in the room thought I was crazy. The people there—mostly men—said they'd never share their passwords with their spouses.

And maybe they've got a point. Family members should be careful about taking extraordinary measures to crack open these encrypted digital crypts, warns Ibrahim Baggili, associate professor of computer science at the University of New Haven and an expert in digital forensics. "This person I knew died, and his wife managed to finally break into his e-mails and iPad and found all sorts of things about him that she did not want to know," says Baggili. "She really loved him, and it changed her whole perspective on him."


MIT Technology Review
Technology is threatening our democracy. How do we save it?
The politics issue
Vol. 121 No. 5
Sep/Oct 2018
$9.99 USD
$10.99 CAD
The "neuropolitics" consultants who hack voters' brains
MIT Technology Review
The Download
Taking control back from hackers
Sustainable Energy
To Feed the World, Improve Photosynthesis
By reworking the basic metabolism of crops, plant scientists hope to forestall devastating food shortages.
by Katherine Bourzac
Rewriting Life
Inside the effort to print lungs and breathe life into them with stem cells
Martine Rothblatt wants to end transplant shortages with 3-D-printed lungs.
2018
35 Innovators Under 35
Meet Jonas Cleveland, CEO of COSY, and the other young inventors, pioneers, and entrepreneurs shaping the future for all of us.

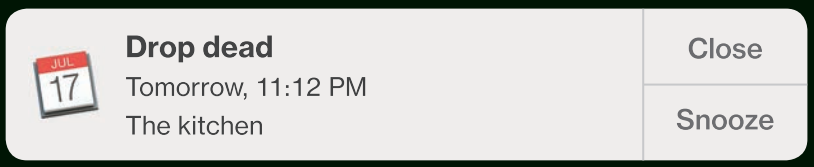
Drop dead
Tomorrow, 11:12 PM
The kitchen
Close
Snooze

Your life span is written in your DNA, and we're learning to read the code.

Want to know when you're going to die?

By Karen Weintraub

It's the ultimate unanswerable question we all face: When will I die? If we knew, would we live differently? So far, science has been no more accurate at predicting life span than a $10 fortune teller. But that's starting to change.

The measures being developed will never get good enough to forecast an exact date or time of death, but insurance companies are already finding them useful, as are hospitals and palliative care teams. "I would love to know when I'm going to die," says Brian Chen, a researcher who is chief science officer for Life Epigenetics, a company that services the insurance industry. "That would influence how I approach life."

The work still needs to be made more practical, and companies have to figure out the best uses for the data. Ethicists, meanwhile, worry about how people will cope with knowing the final secret of life. But like it or not, the death predictor is coming.

The clock

Steve Horvath, a UCLA biostatistician who grew up in Frankfurt, Germany, describes himself as "very straight," while his identical twin brother is gay. So he had a personal interest when, a few years ago, a colleague asked him for help analyzing biological data from the saliva of twins with opposite sexual orientations. The colleague was trying to detect chemical changes that would indicate whether certain genes were turned on or off.

The hypothesis was that these so-called epigenetic changes, which alter the activity of DNA but not the DNA sequence itself, might help explain why two people with identical genes differ in this way. But Horvath found "zero signal" in the epigenetics of the twins' saliva. Instead, what caught his attention was a powerful link between epigenetic changes and

aging. "I was blown away by how strong the signal was," he says. "I dropped most other projects in my lab and said: 'This is the future.'"

Horvath became particularly intrigued by how certain chemical changes to cytosine—one of the four DNA bases, or "letters" of the genetic code—make genes more or less active. Given someone's actual age, looking for these changes in that person's DNA can tell him whether the person's body is aging unusually fast or slowly. His team tested this epigenetic clock on 13,000 blood samples collected decades ago, from people whose subsequent date of death was known. The results revealed that the clock can be used to predict mortality.

Because most common diseases—cancer, heart disease, Alzheimer's—are diseases of aging, the ticking of Horvath's clock predicts how long someone will live and how much of that life will be free of these diseases (though it doesn't foretell which ones people will get). "After five years of research, there is nobody who disputes that epigenetics predicts life span," he says.

Aging eight or more years faster than your calendar age equates to twice the typical risk of dying, while aging seven years slower is associated with half the risk of death, Horvath says. His lab has developed a new version that is such a precise life span predictor they named it after the Grim Reaper: DNAm GrimAge. The epigenetic clock is more accurate the younger a person is. It's especially inaccurate for the very old.

"At this point, we don't have any evidence that it's clinically useful, because there are big error bars," Horvath says. Besides, there's no pill to reverse the effects. But though it will never be perfectly accurate, Horvath and his clock are getting closer than anyone else ever has to answering the question that hangs over us all—and determining whether there is anything we can do to change the answer.

Slow the ticking

As we age, the cytosine at hundreds of thousand of spots in our DNA either gains or loses methyl chemical groups (CH_3). Horvath's insight was to measure these increases and decreases in methylation, find the 300 to 500 changes that matter most, and use those to make his clocks. His findings suggest that the speed of the clock is strongly influenced by underlying genes. He estimates that about 40% of the ticking rate is determined by genetic inheritance, and the rest by lifestyle and luck.

Morgan Levine, who completed postdoctoral research in Horvath's lab and now runs her own lab at Yale, is starting to compare an individual's epigenetic profile with the profile of cells from the lining of a healthy umbilical cord. The more people deviate from that standard, the worse off they are likely to be. She thinks she will eventually be able to compare various epigenetic age measures to predict even in childhood who is going to be at greatest risk of which diseases—when it's still early enough to change that future. "Your genes aren't your fate, but even less so with things like epigenetics," she says. "There definitely should be things we can do to delay aging if we can just figure out what they are."

A few likely contenders are totally unsurprising. Eating a healthy diet including lots of vegetables and fish is associated with slower epigenetic aging. Feel older when you're sleep deprived? It's probably not a coincidence. Horvath has shown that people with insomnia are more likely to show accelerated epigenetic aging. "Everything you associate with a healthy lifestyle does relate to the new biomarkers in the expected way, which is a boring result, but it's scientifically very exciting," he says.

More unexpectedly, he finds that regular exercise won't add much more than a few months to your life. But those measurements are only on the DNA in blood, and Horvath says he'd like to look at changes in muscle, too, to see whether exercise makes a bigger difference there.

Horvath's own clock is not inspiring. He was surprised in analyzing his urine to find that he was epigenetically tracking five years older than his chronological age. A few years later, he tested his blood and was relieved to find the results more in line with his years, but still, he says, "I would say I'm not blessed in terms of epigenetic aging."

At age 50, he says his work is motivated by self-interest—"I'm as desperate as anyone else to find ways of slowing aging." But he also keeps in mind the social and financial costs of an aging population. "We need to find ways to keep people healthier longer," he says.

"After five years of research, there is nobody who disputes

He hopes that refinements to his clock will soon make it precise enough to reflect changes in lifestyle and behavior. Investors and biotech companies are spending hundreds of millions of dollars right now on drugs that might slow aging and defer disease. But how will we know what's effective? Those working on drug discovery can't wait 50 years to find out. Horvath hopes his clock will provide the answers.

The business of death prediction

Companies like Reinsurance Group of America are already looking into using the epigenetic clock to tweak and personalize risk assessments for life insurance. Right now, rates are based largely on demographics—people's gender and age—and a few health metrics, such as whether they smoke. The clock adds another useful data point.

Such personalization raises questions about fairness. If your epigenetic clock is ticking faster through no fault of your own, should you be charged a higher rate for life insurance? The Genetic Information Nondiscrimination Act of 2008—known as GINA—protects against discrimination

DAMON CASAREZ

Horvath, 50, says his work is motivated by self-interest. "I'm as desperate as anyone else to find ways of slowing aging."

on the basis of genes. But it doesn't address epigenetics.

There's also the issue of privacy. Your likely life span or true biological age is information that many consider intensely personal. For now, regulations and privacy policies don't even consider the possibility of such information. But as the science quickly progresses, questions about how to use and protect this data will become ever more pressing.

Can Horvath's clock and other technologies being developed to predict death ever be accurate enough to be truly useful? "I haven't seen any of these purported predictive algorithms be precise in terms of timing of death—to the contrary," says Diane Meier, a professor of geriatrics and palliative medicine at the Icahn School of Medicine at Mount Sinai in New York City. "People live for a really long time with a very high burden of disease and frailty," she says.

Gal Salomon, CEO of Clew Medical, an Israeli company that uses artificial intelligence to identify medical risks in hospitals, says he initially resisted the idea of developing a death predictor, thinking it unethical. Then he realized that doctors could use the technology "to understand where we need to stop." An algorithm Clew developed can help doctors and family members make the decision to switch from aggressive to palliative care, he says, overruling the typical instinct to provide heroic live-saving measures. The system, which for the moment is used only in hospitals, can also alert a family that the end is near, he says.

Atul Butte, a professor at the University of California, San Francisco, who studies quality of care, says the jury is still out about whether this kind of machine learning from patterns of care actually provides better treatment. But there's no doubt, he adds, that medical care is headed in that direction. "Five to 10 years from now, the health system that doesn't use this data to improve their medical delivery is going to be deemed archaic," he says.

Karen Weintraub is a freelance writer based in Cambridge, Massachusetts.

Finally, the drug that keeps you young

Anti-aging pioneer Judith Campisi explains how a recent breakthrough could ward off age-related disease.

By Stephen S. Hall
Portraits by Christie Hemm Klok

Judith Campisi has been a leading figure in the biology of aging since the early 1990s, when her research on the basic mechanisms of cancer revealed an unexpected finding—that cells enter a phase known as senescence that prevents them from becoming cancerous. More than 25 years later, the insight has led to a new kind of drug that may slow or modestly reverse human aging.

Campisi's research is on the role of cellular senescence in cancer and other age-related diseases. Senescent cells undergo a transition into a twilight state where they are still active but no longer dividing; research by Campisi and others showed that this was a strategy to derail incipient cancers, which are characterized by runaway cell division and growth. But she and others also discovered that these senescent cells accumulate as we grow older, secreting an array of molecules that promote the tissue degradation associated with aging.

In the past five years, this insight has led to the pursuit of a new class of drugs known as senolytics, which eliminate senescent cells and, in animal experiments, restore more youthful characteristics. Campisi, a professor at the Buck Institute for Research on Aging in Novato, California, cofounded a company called Unity Biotechnology in 2011, which launched a human trial of its first senolytic drug last July.

She recently discussed her work with Stephen S. Hall, a journalist who has been following anti-aging work for more than two decades.

Why should we suddenly get excited about anti-aging drugs again?
There are now tools available to biomedical scientists that simply didn't exist when I was a graduate student or even a postdoc. So we're finally able to do experiments that were either considered impossible in some cases or were just dreams 20 or 25 years ago. The other thing that has changed is that the field of senescence—and the recognition that senescent cells can be such drivers of aging—has finally gained acceptance. Whether those drugs will work in people is still an open question. But the first human trials are under way right now.

How specifically does senescence contribute to aging?
The correct way to think about senescence is that it's an evolutionary balancing act. It was selected for the good purpose of preventing cancer—if [cells] don't divide, [they] can't form a tumor. It also optimizes tissue repair. But the downside is if these cells persist, which happens during aging, they can now become deleterious. Evolution doesn't care what happens to you after you've had your babies, so after around age 50, there are no mechanisms that can effectively eliminate these cells in old age. They tend to accumulate. So the idea became popular to think about eliminating them, and seeing if we can restore tissues to a more youthful state.

You've suggested that health care could be transformed by senolytic drugs, which eliminate senescent cells. That's a pretty broad claim.
If we think of aging as a driver for multiple age-related pathologies, the idea would be that a new generation of physicians—we call them geriatricians today—will take a much more holistic approach, and the interventions will also be more holistic. That's the idea—it would revolutionize the way we're thinking about medicine nowadays. And just to remind you, 80% of patients in the hospital receiving acute medical attention are over the age of 65. So the idea is that senolytics would be one weapon that geriatricians will have in their arsenal of weapons to treat aging holistically as opposed to one disease at a time.

There is a debate about whether there's a biological limit to the human life span, about 115 years, or whether maximum life span could be extended as long as 130, possibly 150 years. What do you think?
At present, we simply don't know enough to know whether it will even be possible to extend maximum human life span. Average life span? No problem—it's already been done. But maximum life span? We just don't know.

If you look at *C. elegans*, a little worm, the world record for extending the life span of that animal is 10-fold. For humans that would be unbelievable, right? A thousand years. But if you go up the evolutionary scale just a little bit, to the fruit fly *Drosophila*, it's maybe

twofold. And then if you go to a mouse, most of the really high-profile papers extend its life span maybe 20%, sometimes 30%. So think about the difference between a mouse and a human. We're something like 97% genetically identical, meaning we have the same genes. And yet there's a 30-fold difference in our life span.

So it seems to me that in order for evolution to evolve a 30-fold difference in life span with so few really clear genetic differences, evolution maybe had to tweak hundreds, if not thousands, of genes. It's unlikely at the present time that we will find a single drug that's going to be able to do what evolution did.

Some Silicon Valley enthusiasts have been saying that life-span extension up to 500 or 1,000 years is feasible.
Well, it's religion. It's not science. I mean, that's all I can say. It's based on belief, not based on any data. People are certainly welcome to believe whatever they want to believe. But it doesn't make it true!

You've frequently emphasized that aging is a complex process, and that modifying it is not going to be quick or easy. Yet we all yearn for a solution.
Again, don't confuse aging and death. I am optimistic that we will experience medical interventions that will extend—the buzzword now is "health span." I think what terrifies people—certainly what terrifies me—is watching, for example, my mom, who is well into her 90s. She's losing cognitive function, she doesn't walk as well—and she's in pretty good shape! There are lots of people at her age who are confined to wheelchairs. That's aging, and that's terrifying. I am optimistic that we're on the cusp of understanding enough about that process to be able to intervene. And that people like us, who are not at that point, will benefit.

But we're still going to die. I'll remind you of the mouse models, where we eliminate senescent cells. There's a significant increase in median life span, but there's no increase in maximum life span. In a way, the mice died healthier. I think that's the goal, and I think that that's what the venture capitalists are hoping for, because that will be the kind of intervention that will be broadly applicable and will be very desirable. The conflict is with those who think that we're going to live to be 200 or 300 or more years old. That's not realistic at this point.

Let's say we are successful at slowing down or reversing aging, or extending health span. Are there any social or cultural impacts that you have concerns about?
No. In my lifetime, the population of the earth has not quite doubled, but it's getting there. That's unsustainable. The truth of the matter is, not having people die is not going to add much to the population of the earth the way the current rate at which we're producing new people is ruining the earth. So I think that this is ridiculous.

So I really don't see a downside to this. There are problems, but I don't think extending health span is going to exacerbate those problems.

Aging is terrifying. I'm optimistic that we're on the cusp of understanding enough about it to be able to intervene.

Statement of Ownership, Management and Circulation. (Required by U.S.C. 3685.) (1). Publication Title: MIT Technology Review, (2). Publication No. 535-940, (3). Filing Date: 9/19/2018, (4). Issue Frequency: Bi-monthly. (5). No. of Issues Published Annually: 6. (6) Annual Subscription Price: $29.95 (7). Complete Mailing Address of Known Office of Publication: One Main St, 13th Floor, Cambridge, MA 02142 (8). Complete Mailing Address of Headquarters or General Business Office of Publisher: One Main St, 13th Floor, Cambridge, MA 02142 (9). Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor. Publisher: Elizabeth Bramson-Boudreau, MIT Technology Review, One Main St, 13th Floor, Cambridge, MA 02142. Editor: Gideon Lichfield, MIT Technology Review, One Main St, 13th Floor, Cambridge, MA 02142. Managing Editor: Timothy Maher, MIT Technology Review, One Main St, 13th Floor, Cambridge, MA 02142. (10). Owner: Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02138. (11). Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or other Securities: None. (12). Tax Status: The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes: Has Not Changed During Preceding 12 months. (13). Publication Title: MIT Technology Review. (14) Issue Date for Circulation Data Below: July/August 2018. (15). Extent and Nature of Circulation: (a) Total No. of Copies (Net Press Run): Average No. Copies Each Issue During Preceding 12 Months: 207,334; No. Copies of Single Issue Published Nearest to Filing Date: 217,046. (b) Legitimate Paid and/or Requested Distribution (1) Mailed Outside-County Paid Subscriptions Stated on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 154,454, No. Copies of Single Issue Published Nearest to Filing Date: 154,722. (2) In-County Paid/Requested Mail Subscriptions Stated on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 0, No. Copies of Single Issue Published Nearest to Filing Date: 0. (3) Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid Distribution Outside the USPS: Average No. Copies Each Issue During Preceding 12 Months: 11,572, No. Copies of Single Issue Published Nearest to Filing Date: 9,936. (4) Paid Distribution by Other Mail Classes Through the USPS: Average No. Copies Each Issue During Preceding 12 Months: 0, No. Copies of Single Issue Published Nearest to Filing Date: 0. (c) Total Paid and/or Requested Distribution: Average No. Copies Each Issue During Preceding 12 Months: 167,026, No. Copies of Single Issue Published Nearest to Filing Date: 164,658. (d) Nonrequested Distribution: (1) Outside-County Nonrequested Copies included on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 705, No. Copies of Single Issue Published Nearest to Filing Date: 780. (2) In-County Nonrequested Copies included on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 0, No. Copies of Single Issue Published Nearest to Filing Date: 0. (3) Nonrequested Copies Distributed Through the USPS by Other Classes of Mail: 0 Average No. Copies Each Issue During Preceding 12 Months: 0, No. Copies of Single Issue Published Nearest to Filing Date: 0 (4) Nonrequested Copies Distributed Outside the Mail: Average No. Copies Each Issue During Preceding 12 Months: 4,581, No. Copies of Single Issue Published Nearest to Filing Date: 6,203. (e) Total Free or Nominal Rate Distribution: Average No. Copies Each Issue During Preceding 12 Months: 5,286, No. Copies of Single Issue Published Nearest to Filing Date: 6,983. (f) Total Distribution: Average No. Copies Each Issue During Preceding 12 Months: 172,313, No. Copies of Single Issue Published Nearest to Filing Date: 171,641. (g) Copies not Distributed: Average No. Copies Each Issue During Preceding 12 Months: 48,671, No. Copies of Single Issue Published Nearest to Filing Date: 59,657. (h) Total: Average No. Copies Each Issue During Preceding 12 Months: 220,984, No. Copies of Single Issue Published Nearest to Filing Date: 231,298. (i) Percent Paid: Average No. Copies Each Issue During Preceding 12 Months: 96.9%, No. Copies of Single Issue Published Nearest to Filing Date: 95.6%. This Statement of Ownership will be printed in the November/December issue of this publication. I certify that all the information furnished on this form is true and complete. Electronic Copy Circulation: Requested and Paid Electronic Copies: Average No. Copies Each Issue During Preceding 12 Months: 13,651 No. Copies of Single Issue Published Nearest to Filing Date: 14,252 Total requested and paid copies: Average No. Copies Each Issue During Preceding 12 Months: 167,026 No. Copies of Single Issue Published Nearest to Filing Date: 164,658. Total Requested Copy Distribution (Line 15F) + Requested Paid Electronic Copies: Average No. Copies Each Issue During Preceding 12 Months: 172,312 No. Copies of Single Issue Published Nearest to Filing Date: 171,641. Percent Paid and /or Requested (16B divided by 16c x 100): Average No. Copies Each Issue During Preceding 12 Months: 96.9% No. Copies of Single Issue Published Nearest to Filing Date: 95.9%. **Doreen Adger, SVP of Consumer Revenues and Marketing.**

Genes I wish they would find

By Sarah Cooper

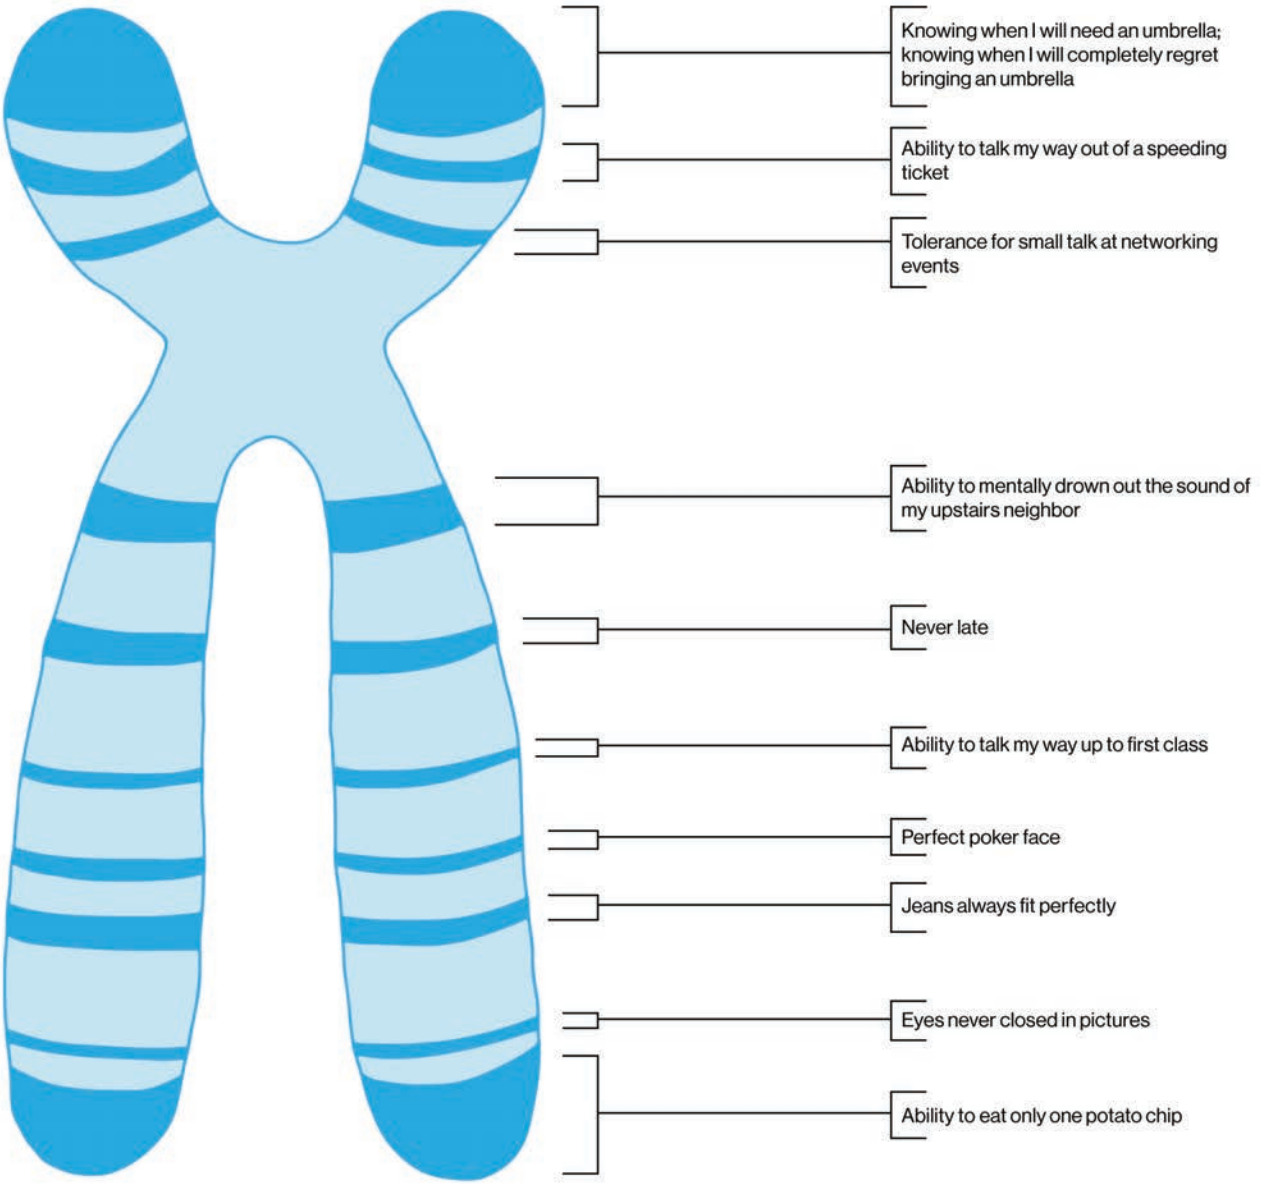

Sarah Cooper is a writer, comedian, and creator of the satirical blog TheCooperReview.com.

MIT Technology Review (ISSN 1099-274X), November/December 2018 issue, Reg. US Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2018. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: Send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the internet at www.technologyreview.com/customerservice. Basic subscription rates: $59.94 per year within the United States; in all other countries, US$69. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051, Fort Erie, ON L2A 6C7. Printed in USA. Audited by the Alliance for Audited Media.

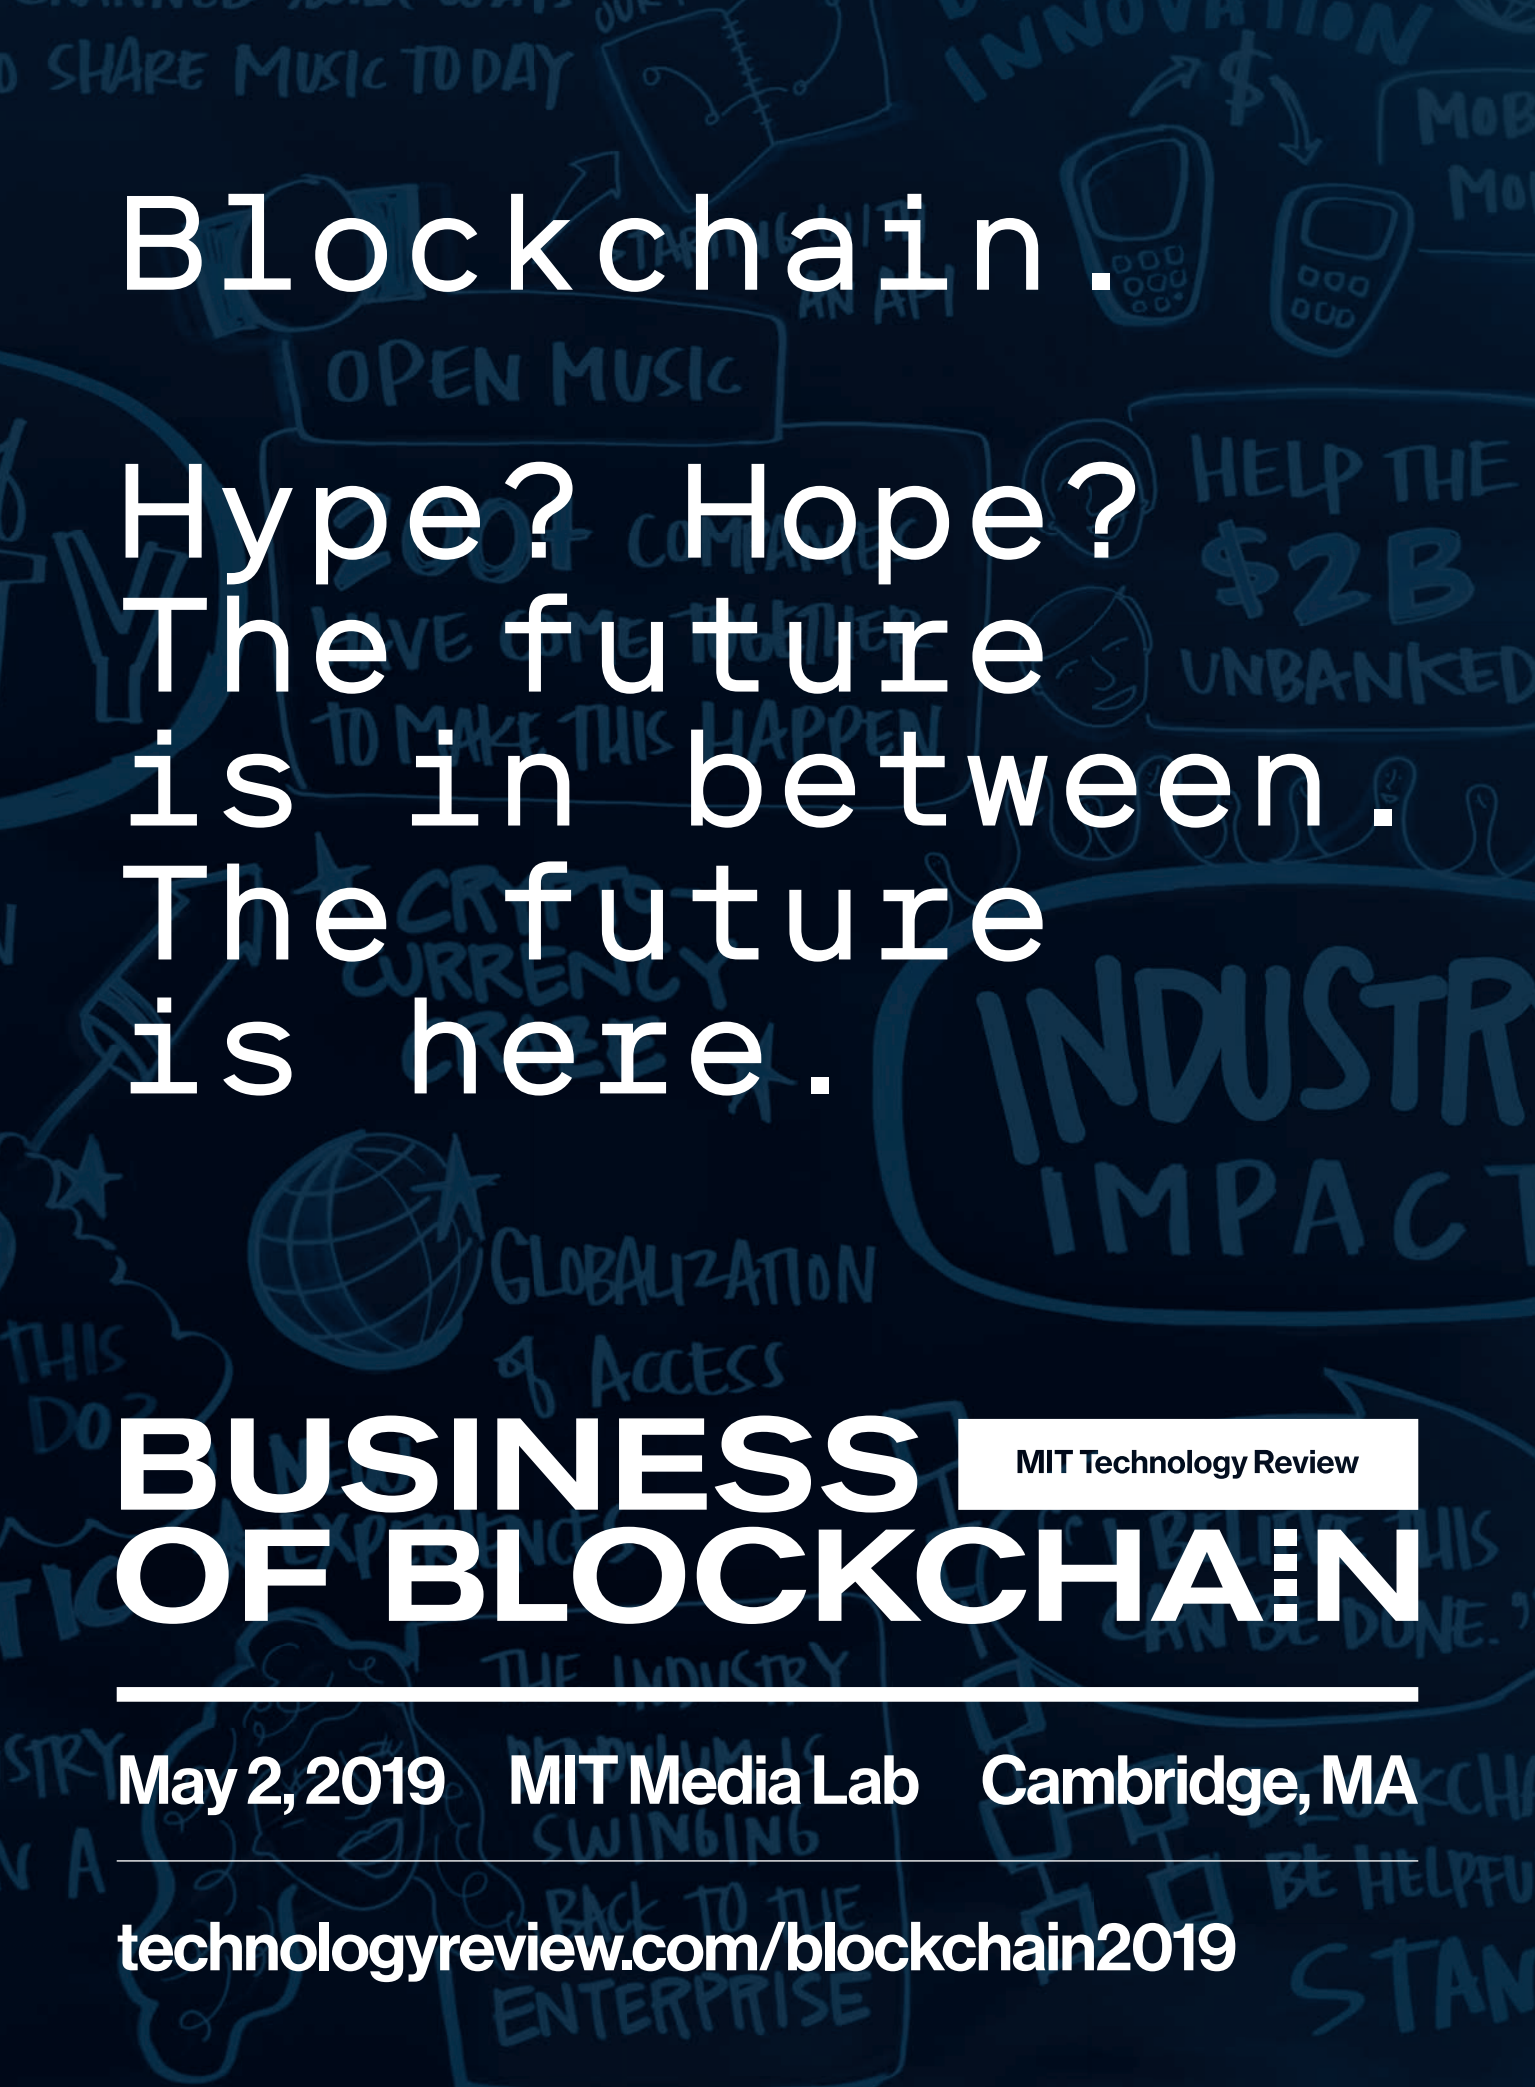
Blockchain.
Hype? Hope?
The future
is in between.
The future
is here.
BUSINESS
OF BLOCKCHAIN
MIT Technology Review
May 2, 2019 MIT Media Lab Cambridge, MA
technologyreview.com/blockchain2019

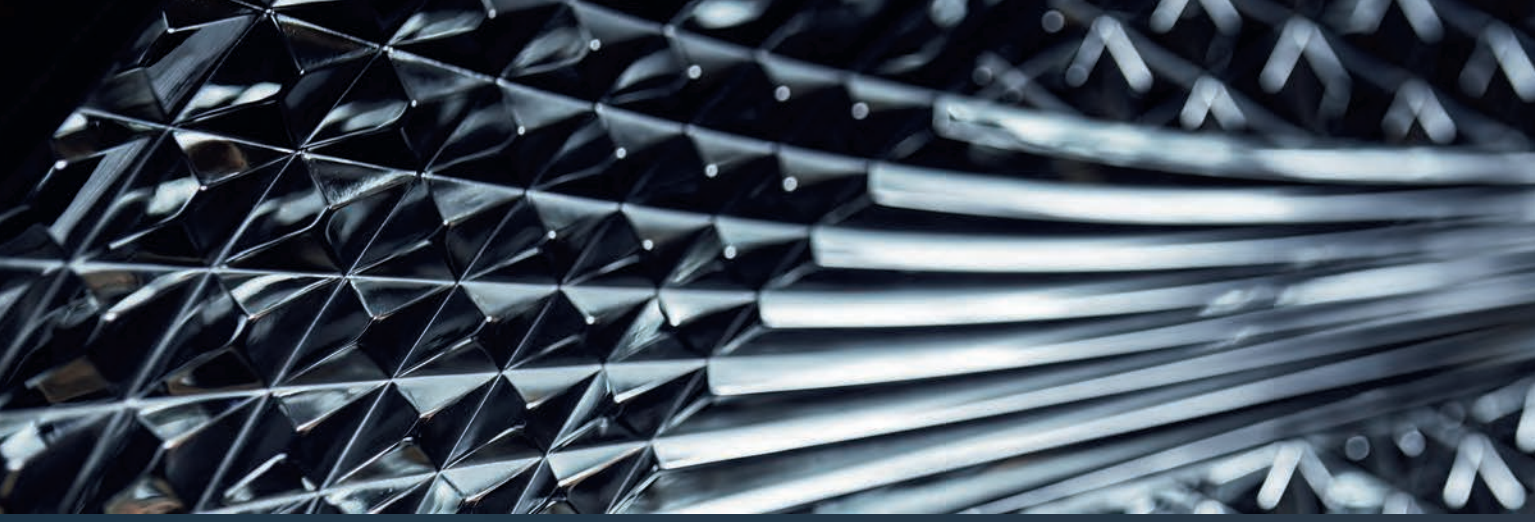

RARE FORM

THE LEXUS LS 500. LIVE IN THE NEW.

We didn't merely try something different. We crafted an entirely original experience. With accents like *Kiriko* glass, designed by hand then delicately etched with thousands of cuts to catch the eye and transform in the light. A 416-horsepower[1] twin-turbo engine paired with a 10-speed Direct-Shift automatic transmission takes you from 0 to 60 in just 4.6 seconds.[1,2] All this is complemented by cutting-edge technology with one of the largest Head-Up Displays in the industry.[3] The Lexus LS 500 isn't simply unique. It redefines what a flagship luxury sedan can be.

LEXUS
EXPERIENCE AMAZING

lexus.com/LS | #LexusLS

Options shown. 1. Ratings achieved using the required premium unleaded gasoline with an octane rating of 91 or higher. If premium fuel is not used, performance will decrease. 2. 2019 LS vs. 2018/2019 competitors. Information from manufacturers' websites as of 5/21/2018. 3. Performance figures are for comparison only and were obtained with prototype vehicles by professional drivers using special safety equipment and procedures. Do not attempt.