
What if that trial failed because we didn't give the drug to the right patients?
What if we closed schools during an influenza outbreak?
What if we held off on that transplant?
What if employers offered more paid sick days?
What if?
What if we redrew the map for organ allocation?
What if vaccination programs had never been implemented?

REALLY HARD QUESTIONS, ANSWERED BY MACHINES

BY ELAINE VITONE AND BRETT MURPHY

ILLUSTRATION BY MICHAEL LOTENERO

A bug goes viral. Inside each unlucky person who takes ill, organs send messenger proteins to one another in crosstalk to fend off infection. Within each of these organs, cells download and duplicate the virus. And all the while, both the cells and the viruses swap data among themselves, gather input from their environments, put it all together, and—most importantly—learn from it. It's reprogram yourself or die.

We biological beings are, at every level, a feedback loop on a mission—an intelligent system.

The machine-learning crowd figured this out decades ago, launching a whole class of techniques and algorithms that took cues straight from the life sciences. Ironically, medicine took a while to warm to the idea of what these two seemingly disparate disciplines, computing and biology, have to offer each other. (Though Pitt's lineage of using computers to solve real-world problems in health care dates back to the 1970s, when Jack Myers, an MD and the late chair of medicine, with Randolph Miller, MD '76, and Harry Pople Jr., created Internist-I, perhaps the first computer-aided diagnostic tool.)

Imagine you want to build a model of a biological process. It's a little bit like perfecting a cake recipe. Say you have 20 ingredients you're considering using. To decide how each variable contributes to the final product, you could go the trial-and-error route, baking Bundt cake after Bundt cake and omitting one ingredient each time. To try changing any two ingredients, you'd have to bake 190 cakes. To change any three, you'd need 1,140. Any four would take 4,845. Or, you could feed all the ingredients into a computer, explaining everything you know about how they interact with one another based on your experience. You could model thousands of what-ifs, coming up with a shortlist of possible recipes—then just bake and taste-test the ones least likely to flop.

In medicine, the "ingredients" for a model might be insights gleaned from the literature, clinical experience, lab experiments, historical records of epidemics, and other data—or some such combination thereof. Researchers run a simulation and check their *in silico* results, as those in the field like to call them, then analyze them for patterns that will inform their "recipe." Then, they gather more data as needed to fine-tune the model and fill in any gaps. Once the model proves viable, they can tweak the dials and test the what-ifs. It's an approach that works well in all sorts of

applications—biomedical device development, disease progression prediction, and so on—and it's holding increasingly more promise as Big Data grows bigger.

Imagine the possibilities with that bug we started with. You could model the molecular process of how it infects cells, duplicates, and spreads throughout the body. You could model disease vectors. (Pitt people have already modeled dengue fever outbreaks at the level of individual mosquitoes.) Take it a few steps further, and you could model how resistance emerges after patients drop various treatments. Drug resistance might then grow to become a population-wide problem. You could model that, too.

Eventually, the medical-computational-modeling community hopes, they'll be able to string all these various pieces together to create one giant *SimCity* of disease, rendering the inner workings of each one of its inhabitants down to the sub-cellular level. Test your what-ifs there, and you could significantly narrow your search for drug candidates, public health interventions, you name it, saving precious time, resources, and lives.

That's the dream. To realize it, Pitt people are delving into difficult questions about health care practice and policy, as well as how the body works. They're building new tools and forging the kind of cross-disciplinary, cross-institutional partnerships it will take to build this *SimCity*. They're asking questions that aren't so easy to ask with a clinical trial. Here are some of the stories behind the work and a few of the intriguing what-ifs these teams are tackling. *—Elaine Vitone*

FIRST, GET THE DATA

In his field, epidemiologist Don Burke, an MD and Pitt Distinguished University Professor of Health Science and Policy as well as professor of medicine, was an early adopter of modeling. His first simulation, which he published in *Nature* in 2004, identified previously unrecognized patterns in Thailand's dengue fever epidemic. He went on to publish similar epidemic analyses for the United States and Central Africa. Following the Sept. 11, 2001 attacks, he designed a smallpox-outbreak model that directly informed U.S. vaccination policy for biodefense preparedness.

Like with all simulation projects, Burke's began with a lot of homework to make sure the model was realistic. "We kept going back to the historical record," he says. (For the Thailand project, his team centralized one province's national reporting on dengue fever going back 30 years.) "And after doing that a number of times, we decided, 'Oh, let's go do it all.'"

By "it all," he meant build a single, centralized, open-source database of all infectious disease cases, everywhere. For as far back as the records go.

A lofty goal, for sure. But by that point—about eight years ago—he was well positioned to build the team that could tackle it. As Pitt's new dean of the Graduate School of Public Health—as well as its associate vice chancellor for global health, health sciences, director of its Center for Vaccine Research (CVR), and UPMC Jonas Salk Professor of Global Health—he'd brought to the University a coveted Models of Infectious Disease Agent Study (MIDAS) grant from the National Institutes of Health. (Pitt has since been named a MIDAS National Center of Excellence.)

To support the MIDAS effort, Burke founded a modeling motherboard of sorts, formally known as the Public Health Dynamics Laboratory (PHDL). A collaboration between Pitt, Carnegie Mellon University, and the Pittsburgh Supercomputing Center, the lab plans to make computational modeling in epidemiology an accessible, everyday tool for students, researchers, public health decision-makers, and anyone else interested. A number of the lab's members are Burke recruits from fields you might not expect in the health sciences—statistical physicists, computer scientists, game theorists, and machine learning experts—whom he proudly calls "hardcore computationalists."

In recent years, PHDL has gone public with that historical database Burke dreamed of. Thus far, Project Tycho, as it's called, includes records for the entire United States, and later this year, the team plans to link it to records from Brazil, Taiwan, and France. The group has also launched FRED, a platform that allows you to simulate the spread of disease from the comfort of your own smartphone. (More on these later—see below and p. 22.) Burke's hope is that the enthusiasm of these early adopters will become … well, infectious. *—EV*

A *WHAT-IF* GENERATOR

Diseases interact with their environment and can't be understood in a vacuum. Like the people who carry them, their reactions differ from scenario to scenario. They're dynamic.

A new modeling platform called FRED—Framework for Reconstructing Epidemiological Dynamics (the acronym honors Fred Rogers)—allows researchers to chart the paths of epidemics and the effects of mitigation strategies, viral evolution, and personal health behavior. "We're trying to tie together things that happen inside human beings and, essentially, the population impact of interventions," says John Grefenstette, a PhD professor of biostatistics in Pitt Public Health and director of the Public Health Dynamics Laboratory at Pitt.

FRED uses census-based, synthetic populations of the entire United States. (What's a synthetic population? Computer-generated data based on actual demographics—"virtually real people without the possibility of a privacy infringement," says Grefenstette.)

The open-source simulator is available online and will be released as an app. It allows you to create a scenario in any U.S. county by controlling "levers"—related to factors like school cancellation days or vaccination rates—that replicate "health-related human behavior based on demographic characteristics."

The tool takes into account the personalities of different places—for instance, an older Pittsburgh population versus a younger Salt Lake City. "Our policies are going to have locally different effects," says Grefenstette.

FRED KNOWS

Here's the difference the length of school closures can make on an influenza epidemic in Allegheny County. In this model, schools are set to close after 10 kids get sick.

The top graphs both show that the number of cases is decreased by school closure but not by as much as you might think, though closures do delay outbreaks. (The "attack rate" is the cumulative percentage of persons infected in the population.) When schools re-open, the epidemic trucks on.

With FRED, you can create animations to see how the epidemic will play out based on the parameters you set.

To try it for your county, go to ... fred.publichealth.pitt.edu/simulator/

FRED: ALLEGHENY COUNTY
INFLUENZA
EFFECTS OF SCHOOL CLOSURE ON INCIDENCE

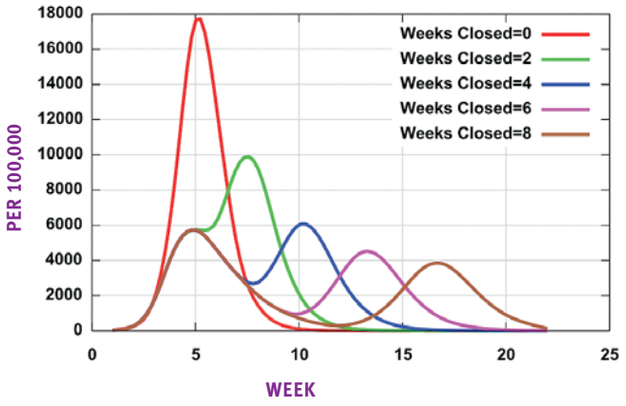

FRED: ALLEGHENY COUNTY
INFLUENZA
EFFECTS OF SCHOOL CLOSURE ON ATTACK RATE

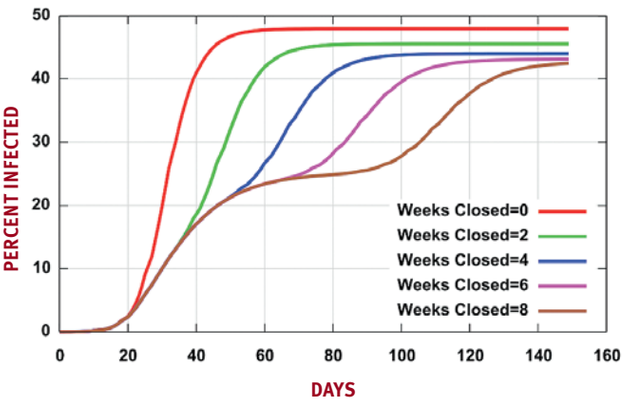

INFLUENZA IN ALLEGHENY COUNTY: NO SCHOOL CLOSURE
DAY (35)

INFLUENZA IN ALLEGHENY COUNTY: 2 WEEK SCHOOL CLOSURE
DAY (35)

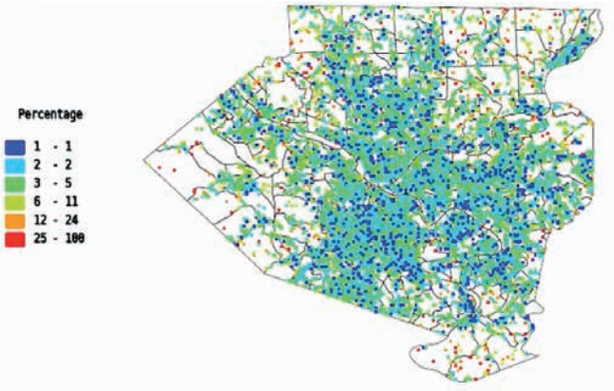

What if we closed schools during an influenza outbreak?

What if employers offered more paid sick days?

The thinking behind FRED was interdisciplinary, he adds. "We had people from the department of health, medical doctors, lawyers, statisticians, and computer people [discussing] what would be the highly relevant questions that we could ask with our models." One practical question the team came up with was: *What if employers offered more paid sick days?* They published their results in the *American Journal of Public Health* last year. Using FRED and data from the U.S. Bureau of Labor Statistics, they concluded that just two extra paid "flu days" would reduce workplace infections by almost 40 percent.

—Brett Murphy

SPECIAL AGENTS

An agent-based model, like what FRED creates, simulates activities of autonomous actors (maybe individual pathogens, people, or organizations) and digests how those goings-on influence the system as a whole. The folks who come up with these kinds of models immerse themselves in fields most people have never heard of—game theory, complex systems, computational sociology, and evolutionary programming. (And they probably like *The Sims* video game.) Their models allow the curious to evaluate a design and its effects on people and places without actually implementing it in the real world—say, what a traffic light might mean for commuters on Main Street, the implications of an invasive species entering the Rhine River basin, or the ripple effect of a novel vaccine.

—Brett Murphy and Erica Lloyd

TOP ROW: BMC PUBLIC HEALTH. 2013 OCT 8;13:940. DOI: 10.1186/1471-2458-13-940.
BOTTOM ROW: COURESY PHDL.

WHAT IF THAT TRIAL FAILED BECAUSE WE DIDN'T GIVE THE DRUG TO THE RIGHT PATIENTS?

CATCHING SWELLS

When physicians talk about sepsis, a word they might use to describe it is *cascade*. But the image that comes to mind for these docs is probably not a gentle waterfall. The physiological response that is sepsis can be every bit as catastrophic as a tsunami. And patient outcomes are all over the map—one severely injured patient who ends up with sepsis (a systemic inflammation tied to infections) can do far better than another with more moderate injuries, for example. The rhyme or reason of it all has eluded scientists.

In the late '90s, a few research groups thought sepsis might respond to a TNF-targeting drug as a possible treatment. TNF (a.k.a., tumor necrosis factor) has been used for decades as a sepsis biomarker, a blood test that signals to physicians when the tide is rising. The drug seemed promising at the outset—animal and preliminary human-trial results were encouraging. But a phase III clinical trial was a dud; it had mixed results. Though many patients benefited, many others were harmed.

Frustrated by these and other dead ends in this confounding condition, Yoram Vodovotz, PhD professor of surgery, Gilles Clermont, MD associate professor of critical care medicine, and mathematician Carson Chow—all of the University of Pittsburgh—hatched a plan for a new approach: to build the first *in silico* model of severe sepsis.

Colleagues told them they were crazy. Sepsis is just too complicated to simulate, they said. But that, Vodovotz recalls, was exactly the point.

"The conscious mind can't handle more than a few things [at once]," he says. "But the unconscious mind can do it quite well. My scientific mentors could integrate huge amounts of information and just go, 'I believe the system plays like this.' Really good, experienced doctors do the same thing. [Modeling gives you] the best of both worlds: the rational process that comes out of your conscious mind, integrated with the ability of your unconscious mind."

After reviewing the literature, the team chose biological parameters that appeared to be important in sepsis: the duration of the precipitating infection or injury, the patient's blood pressure, and the level of dysfunction in patient tissues, among others. Using algorithms designed by Chow, they ran the simulations and watched the resulting changes in endotoxin, cytokine, and other protein levels in the hours, days, and weeks after injury. Vodovotz then validated the model by comparing the simulation results to those of his own follow-up studies of cellular processes in the lab. Then the team ran the simulation again.

It's all about relationships, he says. Instead of focusing on the individual players themselves—the various inflammatory markers and whatever molecular processes might be at work within and among them—first look for patterns in how the players affect one another over time: A inhibits B and C, B inhibits C and A, and so on. That makes the time you spend in the lab much more focused and efficient.

In 2004, the team put their model to the test by re-running the failed anti-TNF-drug study *in silico*—and found comparable results in their simulated patients. The silver lining in all this bad news was that the Pitt study proved that modeling sepsis was an idea that could hold water. And unlike with clinical or laboratory trials, simulated sepsis could be rewound, paused for further pondering, and even altered. Scientists could ask important questions, like: Why was the drug good for some people and bad for others? What separates the two groups of patients? Could the trial have succeeded had the drug been given to a more select group of patients?

A decade later, they're still asking these and other questions about sepsis—and much more, as the scope of their work continues to grow. They're studying a number of other inflammatory "cascades," as well, including liver failure and trauma. (For the latter, Vodovotz and colleagues recently launched a 500-patient study to serve as a data storehouse.)

Their findings are nonlinear. So, in the case of sepsis, yes, high TNF levels are a bad sign, but that doesn't necessarily mean that low TNF is a good thing. Inflammation is more complicated than that—but not unfathomable, says Vodovotz.

In addition to some 70 papers illuminating the vast and highly complex ocean that is acute inflammatory response, the team's "crazy" idea (modeling sepsis, that is) has also led to the founding of a field. The Society for Complex Acute Illness, of which Vodovotz and Clermont are cofounders, now has 150 members. It also led to the founding of a biosimulation company in Pittsburgh. Since 2001, Immunetrics has helped some 20 studies build more successful laboratory and clinical trials. —*EV*

WHAT IF WE HELD OFF ON THAT TRANSPLANT?

MODEL PATIENTS

When a patient receives a new liver, not only is she married to a physiologically taxing regimen of immunosuppressants forever, but she's also opening up a daunting new set of what-ifs: What if the transplant doesn't help? What if that organ could have saved the life of someone else on the transplant list?

Mark Roberts—MD professor and chair of health policy and management in Pitt's Graduate School of Public Health and professor of medicine, of industrial engineering, and of clinical and translational science—has been wrestling with these questions for more than a decade.

His team gathered and analyzed extensive data on disease progression from patients with end-stage liver disease. From these, the team created thousands of virtual people on virtual waiting lists—a model of every member of the U.S. organ allocation system—"each with their virtual physiologies going on," he says. "And now we can say, 'Okay, what would happen if you changed the rules? What if, instead of [allocating an organ to] the sickest person first, you did the person who would benefit the most? Or what if you eliminated the regional preference?'"

Once recovered, donated livers have a shelf life of 18 hours, tops. In his systematic what-iffing, Roberts has shown that more organs might be transplanted—and more lives saved—in time if the regional map for organ allocation were

SPONTANEOUS SURVIVORS

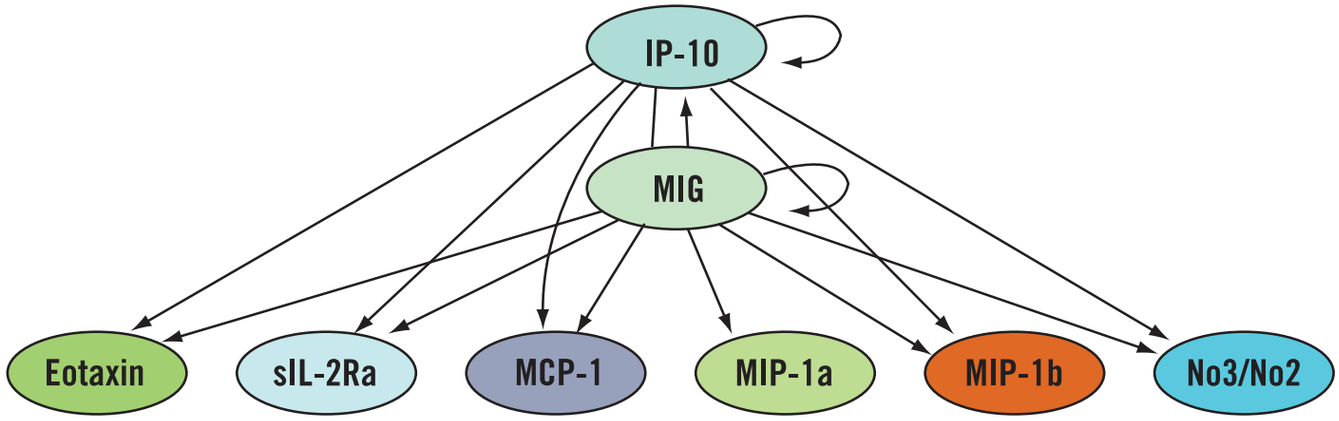

NONSURVIVORS

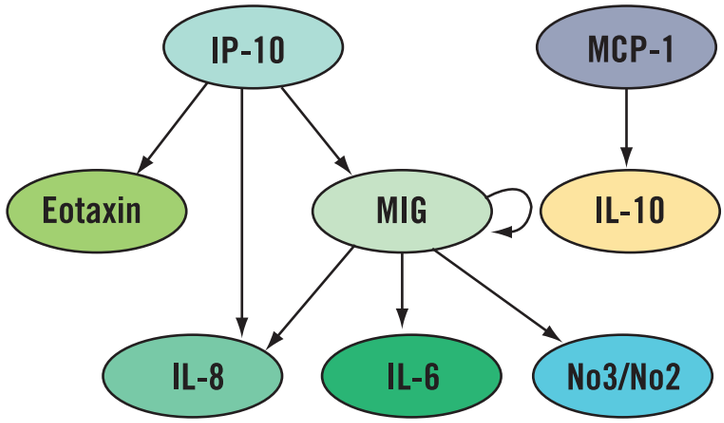

Previously unpredictable disease progression revealed: Children who spontaneously survive acute liver failure share a network of inflammatory responses that's a lot like what's seen in children post-transplant.

LIVER TRANSPLANT

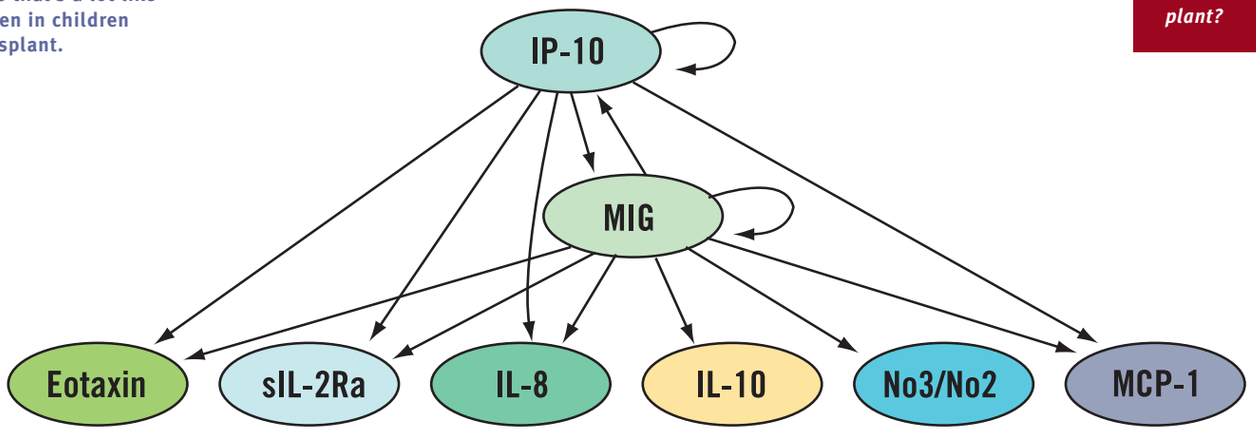

WHAT CHILD IS THIS? Doctors have always had difficulty predicting which children with liver disease would survive without a transplant. Results from a 14-year multisite clinical study give pediatric specialists a new lens. The graphics shown here plot how various inflammatory mediators interact differently among patient groups with different outcomes. Suddenly, says Pitt's Yoram Vodovotz, the researchers "could easily tell the groups apart." The findings are informing a model that allows the Pitt team to get help treating—for now—virtual patients. They can ask questions like: ***Who needs to get on the transplant list today? And who will do well without a transplant?***

GRAPHIC BY MICHAEL LOTENERO, ADAPTED FROM AZHAR N. ET AL. (2013) ANALYSIS OF SERUM INFLAMMATORY MEDIATORS INDENTIFIES UNIQUE DYNAMIC NETWORKS ASSOCIATED WITH DEATH AND SPONTANEOUS SURVIVAL IN PEDIATRIC ACUTE LIVER FAILURE, PLOS ONE 8(11): E78202. DOI:10.1371/JOURNAL.PONE.0078202.

redrawn, among other findings.

Recently, Roberts teamed up with a Pitt group—including Yoram Vodovotz, of surgery—that's exploring another ethical conundrum in transplant medicine, one that arises in cases of pediatric acute liver failure (PALF). This devastating condition can result from poisoning, acetaminophen overdose, infection, or—as is the case with almost half of these kids—for reasons that are never discovered. PALF can take a child from perfect health to the ICU in a matter of weeks, or even days. Without a liver transplant, many will die. And, for reasons no one can explain, many others won't.

Sometimes, a child is put on the transplant list, seemingly at death's door, and then makes a full recovery before a match for an organ can be found. Which raises a delicate question: Are we doing too many transplants?

And the short answer, says Vodovotz, is, *Yes*.

The team didn't come to this conclusion lightly—or easily. It was informed by the culmination of a 14-year clinical study by a multinational consortium. The Pediatric Acute Liver Failure Study Group, as it's called, was funded by the National Institutes of Health and the National Institute of Diabetes and Digestive and Kidney Diseases and led by Pitt's Robert Squires, professor of pediatrics. The study is perhaps the first to consider the distinct outcomes of the disease—survival with native liver, death with native liver, and transplant—separately, says Squires. (Most of the chidren who were part of the transplant group in the study survived.)

After comparing the inflammatory networks of the patient groups, the team arrived at an intriguing finding: The progression of protein interplay seen in the bloodwork of survivors with native livers and that of the transplant recipients (post-transplant) look markedly similar (see p. 17).

Vodovotz explains that taking blood samples to check for levels of inflammatory mediators has never been helpful in predicting which children could survive without a transplant. But the team found that after drawing blood each day, watching how these levels change, and analyzing how these mediators influence one another over time, a new picture emerged.

"If you look at the network representation, which says how mediators are interplaying with one another, it's a night-and-day difference. You could easily tell the groups apart," he says. The study was published last November in *PLOS ONE.*

Vodovotz and Roberts have started a new model: thousands of virtual boys and girls with PALF, each with his or her own virtual physiology and each facing the decision of whether to get on to the virtual liver-transplant waiting list.

"So we can start doing scenarios and say, 'Let's not transplant this virtual child today. Let's wait until tomorrow and see if [she's] any better,'" says Roberts. "And we can test different strategies for listing a child. We can make reasonable predictions about whether we do that child a service by transplanting [her] or not, and when would be the optimal time to list that child for transplantation." *—EV*

A MONTH'S DIFFERENCE

Female sterilization is the second most common contraceptive in the United States, even though Medicaid patients who elect to have the procedure are subjected to a 30-day waiting period. In a study published in the journal *Contraception* and discussed in a recent *New England Journal of Medicine* editorial, Pitt's Sonya Borrero and Kenneth Smith, both MDs in the Department of Medicine, with collaborators, explored what would happen if policymakers were to revise that rule. Women often request to have their "tubes tied" (tubal ligation) while in the hospital after giving birth. The researchers knew, anecdotally, that the mandate could make scheduling the procedure difficult. Patients with private insurance have no such waiting period imposed on them.

So, what if the mandated month-long lag between the request and procedure didn't exist? After building a model, known as a cost-effectiveness decision analysis, based on real Medicaid data (see the brackets on the right), the team concluded that fulfilled sterilization requests would increase by 45 percent.

Here's how the analysis works. All women who request sterilization under Medicaid enter the model. The model then simulates potential outcomes over the course of one year. Researchers can compare what happens with the current policy against a parallel Medicaid universe, which simulates outcomes with an imagined revised-policy branch of the model. Under a revised policy, the probability of women actually receiving the procedure increases with the 30-day barrier removed. Smith says that, annually, such an increase could prevent more than 29,000 unintended pregnancies and save the Medicaid program $215 million by avoiding the costs of childbirth from such pregnancies.

The Medicaid rules also require that women sign a consent form. Yet "assessments of the form's readability indicate that it is overly complicated, and its literacy level is too high for the average American adult," Borrero and coauthors write in the *NEJM*. In a related study, a Borrero team found that 34 percent of the women who read the form did not realize that a tubal ligation was permanent, and many did not realize there were reversible alternatives. Any new policy should have more readable documents to ensure that patients understand their options, the researchers say.

Borrero et al. point out that it is important to be sensitive to the idea that the fertility of the poor seems to be less valued by society. In fact, the Department of Health, Education, and Welfare first established a waiting period in 1976 after numerous troubling reports from that time: Poor women were being pressured into sterilization as part of local or state family planning programs. Health care providers sometimes suggested that welfare and other benefits were tied to sterilization and often didn't get proper patient consents.

Some women are still vulnerable, Borrero notes, pointing out that serious questions have been raised about the sterilizations of 150 women in California prisons between 2006 and 2010.

The *NEJM* editorial authors write, "Although [Medicaid's] policy was designed to protect vulnerable populations, we believe that it does not effectively fulfill that intention—in fact, it restricts the reproductive autonomy of the women it intends to serve." *—BM*

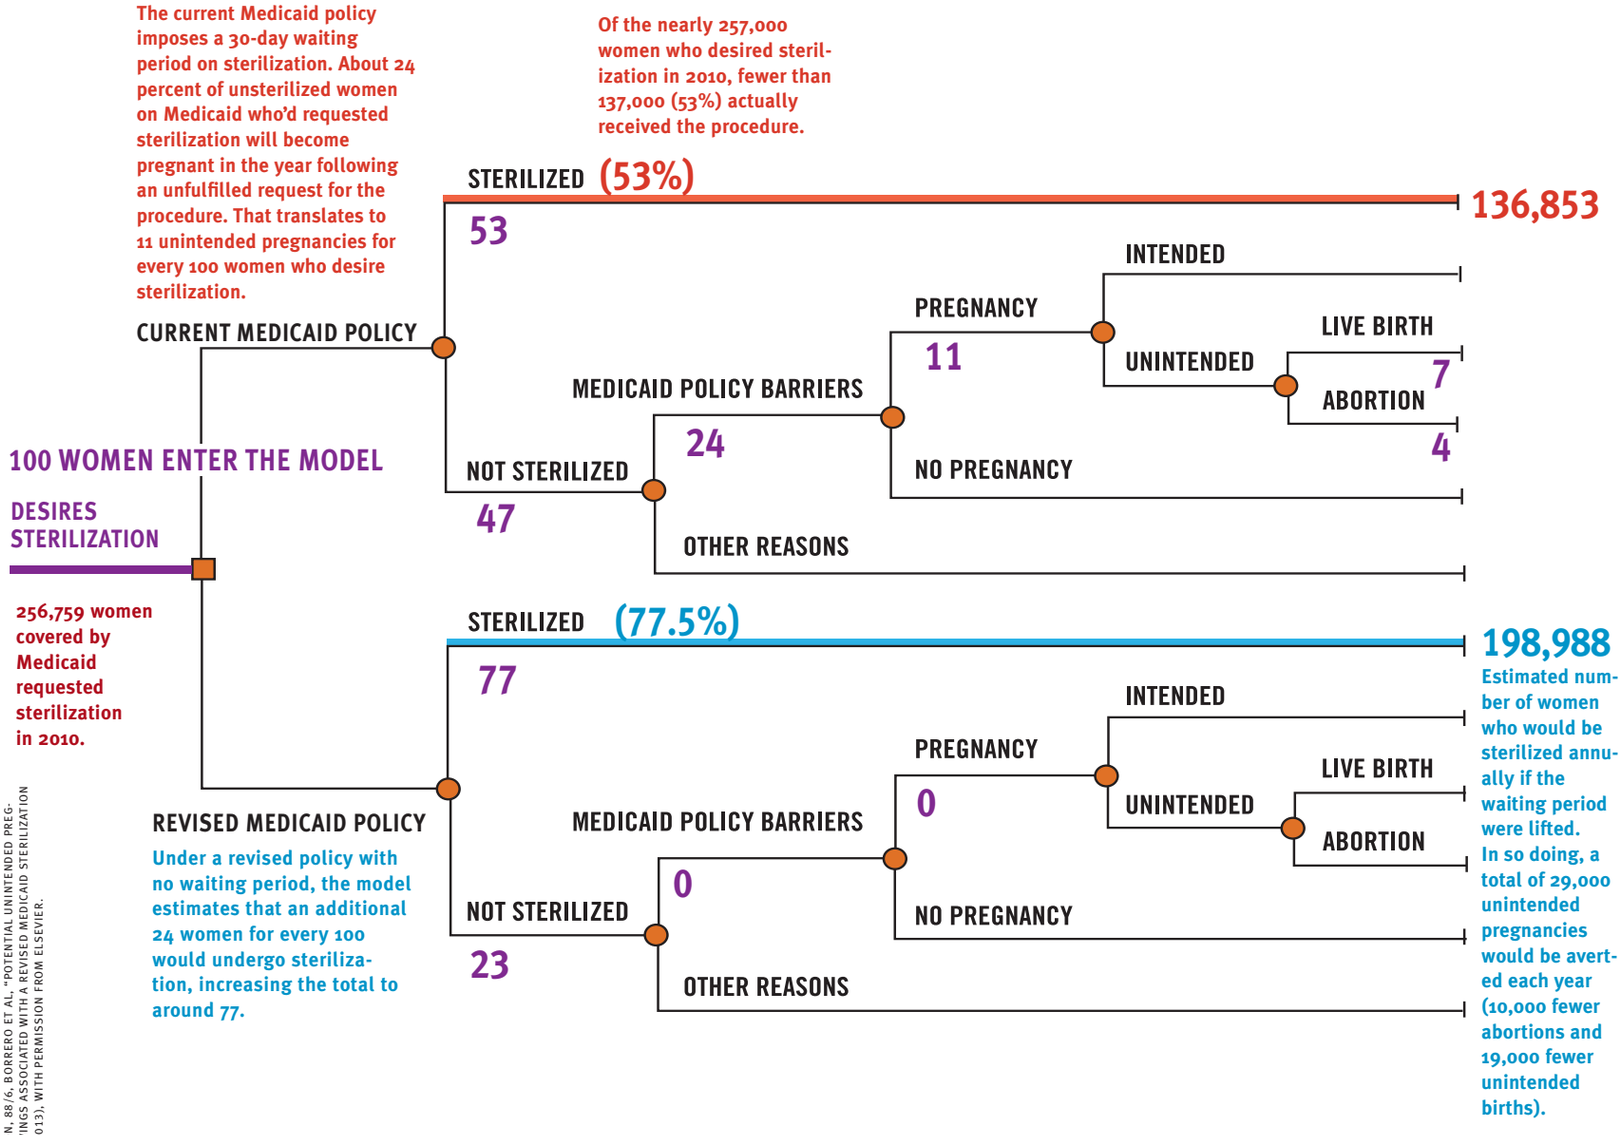

WHAT IF MEDICAID LIFTED THE 30-DAY WAITING PERIOD MANDATED FOR FEMALE STERILIZATIONS?

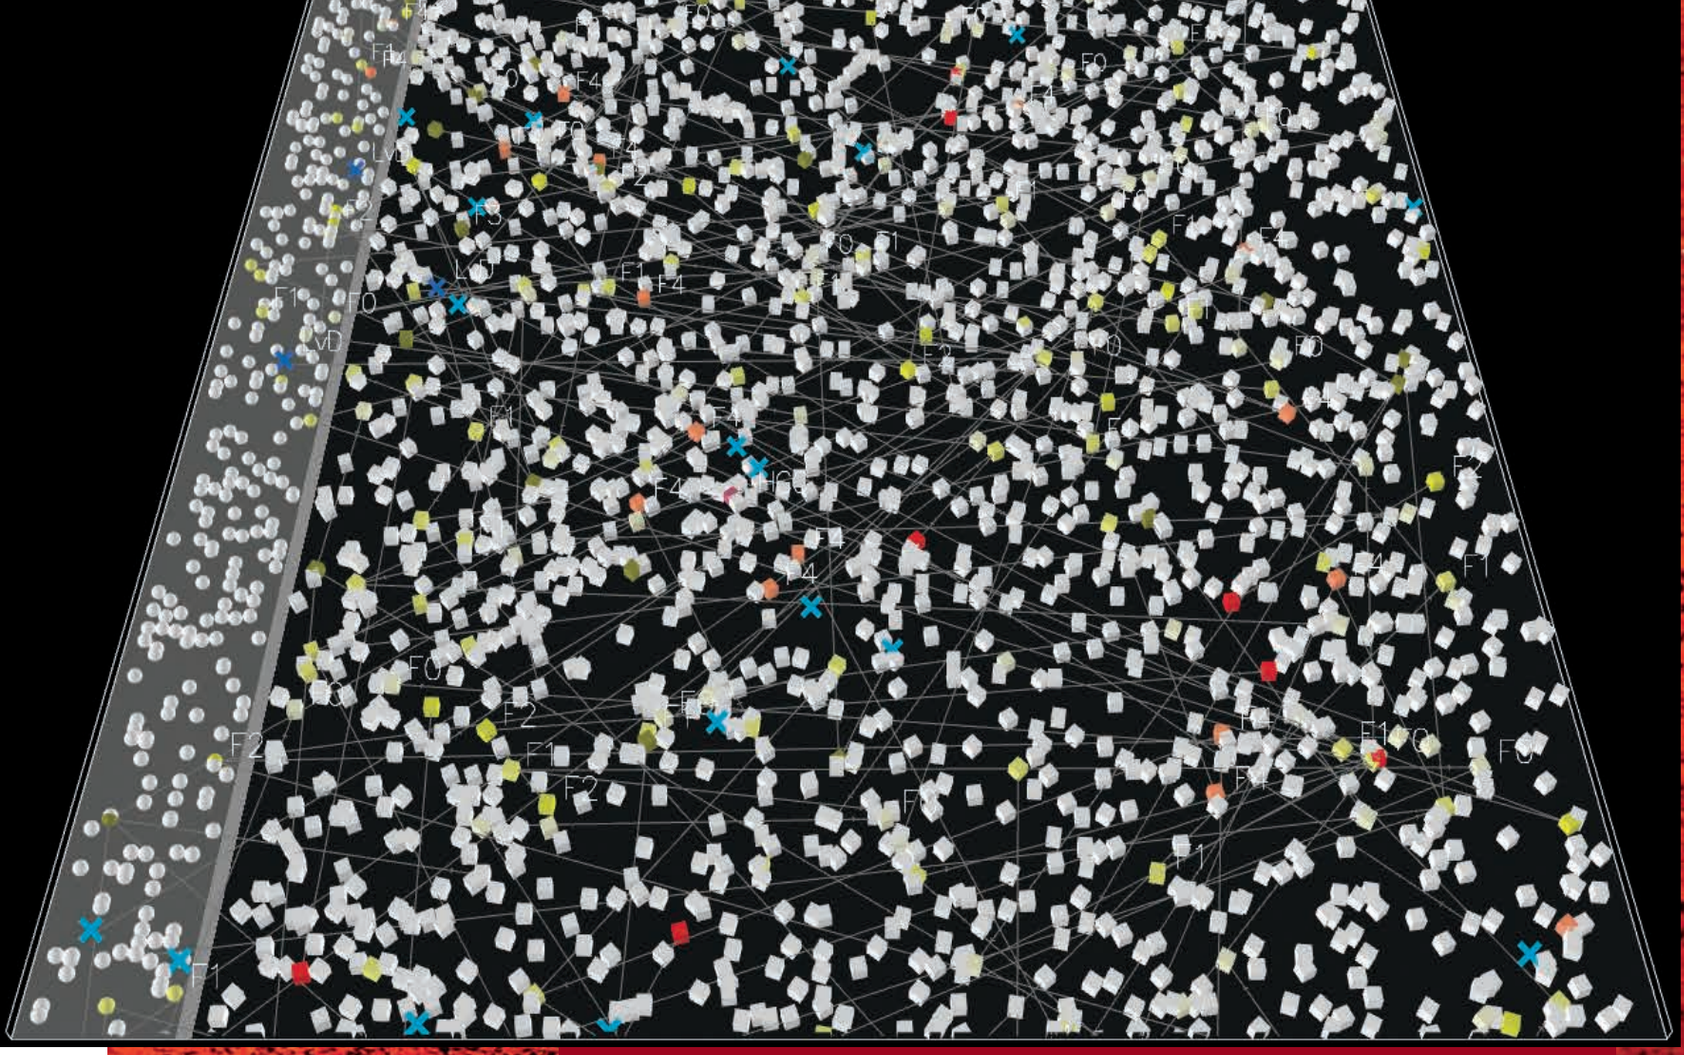

BOTH SIDES OF THE PRISON FENCE: Screening, and when appropriate, treating, inmates for hepatitis C is probably an effective way to save money and protect society at large from the disease, researchers think—even with treatment costs at about $100,000 a patient. The simulation above shows a 1,000-person sample representative of the entire U.S. population. Incarcerated individuals are shown as dots in the shaded region to the left. The blocks are people living freely in the United States. The lines represent infections spreading from person to person.

THINKING INSIDE AND OUT

More than 3 million people in the United States are infected with hepatitis C, a leading cause of chronic liver disease. Between 50 and 75 percent of them don't even know they have it.

Hep C is often transmitted intravenously. And prisons nationwide—with hep C prevalence documented at rates as high as 35 percent, though no standard screening protocols exist—have become a hotbed for the disease, says Jagpreet Chhatwal, a PhD assistant professor at the M.D. Anderson Cancer Center in Houston. While he was an assistant professor of health and policy management and of industrial engineering at Pitt two years ago, Chhatwal teamed up with Pitt's Mark Roberts, MD professor of medicine, Pitt's John Grefenstette, a PhD and director of the Public Health Dynamics Laboratory, and Tianhua He from Tsinghua University in China. They started developing a model that would answer some questions about hep C: What if prisons routinely screened all inmates for hep C and then treated those found to be infected? What would be the cost? What would be the benefits to society at large?

"If we can model the prison system, we can predict the disease impact on intervening while everyone is still inside," Chhatwal says. Many inmates are released unaware they even have hep C.

Using Bureau of Justice statistics, the investigators developed an "agent-based" model (see "Special Agents," p. 15) to simulate people moving between prisons and society and the spread of hep C. "Imagine you're looking at a video game with individuals moving in and out of the [prison] system with certain disease characteristics," Chhatwal says. The model takes into account variables like disease stage, an individual's behavior, access to treatment, and whether a person is aware of the infection.

With the advent of new drugs last year, Chhatwal notes, "the treatment duration has reduced from 48 to 12 weeks." But because it would cost around $100,000 to treat a single

WHAT IF WE SCREENED ALL INMATES FOR HEP C?

patient, and many inmates show no symptoms, prisons have little incentive to change screening policies. (Once prison officials learn of a case of any illness, law requires that the patient be treated.)

Chhatwal estimates that hundreds of thousands of hep C infections could be prevented in the United States throughout the next 10 years if infections in inmates were routinely identified; however, he notes that the team is still validating its conclusions. (The researchers' final estimates will be published this summer as an abstract in *Gastroenterology*.) With their current software, the researchers can simulate up to a 10,000-person sample; that can take several days. They eventually want to translate the model onto the FRED interface to run simulations on the entire U.S. population of 300 million. (See p. 14 to find out what's new in the neighborhood of mass modeling.)

Chhatwal says the model is predicting that people on both sides of the prison fence would benefit from looking out for inmates with hep C: By neglecting the likelihood of infection among this population, he says, "society will bear the burden at some stage." —*BM*

GRAPHIC COURTESY JAGPREET CHHATWAL

OUT WITH THE OLD, IN WITH THE PNEU

Pitt's Kenneth Smith, an MD and professor of medicine, wondered: Is the new pneumococcal vaccine better than the old? And for whom? These vaccines are designed to ward off bacterial pneumonia, bloodstream infections, meningitis, and other infections.

Using national health databases and what's known as a Markov state-transition model, his team found that the older vaccine, usually given to the 65-and-up crowd, ultimately "costs more and had a somewhat smaller spectrum in terms of the types of pneumococcal diseases that it prevented," he says. (The current standard also recommends it for younger persons with high disease risk.)

Published in the *Journal of the American Medical Association* in February 2012, their paper concluded that the new 13-valent pneumococcal conjugate vaccine (PCV13) makes the most economic and health sense for patients over the age of 50, regardless of their medical condition.

The simulations, Smith adds, were sensitive to "herd immunity" caused by children who'd been introduced to the new vaccine. "Kids get the newer vaccine on a routine basis, and that has changed the types of organisms that are causing disease. It's basically cut down the amount of disease the entire population gets." —*BM*

WHAT IF WE USED THIS NEW VACCINE INSTEAD OF THE OLD ONE?

AND WHATNOT

More than a decade ago, Gilles Clermont, MD associate professor of critical care medicine at Pitt, cofounded Immunetrics—a computational modeling software company that's turned what-iffing into a viable Pittsburgh-based biotech enterprise. Immunetrics is now chugging along without him. More recently, he's been exploring ways to use modeling, machine-learning, and other data-driven technology in new smart gadgets in health care.

Big Data, particularly the emerging understanding of biology at the mechanistic level, is opening up opportunities for helping patients. Yet, Clermont cautions, "More data does not necessarily correspond with more knowledge. We're really trying to bridge that gap between data and knowledge in novel ways."

On these projects he collaborates with the likes of associate professor of chemical and petroleum engineering Robert Parker; William Kepler Whiteford Professor of Industrial Engineering Andrew Schaefer; research assistant professor of industrial engineering Louis Luangkesorn; professor of critical care medicine Michael Pinsky; and nursing professor of acute and tertiary care Marilyn Hravnak—all of Pitt. Another collaborator is Artur Dubrawski, senior systems scientist at Carnegie Mellon's Robotics Institute. Here are some of the gizmos they have in the works:

- **An artificial pancreas system that maintains desired blood sugar levels in critically ill patients.**
- **An alert system to help physicians flag possible medical errors at the bedside.**
- **A hospital "air traffic controller" on the lookout for ways to keep patient flow humming along smoothly.**
- **A 15-minute health "forecast" system to give critical care docs a heads-up on which patients are headed for trouble—so the physicians can steer them clear of the storm.**

"The more data we have, the more tools we're going to need to cast it—to reinforce, destroy, or remodel our conceptual framework of how the world works," says Clermont.

"This also applies to finance and economics. It's not unique to health care." —*EV and EL*