

Maia Szalavitz, Backchannel. 08.11.2021 06:00 AM

Online version at:

<https://www.wired.com/story/opioid-drug-addiction-algorithm-chronic-pain>

**The Pain Was Unbearable. So Why Did Doctors Turn Her Away?** A sweeping drug addiction risk algorithm has become central to how the US handles the opioid crisis. It may only be making the crisis worse.

One evening in July of 2020, a woman named Kathryn went to the hospital in excruciating pain.

A 32-year-old psychology grad student in Michigan, Kathryn lived with endometriosis, an agonizing condition that causes uterine-like cells to abnormally develop in the wrong places. Menstruation prompts these growths to shed—and, often, painfully cramp and scar, sometimes leading internal organs to adhere to one another—before the whole cycle starts again.

For years, Kathryn had been managing her condition in part by taking oral opioids like Percocet when she needed them for pain. But endometriosis is progressive: Having once been rushed into emergency surgery to remove a life-threatening growth on her ovary, Kathryn now feared something just as dangerous was happening, given how badly she hurt.

In the hospital, doctors performed an ultrasound to rule out some worst-case scenarios, then admitted Kathryn for observation to monitor whether her ovary was starting to develop another cyst. In the meantime, they said, they would provide her with intravenous opioid medication until the crisis passed.

On her fourth day in the hospital, however, something changed. A staffer brusquely informed Kathryn that she would no longer be receiving any kind of opioid. “I don’t think you are aware of how high some scores are in your chart,” the woman said. “Considering the prescriptions you’re on, it’s quite obvious that you need help that is not pain-related.”

Kathryn, who spoke to WIRED on condition that we use only her middle name to protect her privacy, was bewildered. What kind of help was the woman referring to? Which prescriptions, exactly? Before she could grasp what was happening, she was summarily discharged from the hospital, still very much in pain.

Back at home, about two weeks later, Kathryn received a letter from her gynecologist’s office stating that her doctor was “terminating” their relationship.

Once again, she was mystified. But this message at least offered some explanation: It said she was being cut off because of “a report from the NarxCare database.”

Like most people, Kathryn had never heard of NarxCare, so she looked it up—and discovered a set of databases and algorithms that have come to play an increasingly central role in the United States’ response to its overdose crisis.

Over the past two decades, the US Department of Justice has poured hundreds of millions of dollars into developing and maintaining state-level prescription drug databases—electronic registries that track scripts for certain controlled substances in real time, giving authorities a set of eyes onto the pharmaceutical market. Every US state, save one, now has one of these prescription drug monitoring programs, or PDMPs. And the last holdout, Missouri, is just about to join the rest.

In the past few years, through a series of acquisitions and government contracts, a single company called Appriss has come to dominate the management of these state prescription databases. While the registries themselves are somewhat balkanized—each one governed by its own quirks, requirements, and parameters—Appriss has helped to make them interoperable, merging them into something like a seamless, national prescription drug registry. It has also gone well beyond merely collecting and retrieving records, developing machine-learning algorithms to generate “data insights” and indicating that it taps into huge reservoirs of data outside state drug registries to arrive at them.

NarxCare—the system that inspired Kathryn’s gynecologist to part ways with her—is Appriss’ flagship product for doctors, pharmacies, and hospitals: an “analytics tool and care management platform” that purports to instantly and automatically identify a patient’s risk of misusing opioids.

On the most basic level, when a doctor queries NarxCare about someone like Kathryn, the software mines state registries for red flags indicating that she has engaged in “drug shopping” behavior: It notes the number of pharmacies a patient has visited, the distances she’s traveled to receive health care, and the combinations of prescriptions she receives.

Beyond that, things get a little mysterious. NarxCare also offers states access to a complex machine-learning product that automatically assigns each patient a unique, comprehensive Overdose Risk Score. Only Appriss knows exactly how this score is derived, but according to the company’s promotional material, its predictive model not only draws from state drug registry data, but “may include medical claims data, electronic health records, EMS data, and criminal justice data.” At

least eight states, including Texas, Florida, Ohio, and Michigan—where Kathryn lives—have signed up to incorporate this algorithm into their monitoring programs.

For all the seeming complexity of these inputs, what doctors see on their screen when they call up a patient's NarxCare report is very simple: a bunch of data visualizations that describe the person's prescription history, topped by a handful of three-digit scores that neatly purport to sum up the patient's risk.

Appriss is adamant that a NarxCare score is not meant to supplant a doctor's diagnosis. But physicians ignore these numbers at their peril. Nearly every state now uses Appriss software to manage its prescription drug monitoring programs, and most legally require physicians and pharmacists to consult them when prescribing controlled substances, on penalty of losing their license. In some states, police and federal law enforcement officers can also access this highly sensitive medical information—in many cases without a warrant—to prosecute both doctors and patients.

In essence, Kathryn found, nearly all Americans have the equivalent of a secret credit score that rates the risk of prescribing controlled substances to them. And doctors have authorities looking over their shoulders as they weigh their own responses to those scores.

Even after Kathryn had read up on NarxCare, however, she was still left with a basic question: Why had she been flagged with such a high score? She wasn't "doctor shopping." The only other physician she saw was her psychiatrist. She did have a prescription for a benzodiazepine to treat post-traumatic stress disorder, and combining such drugs with opioids is a known risk factor for overdose. But could that really have been enough to get her kicked out of a medical practice?

As Kathryn continued her research online, she found that there was a whole world of chronic pain patients on Twitter and other forums comparing notes on how they'd run afoul of NarxCare or other screening tools. And eventually she came upon an explanation that helped her understand what might have gone wrong: She had sick pets.

At the time of her hospitalization, Kathryn owned two flat-coated retrievers, Bear and Moose. Both were the kind of dog she preferred to adopt: older rescues with significant medical problems that other prospective owners might avoid. Moose had epilepsy and had required surgery on both his hind legs. He had also been abused as a puppy and had severe anxiety. Bear, too, suffered from anxiety.

The two canines had been prescribed opioids, benzodiazepines, and even barbiturates by their veterinarians. Prescriptions for animals are put under their owner's name. So to NarxCare, it apparently looked like Kathryn was seeing many doctors for different drugs, some at extremely high dosages. (Dogs can require large amounts of benzodiazepines due to metabolic factors.) Appriss says that it is “very rare” for pets’ prescriptions to drive up a patient’s NarxCare scores.

As Kafkaesque as this problem might seem, critics say it's hardly an isolated glitch. A growing number of researchers believe that NarxCare and other screening tools like it are profoundly flawed. According to one study, 20 percent of the patients who are most likely to be flagged as doctor-shoppers actually have cancer, which often requires seeing multiple specialists. And many of the official red flags that increase a person's risk scores are simply attributes of the most vulnerable and medically complex patients, sometimes causing those groups to be denied opioid pain treatment.

The AI that generates NarxCare’s Overdose Risk Score is, to many critics, even more unsettling. At a time of mounting concern over predictive algorithms, Appriss’ own descriptions of NarxCare—which boast of extremely wide-ranging access to sensitive patient data—have raised alarms among patient advocates and researchers. NarxCare’s home page, for instance, describes how its algorithm trawls patient medical records for diagnoses of depression and post-traumatic stress disorder, treating these as “variables that could impact risk assessment.” In turn, academics have published hundreds of pages about NarxCare, exploring how such use of diagnostic records could have a disparate impact on women (who are more likely to suffer trauma from abuse) and how its purported use of criminal justice data could skew against racial minorities (who are more likely to have been arrested).

But the most troubling thing, according to researchers, is simply how opaque and unaccountable these quasi-medical tools are. None of the algorithms that are widely used to guide physicians’ clinical decisions—including NarxCare—have been validated as safe and effective by peer-reviewed research. And because Appriss’ risk assessment algorithms are proprietary, there's no way to look under the hood to inspect them for errors or biases.

Nor, for that matter, are there clear ways for a patient to seek redress. As soon as Kathryn realized what had happened, she started trying to clear her record. She’s still at it. In the meantime, when she visits a pharmacy or a doctor’s office, she says she can always tell when someone has seen her score. “Their whole demeanor has changed,” she says. “It reminds me of a suspect and a detective. It’s no longer a caring, empathetic, and compassionate relationship. It’s more of an inquisition.”

The United States' relationship with opioid drugs has always been fraught. We either love them or we hate them. Historically, periods of widespread availability spur addictions, which lead to crackdowns, which lead to undertreatment of pain—and then another extreme swing of the pendulum, which never seems to settle at a happy medium.

The current anti-opioid climate has its roots in the overmarketing of Purdue Pharma's OxyContin in the mid-1990s. Between 1999 and 2010, opioid prescribing in the US quadrupled—and overdose deaths rose in tandem. To many experts, this suggested an easy fix: If you decrease prescribing, then death rates will decline too.

But that didn't happen. While the total amount of opioids prescribed fell by 60 percent between 2011 and 2020, the already record-level overdose death rate at least doubled during the same period. Simply cutting the medical supply didn't help; instead, it fueled more dangerous drug use, driving many Americans to substances like illegally manufactured fentanyl.

The reason these cuts hadn't worked, some experts believed, was that they had failed to target the patients at highest risk. Around 70 percent of adults have taken medical opioids—yet only 0.5 percent suffer from what is officially labeled “opioid use disorder,” more commonly called addiction. One study found that even within the age group at highest risk, teenagers and people in their early twenties, only one out of every 314 privately insured patients who had been prescribed opioids developed problems with them.

Researchers had known for years that some patients were at higher risk for addiction than others. Studies have shown, for instance, that the more adverse childhood experiences someone has had—like being abused or neglected or losing a parent—the greater their risk. Another big risk factor is mental illness, which affects at least 64 percent of all people with opioid use disorder. But while experts were aware of these hazards, they had no good way to quantify them.

That began to change as the opioid epidemic escalated and demand grew for a simple tool that could more accurately predict a patient's risk. One of the first of these measures, the Opioid Risk Tool (ORT), was published in 2005 by Lynn Webster, a former president of the American Academy of Pain Medicine, who now works in the pharmaceutical industry. (Webster has also previously received speaking fees from opioid manufacturers.)

To build the ORT, Webster began by searching for studies that quantified specific risk factors. Along with the literature on adverse childhood experiences, Webster found studies linking risk to both personal and family history of addiction—not just to opioids but to other drugs, including alcohol. He also found

data on elevated risk from particular psychiatric disorders, including obsessive-compulsive disorder, bipolar disorder, schizophrenia, and major depression.

Gathering all this research together, Webster designed a short patient questionnaire meant to suss out whether someone possessed any of the known risk factors for addiction. Then he came up with a way of summing and weighting the answers to generate an overall score.

The ORT, however, was sometimes sharply skewed and limited by its data sources. For instance, Webster found a study showing that a history of sexual abuse in girls tripled their risk of addiction, so he duly included a question asking whether patients had experienced sexual abuse and codified it as a risk factor—for females. Why only them? Because no analogous study had been done on boys. The gender bias that this introduced into the ORT was especially odd given that two-thirds of all addictions occur in men.

The ORT also didn't take into account whether a patient had been prescribed opioids for long periods without becoming addicted.

Webster says he did not intend for his tool to be used to deny pain treatment—only to determine who should be watched more closely. As one of the first screeners available, however, it rapidly caught on with doctors and hospitals keen to stay on the right side of the opioid crisis. Today, it has been incorporated into multiple electronic health record systems, and it is often relied on by physicians anxious about overprescription. It's "very, very broadly used in the US and five other countries," Webster says.

In comparison to early opioid risk screeners like the ORT, NarxCare is more complex, more powerful, more rooted in law enforcement, and far less transparent.

Appriss started out in the 1990s making software that automatically notifies crime victims and other "concerned citizens" when a specific incarcerated person is about to be released. Later it moved into health care. After developing a series of databases for monitoring prescriptions, Appriss in 2014 acquired what was then the most commonly used algorithm for predicting who was most at risk for "misuse of controlled substances," a program developed by the National Association of Boards of Pharmacy, and began to develop and expand it. Like many companies that supply software to track and predict opioid addiction, Appriss is largely funded, either directly or indirectly, by the Department of Justice.

NarxCare is one of many predictive algorithms that have proliferated across several domains of life in recent years. In medical settings, algorithms have been used to predict which patients are most likely to benefit from a particular treatment

and to estimate the probability that a patient in the ICU will deteriorate or die if discharged.

In theory, creating such a tool to guide when and to whom opioids are prescribed could be helpful, possibly even to address medical inequities. Studies have shown, for instance, that Black patients are more likely to be denied medication for pain, and more likely to be perceived as drug-seeking. A more objective predictor could—again, in theory—help patients who are undermedicated get the treatment they need.

But in practice, algorithms that originate with law enforcement have displayed a track record of running in the opposite direction. In 2016, for example, ProPublica analyzed how COMPAS, an algorithm designed to help courts identify which defendants are most likely to commit future crimes, was far more prone to incorrectly flag Black defendants as likely recidivists. (The company that makes the algorithm disputed this analysis.) In the years since then, the problem of algorithmic unfairness—the tendency of AI to obscure and weaponize the biases of its underlying data—has become an increasingly towering concern among people who study the ethics of AI.

Over the past couple of years, Jennifer Oliva, director of the Center for Health and Pharmaceutical Law at Seton Hall University, has set out to examine NarxCare in light of these apprehensions. In a major recent paper called “Dosing Discrimination,” she argues that much of the data NarxCare claims to trace may simply recapitulate inequalities associated with race, class, and gender. Living in a rural area, for example, often requires traveling longer distances for treatment—but that doesn’t automatically signify doctor shopping. Similarly, while it’s a mystery exactly how NarxCare may incorporate criminal justice data into its algorithm, it’s clear that Black people are arrested far more often than whites. That doesn’t mean that prescribing to them is riskier, Oliva says—just that they get targeted more by biased systems. “All of that stuff just reinforces this historical discrimination,” Oliva says.

Appriss, for its part, says that within NarxCare’s algorithms, “there are no adjustments to the risk scoring to account for potential underlying biases” in its source data.

Other communications from the company, however, indicate that NarxCare’s underlying source data may not be what it seems.

Early in the reporting of this piece, Appriss declined WIRED’s request for an interview. Later, in an emailed response to specific questions about its data sources, the company made a startling claim: In apparent contradiction to its own marketing

material, Appriss said that NarxCare’s predictive risk algorithm makes no use of any data outside of state prescription drug registries. “The Overdose Risk Score was originally developed to allow for ingestion of additional data sources beyond the PDMP,” a spokesperson for the company said, “but no states have chosen to do so. All scores contained within NarxCare are based solely on data from the prescription drug monitoring program.”

Some states do incorporate certain criminal justice data—for instance, drug conviction records—into their prescription drug monitoring programs, so it’s conceivable that NarxCare’s machine-learning model does draw on those. But Appriss specifically distanced itself from other data sources claimed in its marketing material.

For instance, the company told WIRED that NarxCare and its scores “do not include any diagnosis information” from patient medical records. That would seem to suggest, contra the NarxCare homepage, that the algorithm in fact gives no consideration to people’s histories of depression and PTSD. The company also said that it does not take into account the distance that a patient travels to receive medical care—despite a chatty 2018 blog post, still up on the Appriss site, that includes this line in a description of NarxCare’s machine-learning model: “We might give it other types of data that involve distances between the doctor and the pharmacist and the patient’s home.”

These latest claims from Appriss only heighten Oliva’s concerns about the inscrutability of NarxCare. “As I have said many times in my own research, the most terrifying thing about Appriss’ risk-scoring platform is the fact that its algorithms are proprietary, and as a result, there is no way to externally validate them,” says Oliva. “We ought to at least be able to believe what Appriss says on its own website and in its public-facing documents.”

Moreover, experts say, even the most simple, transparent aspects of algorithms like NarxCare—the tallying of red flags meant to signify “doctor-shopping” behavior—are deeply problematic, in that they’re liable to target patients with complex conditions. “The more vulnerable a patient is, the more serious the patient’s illness, the more complex their history, the more likely they are to wind up having multiple doctors and multiple pharmacies,” notes Stefan Kertesz, a professor of medicine and public health at the University of Alabama at Birmingham. “The algorithm is set up to convince clinicians that care of anybody with more serious illness represents the greatest possible liability. And in that way, it incentivizes the abandonment of patients who have the most serious problems.”

To take some of the heat off of these complex patients, Appriss says that its algorithm “focuses on rapid changes” in drug use and deemphasizes people who

have maintained multiple prescriptions at stable levels for a long time. But as ever, the company stresses that a NarxCare score is not meant to determine any patient's course of treatment—that only a doctor can do that.

Doctors, however, are also judged by algorithms—and can be prosecuted if they write more prescriptions than their peers, or prescribe to patients deemed high risk. “I think prescribers have gotten really scared. They are very fearful of being called out,” says Sarah Wakeman, the medical director of the Substance Use Disorder Initiative at Massachusetts General Hospital, an assistant professor of medicine at Harvard, and a doctor who regularly uses NarxCare herself. Research has found that some 43 percent of US medical clinics now refuse to see new patients who require opioids.

Doctors are also, Wakeman says, “just not really sure what the right thing to do is.” A couple of academic surveys have found that physicians appreciate prescription drug registries, as they truly want to be able to identify patients who are misusing opioids. But doctors have also said that some registries can take too much time to access and digest. NarxCare is partly a solution to that problem—it speeds everything up. It distills.

The result of all that speed, and all that fear, says Kertesz, is that patients who have chronic pain but do not have addictions can end up cut off from medication that could help them. In extreme cases, that can even drive some chronic pain sufferers to turn to more dangerous illegal supplies, or to suicide. Among patients with long-term opioid prescriptions, research shows that stopping those prescriptions without providing effective alternative care is associated with nearly triple the risk of overdose death.

“The problem that really infuses the NarxCare discussion is that the environment in which it is being used has an intense element of law enforcement, fear, and distrust of patients,” Kertesz says. “It’s added to an environment where physicians are deeply fearful for their future ability to maintain a profession, where society has taken a particularly vindictive turn against both physicians and patients. And where the company that develops this interesting tool is able to force it onto the screens of nearly every doctor in America.”

As Kathryn became more steeped in online communities of chronic pain patients, one of the people she came into contact with was a 44-year-old woman named Beverly Schechtman, who had been galvanized by her own bad experience with opioid risk screening. In 2017, Schechtman was hospitalized for kidney stones,

which can cause some of the worst pain known to humans. In her case, they were associated with Crohn's disease, a chronic inflammatory disease of the bowel.

Because Crohn's flare-ups by themselves can cause severe pain, Schechtman already had a prescription for oral opioids—but she went to the hospital that day in 2017 because she was so nauseated from the pain that she couldn't keep them or anything else down. Like Kathryn, she also took benzodiazepines for an anxiety disorder.

That combination—which is both popular with drug users and considered a risk factor for overdose—made the hospitalist in charge of Schechtman's care suspicious. Without even introducing himself, he demanded to know why she was on the medications. So she explained that she had PTSD, expecting that this disclosure would be sufficient. Nonetheless, he pressed her about the cause of the trauma, so she revealed that she'd been sexually abused as a child.

After that, Schechtman says, the doctor became even more abrupt. “Due to that I cannot give you any type of IV pain medication,” she recalls him saying. When she asked why, she says he claimed that both IV drug use and child sexual abuse change the brain. “You'll thank me someday, because due to what you went through as a child, you have a much higher risk of becoming an addict, and I cannot participate in that,” she says she was told.

Schechtman says she felt that the doctor was blaming her for being abused. She was also puzzled.

She had been taking opioids on and off for 20-odd years and had never become addicted. Wasn't that relevant? And how could it be ethical to deny pain relief based on a theoretical risk linked to being abused? She wasn't asking for drugs to take home; she just wanted to be treated in the hospital, as she had been previously, without issue.

As would later happen for Kathryn, the experience drove Schechtman onto the internet. “I just became obsessed with researching all of it,” Schechtman says. “I was asking people in these online groups, ‘Have any of you been denied opioids due to sexual abuse history?’ And women were coming forward.”

Schechtman eventually joined an advocacy group called the Don't Punish Pain Rally. Together with other activists in the group, she discovered that the question about sexual abuse history in the ORT unfairly targeted women, but not men. (An updated version of Webster's tool now excludes the gender difference, but the older one seems to live on in some electronic medical record systems.)

She also found many pain patients who said they had problems with NarxCare. Bizarrely, even people who are receiving the gold standard treatment for addiction can be incorrectly flagged by NarxCare and then denied that very treatment by pharmacists.

Buprenorphine, best known under the brand name Suboxone, is one of just two drugs that are proven to cut the death rate from opioid use disorder by 50 percent or more, mainly by preventing overdose. But because it is an opioid itself, buprenorphine is among the substances that can elevate one's NarxCare score—though typically it is listed in a separate section of a NarxCare report to indicate that the person is undergoing treatment. That separation, however, doesn't necessarily prevent a pharmacist from looking at a patient's high score and refusing to offer them prescriptions.

Ryan Ward, a Florida-based recovery advocate, has taken buprenorphine for nearly a decade. He also has a history of severe back pain and related surgeries. In 2018, when his pharmacy stopped carrying buprenorphine, he tried to fill his prescription at a Walmart and was turned away. Then he visited two CVS's and three Walgreens, and was similarly stymied.

"I dress nicely. I look nice. And I would be friendly," he says. "And as soon as they get my driver's license, oh boy, they would change attitudes. I couldn't figure out why."

After panicking that he might plunge into withdrawal—and, ironically, be put at much higher risk of overdose—he changed tactics. He approached a pharmacist at a Publix, first showing her his LinkedIn page, which highlights his advocacy and employment. He described what had happened at the other drugstores.

When she checked the database, she immediately saw the problem: an overwhelmingly high Overdose Risk Score. Unlike her colleagues, however, she agreed to fill the prescription, realizing that it was nonsensical to deny a patient a medication that prevents overdose in the name of preventing overdose. Still, even three years later, if he tries another pharmacy he gets rejected.

Appriss stresses that its data is not supposed to be used in these ways. "Pharmacists and physicians use these scores as indicators or calls-to-action to further review details in the patient's prescription history in conjunction with other relevant patient health information," the company wrote in a statement. "The analysis and associated scores are not intended to work as sole determinants of a patient's risk." Appriss also says that prescriptions for buprenorphine have increased in areas of the country that use NarxCare.

But like the others, Ward has been unable to get his problem fixed. And since most states now require that physicians and pharmacists use these databases, millions are potentially affected. One survey of patients whose providers have checked these systems found that at least half reported being humiliated and 43 percent reported cuts in prescribing that increased pain and reduced quality of life.

Appriss says on its website that it's up to each state to deal with patient complaints. Still, few know where to turn. "The states have made it very difficult," says Oliva. Some don't even allow for error correction. And when Ward tried contacting Appriss directly, he says, he was ignored.

In the early 2010s, Angela Kilby was seeking a topic for her PhD thesis in economics at MIT. When a member of her family, a doctor in the rural South, told her how tough it was to make decisions about prescribing opioids in a community devastated by overdoses, Kilby felt she had found her subject. She decided to study the doctor's dilemma by examining how increased control over opioid prescribing actually affected patients. To track health outcomes, she used insurance claim data from 38 states that had implemented prescription monitoring databases at varying times between 2004 and 2014.

Going into her study, Kilby had been swayed by research and press reports—plentiful in an era of "pill mill" crackdowns and backlash against overprescribing—suggesting that opioids are not only addictive but also ineffective and even harmful for patients with chronic pain. She had predicted that reductions in prescribing would increase productivity and health. "I was expecting to see the opposite of what I saw," she says.

In fact, her research showed that cutting back on medical opioid prescriptions led to increased medical spending, higher levels of pain in hospitalized patients, and more missed workdays. "These are people who are probably losing access to opioids, who are struggling more to return to work after injuries and struggling to get pain treatment," she says.

Intrigued, she wanted to know more. So in the late 2010s, having become an assistant professor at Northeastern University, she decided to simulate the machine-learning model that generates NarxCare's most algorithmically sophisticated measure, the Overdose Risk Score.

Although Appriss did not make public the factors that went into its algorithm, Kilby reverse engineered what she could. Lacking access to prescription drug registry data, Kilby decided to use de-identified health insurance claims data, a source that underlies all of the other published machine-learning algorithms that predict opioid risk. Using roughly the same method that Appriss lays out in

accounts of its own machine-learning work, she trained her model by showing it cases of people who'd been diagnosed with opioid use disorder after receiving an opioid prescription. She sent it looking for resemblances and risk predictors in their files. Then she turned her model loose on a much larger sample, this time with those opioid-use-disorder diagnoses hidden from the algorithm, to see if it actually identified real cases.

What Kilby found was that while NarxCare's model may trawl a different data set, it almost certainly shares an essential limitation with her algorithm.

"The problem with all of these algorithms, including the one I developed," Kilby says, "is precision." Kilby's complete data set included the files of roughly 7 million people who were insured by their employers between 2005 and 2012. But because opioid addiction is so rare in the general population, the training sample that the algorithm could use to make predictions was small: some 23,000 out of all those millions.

Further, 56 percent of that group had addictions before they received their first prescription, meaning that the medication could not have caused the problem—so they had to be excluded from the training sample. (This supports other data showing that most people with opioid addiction start with recreational, rather than medical, use.)

The result was that Kilby's algorithm generated a large number of both false positive and false negative results, even when she set her parameters so strictly that someone had to score at or above the 99th percentile to be considered high risk. In that case, she found, only 11 percent of high scorers had actually been diagnosed with opioid use disorder—while 89 percent were incorrectly flagged.

Loosening her criteria didn't improve matters. Using the 95th percentile as a cutoff identified more true positives, but also increased false ones: This time less than 5 percent of positives were true positives. (In its own literature, Appriss mentions these two cutoffs as being clinically useful.)

Kilby's research also identified an even more fundamental problem. Algorithms like hers tend to flag people who've accumulated a long list of risk factors in the course of a lifetime—even if they've taken opioids for years with no reported problems. Conversely, if the algorithm has little data on someone, it's likely to label them low risk. But that person may actually be at higher risk than the long-term chronic pain patients who now get dinged most often.

"There is just no correlation whatsoever between the likelihood of being said to be high risk by the algorithm and the reduction in the probability of developing

opioid use disorder,” Kilby explains. In other words, the algorithm essentially cannot do what it claims to do, which is determine whether writing or denying someone’s next prescription will alter their trajectory in terms of addiction. And this flaw,

In her paper “Dosing Discrimination,” about algorithms like NarxCare, Jennifer Oliva describes a number of cases similar to Kathryn’s and Schectman’s, in which people have been denied opioids due to sexual trauma histories and other potentially misleading factors. The paper culminates in an argument that FDA approval—which is currently not required for NarxCare—should be mandatory, especially given Appriss’ dominance of the market.

The larger question, of course, is whether algorithms should be used to determine addiction risk at all. When I spoke with Elaine Nsoesie, a data science faculty fellow at Boston University with a PhD in computational epidemiology, she argued that improving public health requires understanding the causes of a problem—not using proxy measures that may or may not be associated with risk.

“I would not be thinking about algorithms,” she says. “I would go out into the population to try to understand, why do we have these problems in the first place? Why do we have opioid overdose? Why do we have addictions? What are the factors that are contributing to these problems and how can we address them?”

In contrast, throughout the overdose crisis, policymakers have focused relentlessly on reducing medical opioid use. And by that metric, they’ve been overwhelmingly successful: Prescribing has been more than halved. And yet 2020 saw the largest number of US overdose deaths—93,000—on record, a stunning 29 percent increase from the year before.

Moreover, even among people with known addiction, there is little evidence that avoiding appropriate medical opioid use will, by itself, protect them. “I think undertreated pain in someone with a history of addiction is every bit, if not more, of a risk factor for relapse,” says Wakeman. She calls for better monitoring and support, not obligatory opioid denial.

Appriss has recognized the need to study NarxCare’s effects on the health and mortality of people flagged by the system—and not just whether it results in reduced prescribing. At a recent webinar, the company’s manager of data science, Kristine Whalen, highlighted new data showing that implementation of NarxCare sped up the decline in opioid prescribing in six states by about 10 percent, compared to reductions before it was used. When asked whether the company was also measuring NarxCare’s real-world effects on patients’ lives, Whalen said, “We’re

actively looking for additional outcome data sets to be able to do what you are describing.”

For Kathryn, at least, NarxCare’s effect on her life and health has been pretty stark. Aside from her psychiatrist, she says, “I don’t have a doctor because of this NarxCare score.” She worries about what she’ll do the next time her endometriosis flares up or another emergency arises, and she still struggles to get medication to treat her pain.

And it’s not only Kathryn’s own pain prescriptions that require filling. Although her dog Moose died in late 2020, Bear continues to need his meds, and Kathryn has since gone on to adopt another medically demanding dog, Mouse. Some states have recognized the problem of misidentified veterinary prescriptions and require NarxCare to mark them with a paw print or animal icon on health providers’ screens. Apparently, though, those prescriptions can still influence the pet owner’s overall scores—and the next busy pharmacist who peers warily at a computer screen.