

Why the FDA Is Embracing Old Math for New Drugs

In the 20th-century statistics wars, Bayesians were underdogs. Now their methods may help speed treatments to market.

March 13, 2026 By Undark and Sara Talpos

Clinical trials for a new drug can take years to complete, and cost up to hundreds of millions of dollars. New [draft guidance](#) from the U.S. Food and Drug Administration aims to make that process faster and cheaper for some studies, by encouraging a tool called Bayesian statistics. The approach dates back more than 250 years, and proponents say its embrace by regulators is overdue, stalled at first by feuding camps of statisticians, then later by a lack of familiarity among trained professionals.

For decades, traditional statistical methods dominated the graduate school curriculum, and today only a small proportion of working statisticians have significant Bayesian training, said Frank Harrell, Jr., a professor of biostatistics at Vanderbilt University School of Medicine who provided input on the guidance in his capacity as an expert adviser to the FDA. Consequently, he continued, “there is a general resistance to change.”

But change is coming nevertheless. In 2022, as part of the [Prescription Drug User Fee Act VII](#), the FDA made a commitment to the pharmaceutical industry to provide a guidance document on Bayesian methods. While this approach was never explicitly prohibited, the industry wanted “to have some consistency, know what to expect if they’re preparing a Bayesian proposal,” said John Scott, a top official in the agency’s Center for Biologics Evaluation and Research.

The proposal has many fans outside the pharmaceutical industry, and Undark spoke with several academic experts in clinical trial design who welcomed the change, though one was decidedly less enthusiastic. Sander Greenland, an emeritus professor of epidemiology and statistics at the University of California, Los Angeles, said he worries the FDA may be handing researchers an opportunity to massage their data in favor of a particular outcome. (“Bayesian statistics is wonderful, until other people start doing it,” reads a slide from a talk he gave to the Royal Statistical Society in London.)

His concern — and others’ enthusiasm — stems from the use of something called a “prior.” When researchers use Bayesian methods for a trial, they can take external information about a therapy, such as results from a previous study, and feed it into the trial’s analysis. In theory, the use of this

prior information makes trials more efficient and intuitive, but experts say it needs to be handled with care and not to unduly sway results.

Bayesian statistics take their name from an English minister, Thomas Bayes. Born in the early 1700s, he belonged to a [nonconformist](#) church, not the Church of England, and therefore would have been banned from the country's most prestigious universities. Bayes attended the University of Edinburgh and eventually became a fellow of the [Royal Society](#).

Bayes' now famous work, "[An Essay toward Solving a Problem in the Doctrine of Chances](#)," was discovered by a friend shortly after his death and published in 1763. It offered a mathematical formula for combining prior data with new data in order to come up with the probability of a hypothesis being correct.

Although the formula is relatively simple, it "has been the source of an immense amount of controversy spanning four centuries now," wrote mathematician Aubrey Clayton in his 2021 [book](#), "Bernoulli's Fallacy: Statistical Illogic and the Crisis of Modern Science." The real challenge came in the form of a competing approach, known as "frequentism," which came to dominate in the 1900s.

Bayesian statistics is wonderful, until other people start doing it.

Here's what frequentism looks like when applied to a clinical trial: Researchers start with a question. For example, does the new drug reduce diabetes symptoms? To ensure that their answer is not just due to randomness or wishful thinking, the research team creates a null hypothesis, something they want to disprove. Typically, they assume "the treatment is ignorable": It doesn't make patient better or worse, said Harrell. Then they look at the data and see if they're surprising, given that assumption.

So, if people receiving the new drug experience fewer diabetes symptoms than those receiving a placebo — and the analysis shows, with 95 percent confidence, that the disparity would not occur with a drug that does nothing — then it's generally treated as an effective therapy.

It's a "really strange tautological way of thinking," said Anna Heath, a statistician who works at the Hospital for Sick Children in Toronto. "I want to show something, so I'm going to pretend the opposite is true."

Nevertheless, this approach came to dominate statistics of the 1900s, in part thanks to a book by the legendary British statistician R. A. Fisher, called "[Statistical Methods for Research Workers](#)," published in 1925. The ideas contained in the book triggered "a complete revolution in the statistical methods employed in scientific research," [wrote](#) one author 25 years later. "There is no field of statistics," he continued, "in which the influence of Fisherian ideas is not profoundly felt."

The book went through 14 editions in a 45-year period, writes Clayton in "Bernoulli's Fallacy," "and it became such the industry standard that anyone not following one of Fisher's recipes would have a hard time getting results published."

Although Bayesians were the underdogs, they were not unheard, and debates raged in the academic literature, said Pavlos Msaouel, an oncologist and researcher at the University of Texas MD Anderson Cancer Center: "The polarization that you see right now in the political space in the U.S. is mild compared to how much they hated each other — the Bayesians and the frequentists of the 20th century — and how much they would openly say it in the academic journals."

But around the turn of the century, things started to calm down. Advances in computing helped make it practical to do the complex calculations required by Bayesian approaches. And a generation of hardcore frequentists started dying off. ("Science progresses one funeral at a time," Msaouel noted in an email to Undark.) People may have realized that there's room for both approaches, he said. This attitude reached regulators as well. In fact, the FDA published guidance for Bayesian methods in medical device trials [in 2010](#), but until now has held back on offering the same for drugs and biologics.

So how does Bayesian analysis work? Rather than having a hermetically sealed experiment, a Bayesian approach combines study data with a prior, which can capture external sources of information like past clinical trials. In this case, there's no null hypothesis. Researchers just answer the question they really care about, said Harrell, which is, "given the data, what's the chance the treatment works?"

Consider the hypothetical trial of the new diabetes drug: Researchers using a Bayesian approach might first look to the scientific literature for relevant high-quality data that could be fed into the prior. They'd then run the trial and combine the prior with the study data to come up with a probability that the treatment works.

"The polarization that you see right now in the political space in the U.S. is mild compared to how much they hated each other — the Bayesians and the frequentists of the 20th century."

Bayesian methods can be helpful in situations where trial data is hard to come by. For example, it can be tricky to do clinical trials with children because they tend to be healthier than adults, so there are fewer potential study participants. This often makes it less practical to run a trial with frequentist methods. In fact, [many pharmaceuticals](#) are only ever approved for use in adults, and then given to children off-label. Borrowing from adult data allows for greater confidence in a trial with fewer participants, said Heath, noting that "some pediatric data is better than no pediatric

data.”

A Bayesian approach can also be faster. In traditional trials, researchers must calculate at the outset the number of people needed in a trial, said Harrell. In most cases, researchers need to run through all the participants before they do their analysis, requiring a significant investment in time and money. In a Bayesian trial, on the other hand, researchers can analyze data every step of the way, sometimes allowing them to reach a determination sooner.

In an interview with Undark, the FDA’s John Scott said that Pfizer and BioNTech took a Bayesian approach when testing their Covid-19 vaccine in 2020. “They looked multiple times during the trial to see if there was early strong evidence of effectiveness,” he said. The early efficacy was very high, allowing it to become the first Covid-19 vaccine available in the U.S. (Scott noted that a frequentist approach could have been used in this case, too, but a Bayesian analysis was easier to interpret.)

Not everyone has embraced the guidance. Sander Greenland, who is among the world’s most cited medical statisticians, used to teach a workshop in Bayesian methods, including giving one at the FDA. But over time he has soured on the approach. The issue, he said, is with how the prior gets implemented. In theory, when there is little available previous evidence, researchers can choose a prior that doesn’t count too heavily, which is known as a “weakly informative” prior. But the field doesn’t have a standard definition of “weakly informative,” he said, and this has opened the door for researchers to stack the deck in favor of their preferred outcome.

Greenland said that he has reviewed studies where the prior selected is described as “weakly informative,” but actually weights the results quite heavily toward a specific outcome. A federal agency could insist on certain guardrails to discourage priors that unfairly bias results, but Greenland noted that the FDA has a track record of making exceptions to its own guidelines.

Others feel more confident that the FDA will be able to put effective guardrails in place. Msaouel said that, from reading the guidance, he thinks the regulators will be able to spot companies gaming the system. “These people know what they’re talking about,” he said. (Msaouel has overseen clinical trials funded by several pharmaceutical companies, and he has also received honoraria from pharmaceutical firms.)

In determining how prior information should be handled, “it’s entirely possible for even two different scientists working in the same area to have somewhat different opinions.”

In an email, Rachael Burden, an FDA communications adviser, wrote that the agency’s review “is the main guardrail” for making sure that decision to greenlight a drug is grounded in solid evidence. “The guidance emphasizes the need for high quality, relevant data and discusses steps necessary to ensure adequate transparency for verification during review,” she added.

Scott said that an increased use of Bayesian statistics may pose some logistical challenges. For instance, there is some judgement involved in determining how prior information should be handled. “It’s entirely possible for even two different scientists working in the same area to have

somewhat different opinions,” he noted. And because the FDA regulates a broad range of medical products, “different divisions may have different answers for what are structurally similar questions, but for different application areas.”

The guidance is [now open](#) for public comment until March 13.

Harrell said the guidance will undoubtedly be updated to reflect feedback. He’s anticipating a lot of excitement as well as pushback. “How that balances,” he said, “I’ll be on the edge of my seat to see that.”

This article was originally published on [Undark](#). Read the [original article](#).