U.S. Health Officials Move Quickly to Deploy Medical A.I. Despite Concerns

Federal projects are rolling out A.I. agents that offer therapy and prescribe medicine, raising concerns about safety and the influence of venture capital investors.

Sept. 14, 2026Updated 1:38 p.m. ET

A demonstration of a Limbic A.I. mental health app.Jeremie Souteyrat for The New York Times

The Trump administration is accelerating efforts to make artificial intelligence an integral part of medical care in the United States, throwing the resources and support of the federal government into projects that deploy A.I. agents to diagnose and prescribe treatments to patients.

The multifront effort at the Department of Health and Human Services has raised concerns among some officials over the past few months that change is moving too quickly, with inadequate evidence of the technology's safety and effectiveness, according to people who have been involved

the discussions. They also worry about outsized influence of Silicon Valley investors that have previously played little role in federal health policy.

Vinod Khosla, a billionaire Silicon Valley venture capitalist whose son has a health care A.I. company, has been particularly influential in conversations with top health department officials, according to people close to the matter. Those officials include Dr. Mehmet Oz, the head of Medicare and Medicaid. Chris Klomp, who faces a confirmation hearing this week to be the second-in-command at the health department, said on a podcast in May that the use of A.I. in patient care was the "holy grail" in the agency's quest to improve the nation's health and reduce costs.

Mr. Khosla has predicted for more than a decade that A.I. would largely replace doctors. "It's over for doctors, human doctors," he said at a July A.I. start-up event. "A.I. is just going to be better."

In an interview, Mr. Khosla trumpeted Curai Health, his son's company, which he has backed financially, saying it could offer an A.I. primary care provider that could give underserved people rapid access to care. "C.M.S. seems to be very excited," Mr. Khosla said, referring to the Centers for Medicare and Medicaid Services.

Concerns about the dangers of A.I., and calls to regulate it, have erupted in recent days as some top executives of A.I. labs have called for slowing the pace of development. In medicine, doctors are increasingly relying on A.I. for help with administrative tasks, but the prospect of A.I. taking over physician roles in assessing and treating patients is unnerving to some.

Vinod Khosla, a billionaire Silicon Valley venture capitalist whose son has a health care A.I. company, has been particularly influential in conversations with top health department officials, according to people close to the matter.Brendan Smialowski/Agence France-Presse — Getty Images

Models meant to render care independently to patients "are not ready for prime time," said Dr. John Whyte, chief executive of the American Medical Association, the nation's largest physicians organization.

He said the evidence that Mr. Khosla and others have cited as showing A.I.'s superiority to doctors — including in a recent essay in an American Medical Association publication whose authors included Mr. Khosla and his son, Neal — was not convincing and largely based on simulations.

"A.I. in health should not be driven by the tech industry," Dr. Whyte said. "It should be driven by the clinical community. And right now, what we have going on is that tech is the tail wagging the dog."

Grace Davis Jamison, a health department spokeswoman, said the department was taking a "responsible, evidence-based approach" to A.I. deployment and "expanding its use as the evidence demonstrates value."

"H.H.S. is embracing cutting-edge American innovation, including A.I. as a practical tool, to better serve patients, families, and clinicians as we work to Make America Healthy Again," she said in a statement.

As part of the effort, federal health officials are working on developing a new payment category for Medicare, the federal health program that covers more than 60 million older Americans, to reimburse companies for A.I. software that supports medical care or diagnosis. Medicare officials have urged private insurers covering about 165 million Americans to reimburse for "technology-supported care," including some A.I. projects, if they improve care and cut costs.

The F.D.A. has approved more than 1,500 medical devices that include A.I. to perform discrete tasks, such as spotting a tumor or a blood clot. But the agency does not have rules built to address newer models, such as chatbots or so-called agentic A.I. that perform tasks that a doctor would do, such as prescribing a drug. Ms. Jamison said the F.D.A. was examining how its rules should change to keep up with A.I. systems "that operate with greater autonomy."

Still, Medicare officials are moving an array of pilot programs forward without F.D.A. approval. Medicare has allowed more than 200 companies to try out pilots that can include A.I.

Through a program aligned with the Medicare effort, the F.D.A. has rejected many of those applicants but has allowed the use of four otherwise-unapproved A.I. models that meet safety criteria and appear likely to benefit patients. They include a project that offers talk therapy with an A.I. chatbot to Medicare beneficiaries with depression or anxiety.

The F.D.A. is still discussing how it should regulate complex A.I. models that change over time and perform actions like prescribing. The task is daunting: The agency would have to manage the risks of large-language models; one was already accused of giving "dangerous" medical advice to a man who delayed care for a life-threatening blood clot.

It would also have to contemplate whether A.I. doctors could go rogue in a manner akin to the OpenAI agents that coordinated a hacking campaign on a company called Hugging Face.

Robert F. Kennedy Jr., the health secretary, has publicly acknowledged the risks, but said during a congressional hearing in April that A.I. was "going to revolutionize medicine."

“A.I. is very dangerous potentially,” he said at the hearing. “But it also has the capacity to bring really great things to humanity, particularly in the realm of human health.”

One of his sons, Finn, started a venture capital fund this year that says it “builds and invests in generational health companies” and was reported to have solicited funds to invest in health care A.I.

Last year, Finn Kennedy worked for 8VC, a venture capital firm run by Joe Lonsdale, a billionaire ally of President Trump’s; the younger Kennedy coauthored articles at the firm, saying that A.I. doctors “won’t work for free,” suggesting that the companies that create the technology should be paid. He laid out the federal overhaul needed to deploy them, which included appointing a person at the F.D.A. to focus specifically on artificial intelligence.

Last week, the health department did just that, announcing that it had elevated Jared Seehafer, a former tech entrepreneur, to a newly created post of deputy commissioner for technology and artificial intelligence. Mr. Seehafer, a political appointee who has been at the agency about a year, has been described by people who worked with him as wanting to pare back the agency’s oversight of

medical A.I.

Mr. Seehafer's appointment drew notice in Silicon Valley investment circles. "Great choice from the H.H.S. team," Sebastian Caliri of 8VC, who co-wrote the A.I. articles with the younger Mr. Kennedy, posted on X. Mr. Lonsdale, his boss, whose venture capital firm has invested in medical A.I., wrote: "Excited to see the right talent going in."

Mr. Seehafer began working at the F.D.A. after he sold his start-up that provided software to help companies comply with F.D.A. rules. Soon he began to clash with Dr. Marty Makary, then the agency's commissioner, and with the director of medical devices, whose division oversees software that directs patient care.

Mr. Seehafer proposed a plan to centralize A.I. oversight at the health department or in the F.D.A. commissioner's office, raising concerns that decisions would be more influenced by the tech sector than by medical considerations, according to people familiar with the plan who were not authorized to discuss it publicly.

The plan met with resistance from Dr. Michelle Tarver, the top medical devices official overseeing A.I. software, favoring leadership from expert scientists with a focus on patient safety.

Dr. Makary, who had been outspoken about the influence of Silicon Valley investors, then got word that health department officials wanted him to fire Dr. Tarver, according to people with knowledge of the matter. He refused.

The F.D.A. said in a statement, “The work, expertise, and involvement of career staff is a vital component of F.D.A. successfully delivering public health for the American people.”

The statement added that Mr. Seehafer “has nothing but respect for Dr. Tarver” and her division’s “long history of facilitating innovation while ensuring the safety and effectiveness of medical devices.”

Within the administration, conversations about A.I. in medicine have continued, including on whether to pay AI doctors operated by tech companies and start-ups as much as 60 to 80 percent of what human doctors earn for the same service, according to a person involved.

Other projects offer payment based on whether they demonstrate benefits to patients and cost savings to Medicare. One of them is from the company Limbic, offering cognitive behavioral therapy to Medicare beneficiaries who were diagnosed with depression or anxiety.

The company, which was backed by the elder Mr. Khosla,

was selected by the F.D.A. to participate in a Medicare A.I. pilot.

In it, an A.I. counselor, which has a female voice and uses the name "Hope," offers an hourlong session each week with check-ins as needed. A study by company scientists and engineers that was published in the journal Nature Medicine found that clinicians who were unaware if they were reviewing a therapy transcript led by a human or a chatbot concluded that the Limbic A.I. outperformed humans and a general chatbot.

Ross Harper, Limbic's chief executive, said that one clinician would supervise thousands of A.I. therapists.

"It is a moral issue to get this out into the hands of people who can't access care," said the Limbic executive, who holds a doctorate in computational neuroscience. He added: "We have to remind ourselves: What are we comparing this against? We're comparing it against no treatment, a long wait list."

Another medical A.I. pilot with about $60 million in funding is moving forward at the health department's Advanced Research Projects Agency. Dr. Haider Warraich, a former F.D.A. official and cardiologist, said the Advocate project is meant to ultimately deploy autonomous A.I. to manage care for heart failure patients, giving them access to a chatbot

that can prescribe and manage medications.

Dr. Warraich said the project was inspired by his experience as a cardiologist managing heart failure patients during the Covid pandemic. He said his team had routine phone contact with patients, many of whom were afraid to enter the hospital and catch the virus, and found that patients improved with the regular check-ins.

He said the A.I. effort, put into practice by teams at Stanford and Duke, among other groups, would include remote monitoring or wearable sensors linked to the A.I., which could adjust medications if a patient's condition began to deteriorate. (Given current laws and rules, a doctor will oversee the A.I.)

The goal, he said, "is to really bring safe and effective clinical A.I. to the bedside, and build a regulatory pathway for this technology, so that we can really help unlock the potential of this technology — not just for heart failure, which is our main focus, but really for all conditions."

Other companies accepted into Medicare's pilot programs include Slingshot AI and Devoted Health, organizations backed by Andreessen Horowitz, a venture capital firm that has heavily supported Mr. Trump.

Neal Khosla's company, Curai Health, was selected by

Medicare and has an application pending with the F.D.A. to deploy its autonomous A.I. primary care doctors. In a pre-published study by company scientists and engineers, his team found that clinicians concurred with the AI doctor's diagnosis about 91 percent of the time across about 2,300 cases. However, in three cases, the A.I. suggested a telehealth visit in situations where a human doctor advised an urgent or emergency evaluation.

Given the promise to vastly expand affordable care, the younger Mr. Khosla said he was surprised by the vitriol generated by the recent editorial that he and his father and Dr. Ezekiel Emanuel, a health policy official in the Obama administration, wrote with Abe Baker-Butler in a JAMA journal. The article argued that autonomous A.I. outperformed human doctors in several areas and would be in widespread use by 2030.

Dr. Emanuel said that he got feedback from peers reminding him of the role doctors play in consoling patients. Overall, he said the response has been about "40 percent supportive."

Kitty Bennett contributed research.

Christina Jewett covers the Food and Drug Administration, which means keeping a close eye on drugs, medical devices, food safety and tobacco policy.