State of California
Office of the Attorney General

ROB BONTA
ATTORNEY GENERAL

February 27, 2026

Submitted via Federal eRulemaking Portal

The Honorable Robert F. Kennedy, Jr., Secretary
U.S. Department of Health and Human Services
200 Independence Avenue, SW, Washington, DC 20201

The Honorable Thomas Keane, M.D., M.B.A.
U.S. Department of Health and Human Services
Assistant Secretary for Technology Policy
Office of the National Coordinator for Health Information Technology
Mary E. Switzer Building, Mail Stop: 7033A
330 C Street SW, Washington, DC 20201

RE: "Health Data, Technology, and Interoperability: ASTP/ONC Deregulatory Actions To Unleash Prosperity," 90 Fed. Reg. 60,970 (RIN 0955-AA09)

Dear Secretary Kennedy and Assistant Secretary Keane:

I write to comment on the U.S. Department of Health and Human Services (HHS) and the Assistant Secretary for Technology Policy/Office of the National Coordinator for Health Information Technology's (ASTP/ONC) proposed rule entitled "Health Data, Technology, and Interoperability: ASTP/ONC Deregulatory Actions To Unleash Prosperity," 90 Fed. Reg. 60,970 (Dec. 29, 2025) (the "Proposed Rule").[1] On behalf of the California Department of Justice and State of California, I urge the federal government to reconsider its decision to remove the model card requirement from the certification program as it would result in a step backwards for accountability, transparency, and explainability in the rapidly evolving industry of health information technology.[2]

[1] Available at https://www.federalregister.gov/documents/2025/12/29/2025-23896/health-data-technology-and-interoperability-astponc-deregulatory-actions-to-unleash-prosperity.

[2] As ONC has described, *explainability* or *interpretability* refers to the ability for developers to explain in plain language how their data models are intended to be used. *See* Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing, 88 Fed. Reg. 23,746; 23,793 (Apr. 18, 2023) (to be codified at 45 C.F.R. pts. 170, 171).

California's comments are limited to the provisions of the Proposed Rule that eliminate transparency requirements in the use of decisionmaking tools outlined in Section III(A)(2)(e). 90 Fed. Reg. 60,986–87 (defined as "decision support interventions"). Transparency is the building block for accountability and serves the goal of delivering safe, effective, and equitable access to healthcare services. California supports the transparency elements of the prior HTI-1 Rule as they shine light on the validity of these decisionmaking tools and the data used to build predictive models. We believe it is of critical importance to ensure the use of decisionmaking tools continues to promote equitable access to healthcare and urge transparency that will assist in the regulation and enforcement of these emerging and complex technologies.

I. The Proposed Rule Seeks to Eliminate the Model Card Requirement for ONC Certification

ASTP/ONC launched the ONC Health IT Certification Program in 2010 to establish standards for electronic health records to meet certain markers for data accuracy, security, and interoperability. The program was created to make electronic health records more accessible for use in various federal, state, and private programs. As new technologies and needs arose, ASTP/ONC released updated versions of the certification criteria including the HTI-1 rule, which introduced requirements in response to the proliferation of automated decisionmaking tools trained on electronic health records.

With the Proposed Rule, ASTP/ONC seeks to eliminate 34 and revise seven of the existing 60 certification criteria.[3] Among those is the model card requirement established by HTI-1.[4] Currently, health IT developers must provide detailed source attribute information for "predictive decision support interventions"—including Artificial Intelligence or AI-based systems. These source attributes, referred to by ASTP/ONC as the "AI model card" requirements function like a nutrition label by including information such as the intervention's intended use, the target patient population and end users, known risks, the nature of the training data and inputs, and the processes used for external validation, among other specified details.[5] Model cards reveal whether an AI system has the potential to produce biased or incorrect conclusions by, for example, indicating whether a given tool was trained on data that reflect patients with

[3] Assistant Secretary for Technology Policy, HTI-5 Proposed Rule Chart, https://healthit.gov/resources/hti-5-proposed-rule-chart/ (last visited Feb. 23, 2026).

[4] Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing, 89 Fed. Reg, 1192 (Jan. 9, 2024), https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and; *see also* 45 C.F.R. § 170.315(b)(11).

[5] Office of the National Coordinator for Health Information Technology, Decision Support Interventions (DSI) Fact Sheet (2023), https://healthit.gov/wp-content/uploads/2023/12/HTI-1_DSI_fact-sheet_508.pdf.

different backgrounds, income levels, and geographic locations.[6]

The HTI-5 Proposed Rule would remove the requirement for developers to supply this source attribute information. This would eliminate one of the most significant guardrails currently in place on a federal level for the use of AI in healthcare.

II. Algorithmic Discrimination in Healthcare

The importance of maintaining existing federal transparency guardrails for healthcare AI is demonstrated by evidence of algorithmic discrimination occurring in the industry. Research reveals concerning examples in which the reality of algorithmic decisionmaking falls short of our expectations of fair and equitable delivery of healthcare services. In particular, algorithmic decisionmaking runs the risk of replicating and even amplifying human biases, particularly those affecting historically disadvantaged (and legally protected) groups.[7] Algorithmic discrimination occurs when these automated systems contribute to unjustified different treatment or impacts disfavoring people based on their race, color, ethnicity, sex (including pregnancy, childbirth, and related medical conditions, gender identity, intersex status, and sexual orientation), religion, age, national origin, disability, veteran status, genetic information, or any other classification protected by law.[8] Depending on the specific circumstances, algorithmic discrimination may violate federal and/or state laws, including civil rights laws.[9]

A groundbreaking study published by *Science* in 2019 revealed that a healthcare prediction algorithm used by hospitals and insurance companies throughout the country to identify patients in need of "high-risk care management" was far less likely to nominate Black patients.[10] The tool recommended White patients for high-risk care management programs, which provide

---

[6] Casey Ross, Trump administration to scrap federal rule requiring transparency into health AI tools, STAT (Dec. 22, 2025), https://www.statnews.com/2025/12/22/hhs-proposes-scrapping-ai-model-cards-transparency-rule/.

[7] *See, e.g.*, Executive Office of the President, Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights (2016), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf (last visited Feb. 23, 2026).

[8] The White House Office of Science and Technology Policy, Blueprint for an AI Bill of Rights 5 (2022), https://www.govinfo.gov/content/pkg/GOVPUB-PREX23-PURL-gpo193638/pdf/GOVPUB-PREX23-PURL-gpo193638.pdf.

[9] *See* Legal Advisory, State of California Department of Justice Office of the Attorney General, California Attorney General's Legal Advisory on the Application of Existing California Law to Artificial Intelligence in Healthcare (2025), https://oag.ca.gov/system/files/attachments/press-docs/Final%20Legal%20Advisory%20-%20Application%20of%20Existing%20CA%20Laws%20to%20Artificial%20Intelligence%20in%20Healthcare.pdf.

[10] *See* Ziad Obermeyer et al., Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations, Science, Oct. 25, 2019, at 447-53, https://www.science.org/doi/10.1126/science.aax2342; *see also* Casey Ross, Epic Overhauls Popular Sepsis Algorithm Criticized for Faulty Alarms, STAT (Oct. 3, 2022), https://www.statnews.com/2022/10/03/epic-sepsis-algorithm-revamp-training/; Andrew Wong et al., External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients, 181 JAMA Internal Med. 1065 (2021).

intensive—and often expensive—services to people with complex health needs.[11] The tool used higher healthcare costs as a proxy for medical need. Because Black patients are less likely to have access to care, even if they are insured, their healthcare costs tend to be lower. As a result, the algorithm systematically underestimated the health needs of Black patients and excluded them from high-risk care programs.[12]

Recent studies further illustrate how biases can be perpetuated or amplified by algorithms.[13] For example, a study led by Cedars-Sinai found a pattern of racial bias in treatment recommendations generated by leading artificial intelligence platforms for psychiatric patients.[14] Two large language models omitted medication recommendations for an attention-deficit/hyperactivity disorder case when race was explicitly stated, but they suggested them when those characteristics were missing from the case. Another large language model suggested guardianship for depression cases with explicit racial characteristics. The study concluded that the models showed racial bias because they reflected bias found in the extensive content used to train them.

In California, we have seen models integrated directly into existing electronic health record programs produce disappointing results. The University of California San Francisco reported bias in an algorithm used to identify potential appointment no-shows to facilitate double booking for that appointment.[15] In practice, the model resulted in low-resourced and marginalized populations being double-booked more often than others.[16] This seemingly innocuous program produced an output that reflected underlying structural inequalities and calls attention to bias in the labeling of data. These examples highlight how algorithmic tools that can create feedback loops that worsen discrimination.

The responsible use of data and technology by the healthcare industry has the potential to improve patient outcomes and increase efficiency. However, as evidenced by these examples, the

---

[11] Obermeyer, *supra* note 10; Carla K. Johnson, Racial Bias in Health Care Software Aids Whites Over Blacks, The Seattle Times, Oct. 25, 2019, https://www.seattletimes.com/seattle-news/health/racial-bias-in-health-care-software-aids-whites-over-blacks/.

[12] Obermeyer, *supra* note 10 (according to the author, the healthcare costs of unhealthier Black patients were on par with the costs of healthier White people, which meant they received lower risk scores, even when their needs were greater).

[13] *See, e.g.*, Kelley Tipton, M.P.H. et al., Agency for Healthcare Research and Quality, Impact of Healthcare Algorithms on Racial and Ethnic Disparities in Health and Healthcare, Dec. 2023, https://www.ncbi.nlm.nih.gov/sites/books/NBK598802/pdf/Bookshelf_NBK598802.pdf.

[14] Press Release, Cedars Sinai, Cedars-Sinai Study Shows Racial Bias in AI-Generated Treatment Regimens for Psychiatric Patients (Jun. 30, 2025), https://www.cedars-sinai.org/newsroom/cedars-sinai-study-shows-racial-bias-in-ai-generated-treatment-regimens-for-psychiatric-patients/.

[15] Brandie M. Nonnecke, et al., University of California Presidential Working Group on AI, Responsible Use of Artificial Intelligence: Recommendations to Guide the University of California's Artificial Intelligence Strategy, Univ. of California (Oct. 2021), https://www.ucop.edu/ethics-compliance-audit-services/compliance/uc-ai-working-group-final-report.pdf.

[16] *Id*.

use of predictive algorithmic models also presents risks that could lead to adverse outcomes, including bias, unfair discrimination, or other unconscionable impacts. As such, knowing how predictive algorithmic models are developed, trained, tested, and evaluated for bias and fairness is more crucial than ever.

III. California's Leadership Investigating Algorithmic Bias in Healthcare Reveals Hospitals Require More Transparency

As the primary enforcer of consumer laws in California, I have made the protection of Californians' access to high quality, affordable, and equitable healthcare a chief priority. In this regard, my office has a unique vantage point to evaluate the scope of issues presented by predictive modeling tools in the healthcare industry.

On August 31, 2022, my office initiated a first-of-its-kind inquiry focused on health equity, examining how healthcare providers address racial and ethnic disparities in their utilization of commercially available decisionmaking technologies.[17] Our initial analysis reinforced the risks associated with the use of algorithmic tools and the barriers to their effective and fair deployment. In California, these types of decisionmaking tools are now regularly used by hospitals to make judgments about patients across many contexts, ranging from medical treatments to managing revenue. Yet, many hospitals we queried reported three problematic issues. First, hospitals relied on analytical tools procured from third-party vendors and reported that they relied on the vendor's assessment that the tools they use are ethical and unbiased. Second, hospitals generally lacked insight into vendor's data modeling and reported that important information is often shrouded behind proprietary walls. Third, while some hospitals reported de-identifying data and/or removing race as a data coefficient, these same hospitals did not report any internal assessment as to whether altering data sets have an impact perpetuating bias. As multiple real-world examples have revealed, analytical decisionmaking tools can be disturbingly inaccurate, even when race is removed from data sets.[18]

One thing is clear; hospitals struggle to ensure that tools they procure from vendors can be used in a manner that is truly ethical and robust. Our findings inform our opposition to the Proposed Rule as addressed in detail below.

IV. ASTP/ONC's Justifications for Removing the Model Card Requirement Are Unsupported

To justify its removal of the model card requirement, ASTP/ONC cites to several Executive Orders. *See* 90 Fed. Reg. 60,987 (citing Executive Order 14,719 and Executive

---

[17] Press Release, State of California Department of Justice Office of the Attorney General, Attorney General Bonta Launches Inquiry into Racial and Ethnic Bias in Healthcare Algorithms (Aug. 31, 2022), https://oag.ca.gov/news/press-releases/attorney-general-bonta-launches-inquiry-racial-and-ethnic-bias-healthcare.

[18] Obermeyer, *supra* note 10.

Order 14,110). This alone does not justify the agency's actions. *See State v. Su*, 121 F.4th 1, 16 (9th Cir. 2024) (the Administrative Procedure Act (APA) does not allow presidential administrations to "issue agency regulations that evade APA-mandated accountability by simply issuing an executive order first"); *City of Fresno v. Turner*, 2025 WL 2721390 at *8 (N.D. Cal. Sept. 23, 2025) ("[T]he fact that an agency's actions were undertaken to fulfill a presidential directive does not exempt them from arbitrary-and-capricious review.") (quoting Kingdom v. Trump, No. 25-cv-691, 2025 WL 1568238, at *10 (D.D.C. June 3, 2025)).

In addition to the Executive Orders, ASTP/ONC notes four factors that support its proposal to eliminate the model card certification criteria: (1) that providers lack the time or technical background to meaningfully engage with model card information, (2) the anticipated savings and benefits of transparent predictive or generative AI application has not materialized, (3) the model card criteria is not required of other tech companies working in healthcare and has led to inconsistent oversight, and (4) removal of the model card requirement will promote innovation of AI in healthcare by reducing burden. 90 Fed. Reg. 60,987.

First, ASTP/ONC states that "most clinical users lack the time or technical background to engage with source attribute information" and that "no publicly available evidence indicating a single doctor, nurse, or administrator has accessed, recorded, or modified a single source attribute." 90 Fed. Reg. 60,987. This is contrary to the findings of our investigation, where we have spoken to providers who have affirmed the overall importance of getting understandable and comprehensive information from health IT vendors. And, as stated above, we have heard from many hospitals that they are struggling to ensure that tools they procure from health IT developers can be used in an ethical and robust way because they generally lack insight into developers' data modeling. Model cards are an important first step in understanding and evaluating that data.

Our findings have been corroborated by the Coalition for Health AI (CHAI), an organization whose public-private partnership includes both healthcare systems and tech start-ups committed to using artificial intelligence to improve clinical practice and foster innovation. In March 2025, CHAI collected over 150 responses from its members—including health systems, the tech industry, academia, and professional services—and reported that approximately 90% of respondents rated AI transparency as important, with particularly high support for public disclosures of AI model limitations and risks and real-world bias testing and demographic fairness evaluations.[19] Later in the year, CHAI deployed applied model cards in 36 health systems across the country including Kaiser Permanente, UMass Memorial, and University of Texas Health System. [20] In short, there is demand from providers for transparency in health IT products that use predictive algorithmic models.

---

[19] Coalition for Health AI, AI Action Plan: Mar. 2025 (2025), https://mcusercontent.com/14dfd11830f002586004651e1/files/82fe8466-566c-70e3-0727-b4fe6cc5a766/AI_Action_Plan_Insights_Final.pdf.

[20] Coalition for Health AI, Health AI "Nutrition Label" Advances, CHAI Blog (Mar. 3, 2025), https://www.chai.org/blog/health-ai-nutrition-label-advances.

To the extent that some providers have difficulty evaluating products using model cards, this warrants revisiting, rather than retreating from, the model card criteria to determine if they should be revised to be more impactful.[21] Model cards that contain inadequate, or even misleading, information would make it prohibitive for providers to meaningfully use them. This is not conjecture. Health IT developers have been in the news for misrepresenting the capabilities of their products. In a recent example, the Texas Attorney General Office settled with a healthcare AI developer that was making false claims about its products to providers.[22]

Second, ASTP/ONC claims that transparency may not have the savings and benefits the agency anticipated. ASTP/ONC cites to their regulatory impact analysis from the HTI-1 rulemaking, but cites no evidence aside from reiterating the claim that there is "no supporting evidence or published research indicating source attribute transparency requirements . . . have improved health care delivery organizations' evaluation and effective use of AI." 90 Fed. Reg. 60,987. As stated above, this does not merit abandoning the mission of transparency when there is so much support by providers for this type of information to enhance healthcare delivery.

Third, ASTP/ONC states that according to certified health IT developers, it is "unacceptable" that other tech companies producing similar, health-based AI systems are not subject to the same oversight. 90 Fed. Reg. 60,987. Although the certification program is voluntary, obtaining the accreditation from ONC confers a marker of trustworthiness that leads to more widespread adoption of the product. According to a recent ONC 2022 Report to Congress, nearly all non-federal acute care hospitals and about three-quarters of office-based physicians use health IT certified via the ONC Health IT Certification Program.[23] So even if ASTP/ONC does not regulate other tech companies working in this space, its authority to set certification standards still has implications for the entire industry. The prevalence and value of ONC-certified health IT in the industry shows the importance of this standard-setting to healthcare providers and other procurers. ONC should not reduce transparency for crucial health technology according to a "lowest common denominator" theory.

Fourth, ASTP/ONC asserts that in order to promote innovation in the health AI space, it must remove the transparency requirements established by HTI-1. 90 Fed. Reg. 60,987. The

---

[21] As we noted in our comment letter on HTI-1, there are limitations to increased transparency and certification of health IT, including the continuing obligations of users to ensure that their use of certified technology complies with all applicable laws, and difficulties ensuring transparency given the number of different actors and the complexity of the industry. *See* Letter from Rob Bonta, Attorney General of California, to Xavier Becerra, Secretary of the U.S. Department of Health and Human Services (Jun. 20, 2023), https://oag.ca.gov/system/files/attachments/press-docs/FINAL-%20Letter%20to%20Secretary%20Xavier%20Becerra%20and%20Dr.%20Tripathi%206-20-2023.pdf.

[22] Press Release, Attorney General of Texas, Attorney General Ken Paxton Reaches Settlement in First-of-its-Kind Healthcare Generative AI Investigation (Sept. 18, 2024), https://www.texasattorneygeneral.gov/news/releases/attorney-general-ken-paxton-reaches-settlement-first-its-kind-healthcare-generative-ai-investigation.

[23] Office of the National Coordinator for Health Information Technology, 2022 Report to Congress, 2022, https://healthit.gov/wp-content/uploads/2023/02/2022_ONC_Report_to_Congress.pdf.

Proposed Rule focuses on the perspective of developers without considering the perspective of providers and regulators. As a regulator of the healthcare industry in California, I can confirm that my office values the transparency that was afforded by the prior HTI-1 rule. As stated above, there is demand from providers for transparency in health IT products.

ASTP/ONC also does not take into consideration the significant burden it is placing on health providers by removing the model card requirement. For example, healthcare providers' compliance with Section 1557 of the Affordable Care Act becomes much more difficult without model card requirements. Section 1557 prohibits providers from discriminating based on a patient's protected status.[24] It also requires that providers make reasonable efforts to identify uses of patient care decision support tools in its health programs or activities that employ input variables or factors that measure race, color, national origin, sex, age, or disability.[25] The Proposed Rule's removal of the model card requirement would thus eliminate a critical tool for providers to ensure that they are providing nondiscriminatory healthcare in compliance with federal law.

The Proposed Rule will make it harder for healthcare providers to comply with state laws as well. I issued an AI advisory at the beginning of 2025 about the application of California law to AI in healthcare.[26] Examples of what could violate California law include:

- Denying health insurance claims using AI or other automated decisionmaking systems in a manner that overrides doctors' views about necessary treatment which carries a potential risk of harm to patients.
- Using generative AI or other automated decisionmaking tools to draft patient notes, communications, or medical orders that include erroneous or misleading information, including information based on stereotypes relating to race or other protected classification.
- Determining patient access to healthcare using AI or other automated decisionmaking systems that make predictions based on patients' past healthcare claims data, resulting in disadvantaged patients or groups that have a history of lack of access to healthcare being denied services on that basis while patients/groups with robust past access being provided enhanced services.
- Double-booking a patient's appointment, or creating other administrative barriers, because AI or other automated decisionmaking systems predict that patient is the "type of person" more likely to miss an appointment.
- Conducting cost/benefit analysis of medical treatments for patients with disabilities using AI or other automated decisionmaking systems that are based on stereotypes that undervalue the lives of people with disabilities.

[24] 45 C.F.R. § 92.210.
[25] *Id*.
[26] Legal Advisory, *supra* note 9.

As these examples illustrate, the transparency afforded by model cards is crucial to providers to understand how decisionmaking tools work and to decide whether they are safe to use for their patients.

Finally, ASTP/ONC's proposal to remove the model card certification requirement is short sighted because transparency can lead to more, not less, innovation. Model cards can help to simplify the process of sharing validation and testing results, as well as accelerate the development of models through trusted, independent entities.[27] Moreover, unbiased AI is more effective AI. AI that is trusted to deliver the results it promises without coming at the expense of patients' health can lead to more widespread adoption, thus serving dual purposes of promoting innovation while eliminating bias in our healthcare systems.

V. Conclusion

For the reasons stated in this letter, we urge the federal government to maintain the model card requirement the HTI-5 Proposed Rule seeks to eliminate.

Sincerely,

ROB BONTA
California Attorney General

[27] CHAI Blog, *supra* note 20.