



Elizabeth G. Taylor
Executive Director

Board of Directors

Ann Kappler
Chair
Prudential Financial, Inc.

William B. Schultz
Vice Chair
Zuckerman Spaeder LLP

Shamina Sneed
Secretary
TCW Group, Inc.

Nick Smirensky, CFA
Treasurer
New York State Health Foundation

L.D. Britt, MD, MPH
Eastern Virginia Medical School

Jeanna Cullins
Fiduciary & Governance Practice Leader
(Ret.)

Ian Heath Gershengorn
Jenner & Block

Michele Johnson
Tennessee Justice Center

Arian M. June
Debevoise & Plimpton LLP

Jane Preyer
Environmental Defense Fund (Ret.)

Lourdes A. Rivera
Pregnancy Justice

Donald B. Verrilli, Jr.
Munger, Tolles & Olson

Stephen Williams
Houston Health Department

Ronald L. Wisor, Jr.
Hogan Lovells

Senior Advisor to the Board
Rep. Henry A. Waxman
Waxman Strategies

General Counsel
Marc Fleischaker
Arent Fox, LLP

September 22, 2023

VIA ELECTRONIC SUBMISSION

Senator Bill Cassidy
Ranking Member
Health, Education, Labor & Pensions Committee
455 Dirksen Senate Office Building
Washington, D.C. 20510

RE: Exploring Congress' Framework for the Future of AI

Dear Ranking Member Cassidy,

The National Health Law Program (NHeLP) is a public interest law firm that fights for equitable access to quality health care for people with low incomes and underserved populations and for health equity for all. For over fifty years, we have litigated to enforce health care and civil rights laws, advocated for better federal and state health laws and policies, and trained, supported, and partnered with health and civil rights advocates across the country. We believe that all people should have access to the health care they need, regardless of geography, race, ethnicity, language, income, disability, sex, gender identity, sexual orientation, immigration status, or other factors.

Throughout much of our history, NHeLP [has fought against](#) faulty and harmful computer systems and processes that rely on automated decision making systems (ADS), including Artificial Intelligence (AI), such as assessment tools that unlawfully denied people benefits to which they are entitled. We have also partnered with [Upturn](#) and [Legal Aid of Arkansas](#) to form the [Benefits Tech Advocacy Hub](#) to give advocates tools to fight harmful benefits technology and force greater transparency so that harm to individuals can be identified, prevented, or reduced earlier in the technology's lifecycle. NHeLP has also released our [Principles for Fairer, More](#)

[Responsive Automated Decision-Making Systems](#), which reflect our years of work regarding ADS, including AI, and what features and protections are needed in responsible ADS. We have real-world experience fighting the harm caused by technology in public benefits and have practical experience about the how and why such harms occur. We are also experienced advocates regarding [data collection](#) to promote equity, while protecting the identity and private information of those from whom the data is collected.

NHeLP's decades of experience fighting to protect the rights of low-income and underserved people impacted by AI give us a different perspective on the policies that are needed. We understand what the systems look like on the ground and how they impact people, including protections that should be in place to minimize harm.

I. The AI Framework Needs to Broadly Define AI to Protect Against Harm.

While we welcome the Republican Committee members exploring Congress' AI framework, we urge you to expand your framework to include a broader definition of AI and directly address the issues most impacting public benefits programs. Most current AI efforts focus on machine learning and sophisticated generative AI, but any proposed legislation needs to include all technology used to make decisions or assist in making decisions, including rules driven technology. We have seen the harmful impact of different types of ADS, even those that are based on more traditional algorithmic programs. We also urge you to balance innovation against the harms of such systems and set requirements for transparency, effectiveness, and liability of AI.

II. Innovation Can Be Supported While Providing Appropriate Safety Net Protections for those Affected.

Efforts to foster innovation in AI must be balanced against its real and potential harms. Protection against harmful AI has long been required in public benefits programs. When such safeguards are fully enforced, they can support innovation by helping to promote awareness and identify issues, including discrimination and unintended impacts. AI has a long and ongoing history of harm in public benefits across all programs as seen in the Benefits Tech Advocacy Hub's [case study library](#). Importantly, public benefits are protected by constitutional due process protections that require people understand the decisions made about their services and provide them an opportunity to contest decisions. The connection of public benefits to fields in which AI is frequently used is vast. For example, decisions made in hospitals about health care for individuals who are covered by Medicaid are covered by constitutional due process. Similarly, these basic protections cover decisions made or recommended by AI in various other spheres in which a person is receiving services paid for through public benefits programs. It is critically important to recognize that most AI is based on



statistical analysis or other ways of predicting outcomes based on data inputs, so it is highly likely that the AI will not always make proper predictions about people whose individual circumstances do not neatly fit into predictive models.

The protections of constitutional due process provide a safety net to protect against the wrongful denial or termination of important services. Many AI accountability frameworks reflect the straightforward requirements of due process, but sometimes make them mere recommendations rather than requirements. Because of how deeply public benefits and AI are intertwined, any Congressional AI framework should use due process protections as the base level requirement. Therefore, safeguards are required that ensure people understand when a decision has used AI, the criteria the AI uses, the standard being applied, how to appeal the decision and request an exception, and any other information and procedures necessary to ensure that a person understand and can appeal a decision made by AI.

The harms of AI have been prevalent during the Medicaid continuous coverage unwinding. Due to the volume of redeterminations in unwinding, [CMS has recently become aware of a years-long issue](#) where many states' eligibility systems incorrectly conduct ex parte redeterminations at the household level instead of the individual level, resulting in many eligible individuals losing Medicaid coverage. This has especially impacted children who are eligible, often at higher income levels, and could be renewed ex parte, but whose parents or caregivers could not. When other household members are not able to be renewed and do not respond to a request for information, eligible children lose coverage. In some instances, this is due to incorrect state policy, but in others states were unaware the eligibility computer systems were functioning this way. Similarly, eligibility systems during the unwinding have shown systemic problems with not allowing people the minimum required time to respond, terminating people without evaluating them for other categories of benefits, or failing to evaluate for certain categories of benefits once an individual is found ineligible under one category. Such problems exemplify the equity issues, where marginalized populations such as children are disproportionately impacted, entwined in the real harms caused by AI use in benefits decision-making.

The harms from wrongful denial of services are significant and must be weighted heavily in cost-benefit analyses of AI innovation and implementation. Innovation cannot come at the cost of people in need of the services they are supposed to be receiving, and Congress can support implementation of new technology without sacrificing protection from harm. Primarily, Congress needs to ensure that any minimal protections recognize that due process protections apply to certain uses of AI and that the Congress's recommendations do not disturb those protections. In addition, the AI framework should require testing and protections for individuals—including ongoing monitoring of devices and tools to ensure they are accurate, reliable, and non-discriminatory—especially considering the faith people have in technology

being used for critical health care decisions. People want to trust computers and generally think they are more reliable than human judgment. To support such faith, technology must be constantly tested, reviewed, and adjusted with corrective action to remediate previous errors or harms.

III. Transparency and Accountability Must Be Required Throughout the AI Lifecycle.

The development of effective AI requires transparency and mechanisms for enforcement and liability throughout the [technology lifecycle](#). The potential for discrimination and harm can enter at any step of the process from how the question to be answered by AI is formulated to how the technology is used in practice. Inequity is often built into AI through the data sets used to train it. When training data is pulled from systems that have institutional biases, such biases become part of the AI's decision-making process and lead to inequitable and discriminatory results. The oft-cited [Obermeyer study](#) finding bias in tools that determine health care needs based on data measuring health care usage is one example of how certain fields, such as health care, need particular scrutiny when incorporating data or studies from the field into AI. Further, when certain populations are under-represented in training datasets, AI tools are not valid for use on these populations. For example, it is now [well-known](#) that the race-based corrective factor in analyzing eGFR to determine the presence of kidney disease leads to under diagnosis and reduced care for Black patients, including exclusion from kidney transplant qualification.

Those who use the technology, such as eligibility workers and clinicians, and those who the technology is used on, such as patients, need to understand how the technology is being used and its limitations. All parties need to understand how data and privacy are protected, and be provided with enough information about any decisions made such that they can challenge the decision. These protections are about more than safety standards. Protection from AI harms includes non-discrimination and meeting constitutional due process requirements, where applicable.

IV. Data and Privacy Concerns Must Be Prioritized.

Privacy and data accuracy concerns also deserve more attention. For example, Electronic Health Records (EHRs) are not always complete or fully accurate, but the information they contain can be used by AI tools in health care decision-making. EHRs often do not include critical information that may be in a provider's notes and they can lack inclusive selection options for demographic markers that impact care. For example, EHRs do not always have a full range of gender identity options, leading to improper care provided to gender diverse individuals. This can be seen when a person who identifies as male but needs access to

gynecological services is denied those services by AI because the EHR does not identify them as female. EHRs can even be a source of bias that leads to discrimination such as when a clinician guesses at a patient's race or ethnicity or inputs their own assumptions and judgments. EHR data introduces another place where human bias can influence an AI tool.

Further, discrimination must be considered in AI data privacy protections. These concerns go beyond HIPAA requirements. Individuals who seek care that is highly politicized or stigmatized, such as abortion, substance use disorder treatments, mental health care, and gender affirming care, need to know their health care data is protected. Otherwise, individuals may be deterred from seeking out care that is stigmatized or may experience retribution for seeking certain treatments.

V. Conclusion.

As mentioned previously, we are including as part of our comments NHeLP's [Principles for Fairer, More Responsive Automated Decision-Making Systems](#) because much of what we believe needs to be in an AI Framework is set forth in these principles. This document reflects much of our experience with ADS and protections that are necessary, including citations to other related NHeLP work product. While building an AI framework, individuals' rights cannot be forgotten, especially since many of the systems the Ranking Member's paper discusses are used to make decisions about public benefits.

Even with a strong framework and incorporation of our principles, incorrect decisions may not be completely preventable. Implementation of all AI must also include an easily accessible exceptions process that has clear guidelines to prevent and redress harms. Perfecting AI may not be possible, but protecting people is.

Thank you for your attention to our response and attached principles. If you have any questions, please reach out to Elizabeth Edwards at edwards@healthlaw.org or Cassandra LaRose at larose@healthlaw.org.

Sincerely,

Cassandra LaRose
Staff Attorney

Attachment: NHeLP, Principles for Fairer, More Responsive Automated Decision-Making Systems

