

Symposium

The Need for Prospective Integrity Standards for the Use of Generative AI in Research

Kayte Spector-Bagdady 

University of Michigan, Ann Arbor, MI, United States

Abstract

The federal government has a long history of trying to find the right balance in supporting scientific and medical research while protecting the public and other researchers from potential harms. To date, this balance has been generally calibrated differently across contexts – including in clinical care, human subjects research, and research integrity. New challenges continue to face this disparate model of regulation, including novel Generative Artificial Intelligence (GenAI) tools. Because of potential increases in unintentional fabrication, falsification, and plagiarism using GenAI – and challenges establishing both these errors and intentionality in retrospect – this article argues that we should instead move toward a system that sets accepted community standards for the use of GenAI in research as prospective requirements.

Keywords: integrity; ethics; research

Introduction

The federal government has struggled for many decades to find the right balance in funding and encouraging scientific and medical research while also mitigating risk, preventing harm, and punishing wrongdoing by researchers themselves. By and large, the default approach has been to tailor governance models in different ways depending on context, with important regulatory differences in clinical care versus human subjects research versus research misconduct.

This disparate model of regulation continues to face new challenges, including the rapid expansion of studies focused on health data. In addition to big data research, the use of generative artificial intelligence (GenAI) and other machine learning tools raises novel issues for data integrity.¹ For example, because of GenAI's proclivity to copy work without attribution and “hallucinate” by creating false information, unintentional falsification, fabrication, and plagiarism by researchers might become more common. In addition, the clues that others in the scientific community use to recognize that human-generated data are untrustworthy may be upended in potentially irreversible ways, making it increasingly difficult to assess data integrity or even define its scope.

These novel data-driven research challenges call into question whether a retrospective system of response to research misconduct is still the appropriate standard with which to approach responsible conduct of research requirements. This article argues that we should instead move toward a system that sets accepted community standards for the use of GenAI in research as *prospective*

requirements that researchers will be held accountable to — without having to establish that fabrication, falsification, or plagiarism actually occurred.

Human Subjects Research Regulations

Our current human subjects research governance structure began with the end of the infamous US Public Health Service's (PHS) Syphilis Study in Tuskegee, AL (1932–72), during which Black men were deceived, studied, and prevented from receiving treatment for their syphilis.² While the study was known widely within the syphilology and broader “venereal disease” community,³ the general public was not aware of these experiments until 1972 when a whistle blower collaborated with the *New York Times* to expose them.⁴ The resultant article caused a scandal large enough to finally put an end to the experiments (notably over the ongoing objections of some representatives of the US Centers for Disease Control and Prevention (CDC)).⁵ This new regulatory structure built in response to the PHS Syphilis Study represented a professional and legal transition from the tort law approach of holding physicians retrospectively liable for an actual injury to their patient in clinical care,⁶ to a system of prospective regulation for physician researchers (and other scientists) to *prevent* harm to participants in research in the first place.⁷

These “Human Subjects Research Regulations,” the first part of which is called the “Common Rule,” protects research participants and their identifiable data and biospecimens in federally funded research.⁸ Many academic medical centers also extend these protections to all human subjects research at their institutions,⁹ and the US Food and Drug Administration (FDA) has similar protections for research with an investigational-only product¹⁰ or if researchers submit such data to the FDA.¹¹ Under these regulations, a prospective

Email: kaytesb@med.umich.edu

Cite this article: K. Spector-Bagdady. “The Need for Prospective Integrity Standards for the Use of Generative AI in Research,” *Journal of Law, Medicine & Ethics* (2025): 1–6. <https://doi.org/10.1017/jme.2025.41>

participant generally needs to be clearly notified that they are being asked to enroll in research and they have the option to decline or withdraw. Researchers either must obtain informed consent to engage with participants directly or secure an Institutional Review Board (IRB) exemption or waiver for low-risk research, including research with health data or biospecimens previously collected for a different purpose, such as clinical care.¹²

Research Misconduct Regulations

Intentional misconduct

Along with the response to the PHS Syphilis Study, the 1970s and early 80s were a time of intense scrutiny of research misconduct, and ways to respond when it was discovered. It was widely believed that the standard process of co-authorship, peer review, and post-publication replication would root out general misbehaviors.¹³ Despite Representative Al Gore holding the first hearing on research misconduct in 1981,¹⁴ then-National Institutes of Health (NIH) Director, Donald Frederickson, argued that congressional mandates for research regulation were unnecessary as “the natural sciences contain ultimate correctives for any debasement of the knowledge derived from research.”¹⁵ A committee of the Association of American Medical Colleges the following year agreed: “The principal deterrent in research fraud is the overwhelming probability that fraudulent data will be detected soon after their presentation.”¹⁶

Even the infamous case of John Darsee, a Harvard cardiologist caught intentionally fabricating data in many publications, could not shake scientists’ dogged belief that the best way to handle intentional misconduct was post-hoc detection and response.¹⁷ The *New England Journal of Medicine (NEJM)* had published two of Darsee’s retracted articles. *NEJM* Editor Arnold Relman responded to these retractions in an editorial: “It seems paradoxical that scientific research, in many ways one of the most questioning and skeptical of human activities, should be dependent on personal trust” but that “editors and referees of scientific papers ... have *no choice* but to assume that the authors have honestly reported what they did and what they observed (emphasis added).”¹⁸ Relman even went so far as to argue that “A request from an editor for primary data to support the honesty of an authors’ findings in a manuscript under review would probably poison the air and make civil discourse between authors and editors even more difficult than it is now.”¹⁹ Thus, despite the case in point, Relman felt that trust in the individual scientist was still the best path forward to ensure research integrity: “The damage done to morale and the free exchange of ideas may in the long run be far more costly than even the deprivations of an occasional Darsee. In science, as in any other human activities, trust has its risks, but they are far exceeded by the benefits.”²⁰

Unintentional misconduct

Cases of intentional fabrication and falsification were also accompanied by other cases of research misconduct based, ultimately, on negligence and/or poor record keeping. One of the most important cases at the time, dubbed “The Baltimore Case,”²¹ involved scientists Thereza Imanishi-Kari and David Baltimore. In 1986 they, with colleagues, had published a groundbreaking NIH-funded article on murine immunology in *Cell*.²² That same year, Margot O’Toole, a junior molecular biologist in Imanishi-Kari’s laboratory, started challenging what she saw as sloppy and erroneous data in

that publication.²³ While the authors agreed that there were minor errors in the paper, they believed that it did not warrant correction because the errors did not impact the paper’s central findings.²⁴

O’Toole later began to claim that Imanishi-Kari had *intentionally* manipulated the data, and the matter was eventually elevated to an NIH investigation and a 1988 congressional investigation led by Representative John Dingell of Michigan. Rep. Dingell was particularly concerned that the NIH could not “afford to divert precious dollars into areas of meaningless or fraudulent work.”²⁵ The investigatory team ultimately found that Imanishi-Kari’s records were “disparate, unorganized, and scattered around her laboratory.”²⁶

Throughout this period, and heightened by the Rep. Dingell hearings, the scientific community generally strongly opposed any legislative interest in policing science.²⁷ Baltimore himself argued that the investigation into Imanishi-Kari was a “warning [to all scientists] to be vigilant to such threats, because our research community is fragile, easily attacked, difficult to defend, easily undermined.... What is now my problem could become anyone else’s.”²⁸

During this case, in 1989, the PHS published the “Responsibility of PHS Awardee and Applicant Instructions for Dealing with and Reporting Possible Misconduct in Science” to set “uniform policies and procedures for investigating and reporting instances of alleged or apparent misconduct ... supported with funds made available under the PHS Act.”²⁹ It defined research misconduct as “fabrication, falsification, plagiarism, or other practices that seriously deviate from those that are commonly accepted within the scientific community ...” and excluded “honest error or honest differences in interpretations or judgements of data.”³⁰ Sanctions for violations of integrity were generally left up to awardee institutions, although the Office of Scientific Integrity reserved the right to “impose sanctions of its own upon investigators or institutions ... if such action seems appropriate.”³¹ Therefore, while the Baltimore *Cell* article was ultimately retracted in 1991 for erroneous data due to poor record-keeping,³² the researchers were eventually exonerated of research *misconduct* due to a lack of intentionality.

Research misconduct regulations

The US Department of Health and Human Services (HHS) revised the research integrity rules in 2005, reflecting standards implemented by a 2000 memo from the White House Office of Science and Technology Policy.³³ They continued to exclude “honest error” or “difference of opinion” from the definition of a violation of research integrity but added that misconduct must be engaged in with the appropriate *mens rea*, or intentionality to commit, such that the misconduct be “committed intentionally, knowingly, or recklessly” (i.e., not “honest error”) and that the misconduct must represent “a significant departure from accepted practices of the relevant research community” (i.e., not “a difference of opinion”).³⁴ HHS did not define any of these terms. HHS also said it would consider factors such as whether actions were part of a pattern, their impact on research or the public health, acceptance of responsibility, and whether the respondent had retaliated in any way.³⁵

While the terms “intentionally” and “knowingly” are more readily understood, the scope of the term “reckless” in this space has continued to be a matter of debate.³⁶ Legally, “reckless” is generally understood to be something less than “intentional” or “knowing,” but worse than “negligence” (i.e., “the failure to behave with the level of care that a reasonable person would have exercised under the same circumstances”).³⁷ But it was not until the 2018 case of *Office of Research Integrity v Kreipke* that an

Administrative Law Judge (ALJ) formally adopted the Black's Law Dictionary definition of the terms including that “*intentionally* means one acts with the aim of carrying out the act. *Knowingly* means that one acts with knowledge and information and awareness of the act. *Recklessly* means one acts without proper caution despite a known risk for harm.”³⁸ And, unlike the Baltimore case before the 2005 regulations, here — because Kreipke “was aware of disorganization and lack of record keeping when he joined the laboratory” — the ALJ argued that it was *reckless and therefore a violation of research integrity* for him to “simply assume that materials placed in his grants, articles, and posters were reliable.”³⁹ Therefore, in this case, it seems that the justification that Kreipke's actions were reckless was: (1) he had *knowledge* of the disorganization of his lab, and (2) to publish or share data generated by that lab *in the context of that knowledge* was a reckless assumption of their validity.

HHS most recently updated these regulations in September of 2024 (effective as of January 1, 2026) due to “policy developments and technological changes applicable to research misconduct” including NIH's 2023 Data Management and Sharing Policy and “the shift to saving data on the cloud” and “the ability to use artificial intelligence to detect image falsification....”⁴⁰ HHS also cited “increasing public concerns about research integrity in science” as well as institutional questions regarding misconduct review as motivators.⁴¹

The new rule does not change the standards for finding research misconduct, but does add a definition for “accepted practices of the relevant research community” to be, “commonly accepted professional codes or norms within the overarching community of researchers and institutions that apply for and receive PHS awards.”⁴² While the regulations adopted the *ORI v Kreipke* definitions of intentionally and knowingly, it slightly shifted the definition of reckless from the previous “acts *without proper caution* despite a known risk” to “*indifference* to a known risk....”⁴³ Whether that rhetorical shift will have implications in application remains to be seen.

Conflicting systems of governance

Thus, we see the differences and tensions between the different models of governance described above. When a physician is acting as a clinician in the setting of a *doctor patient relationship*, harms are generally assessed through a system of retrospective liability via a medical malpractice claim. This tort-based framework requires that a measurable harm befall a patient, for example an injury, that was caused by the physician's failure to adhere to the standard of care.⁴⁴ In the area of *human subjects research*, governance has generally taken the form of prospective regulation, or requiring compliance with a set of rules written to mitigate risk and prevent harm before it can happen. Under this system, if the same physician is interacting with patients as research participants, the assumption is that the risk of the harm to participants is so high, that not following proscribed steps to protect them should be punishable in and of itself — even if no participant gets hurt.⁴⁵ As John Lantos has argued, “It should not be controversial to claim that our system of research regulation in the United States today is based on a deep distrust of researchers and the entire research enterprise.”⁴⁶ This observation is astute when applied to the transition of a physician engaging with patients to a physician *researcher* engaging with their patients as participants. But the governance of researchers on the issue of integrity is approached differently yet again. In *research*

misconduct, this governance has generally taken the form of retrospective liability, or punishing wrongdoing once it has already occurred. The most punitive actions are also generally targeted for those who intentionally violate the regulations.⁴⁷

Therefore, while clinical malpractice is essentially a “no harm, no foul” system of retrospective governance, and human subjects research a prospective regulatory system, research integrity can be seen as primarily a “no foul, no harm” approach. Indeed, its very name — research *integrity* — clearly indicates that if the action does not lack integrity, it is not a violation.

And yet, intentionality does not alleviate the harm research misconduct can cause, it is just an affirmative defense against serious punishment. One could even argue that the integrity regulations allow unintentionally erroneous data in grants or publications to be considered a “negative externality” — an economic concept for a situation where predicted benefits redound to one party but risks are born by another.⁴⁸ Using the Baltimore case as an example, while originally publishing in *Cell* was a major benefit to the researchers promoting their scientific reputation, there was a serious negative impact on the junior biologist who spent a year trying to replicate the work or the other researchers who cited and built on erroneous data. And just because Imanishi-Kari was later found to not have the intentionality required for a violation of research integrity, that did not alleviate those burdens for others. While intentionality is certainly foundational to the concept of integrity, it does not otherwise rectify harms caused by erroneous data produced without it.

Research Integrity for Generative AI Technologies

In the time since research misconduct regulations were originally conceptualized, much has advanced in scientific and medical research that make unintentionally erroneous data easier to produce, and intentionally erroneous data harder to establish. For example, large-scale collections of health data now allow researchers to analyze correlations between genetic variants, behavior, environment, and health outcomes to piece together what outcomes individuals can modify, or clinicians can treat, and subsequently improve.⁴⁹ This includes promising advances in big data research, the use of AI, GenAI, and other machine learning tools.⁵⁰ Methods of tracking and recording data have also changed dramatically with movement toward fully electronic datasets, cloud computing, and data enclaves.⁵¹

With the promise of new data-driven technologies, however, come novel challenges for regulation and education regarding the responsible conduct of research.⁵² As an example, consider the emerging technology of GenAI, which is the “simulation of human intelligence by machines.”⁵³ It includes innovative Large Language Models which can analyze data, including from the internet and electronic medical records, to generate its own derivative synthetic content such as text and images. One GenAI developer goal is to make healthcare more efficient and cost-effective.⁵⁴ For example, Epic, the electronic medical record (EMR) platform supporting most US hospitals, is integrating generative pre-trained transformers (GPT-4) into academic medical centers across the country.⁵⁵

But the use of GenAI in health research poses new and distinct research integrity challenges. It produces its own derivative health information, which raises questions about how to characterize this information for use, sharing, and assessments of integrity.⁵⁶ GenAI can also hallucinate or copy existing information in ways that average researchers might not realize or be able to control.⁵⁷

Recently, Open AI's tool "Whisper," an auto-transcription (automatic speech recognition) system sometimes described as "ambient listening," entirely made up portions of doctor-patient conversations including hallucinating "racial commentary, violent rhetoric and even imagined medical treatments."⁵⁸ In another recent study, researchers found that using machine learning techniques tied to genomic information in order to predict otherwise missing health outcome data generated unreliable results.⁵⁹

Humans, obviously, can also intentionally or unintentionally make informational errors or manipulate data. As discussed above, many have used the ability for other researchers to notice this kind of erroneous data as a justification for the current, retrospective, system. But the clues that others use to recognize that human generated data are untrustworthy will be upended by GenAI in potentially irreversible ways.⁶⁰ For example, Imanishi-Kari was eventually exonerated and Kreipke was eventually sanctioned after examination of written lab notebooks. There have also been more recent cases in which skeptics requested fabricated electronic data that could not be produced.⁶¹ For example, in 2020 a research team published articles in both *The Lancet* and *NEJM* finding increased mortality rates associated with the use of hydroxychloroquine and blood pressure medications to treat COVID-19, allegedly based on data from 671 hospitals across six continents.⁶² The World Health Organization even paused clinical trials around the world involving thousands of COVID patients in response. When questioned by other researchers, however, the lead authors admitted that they had neither accessed nor reviewed the data upon which their findings were allegedly based. Those data ended up being entirely fabricated.⁶³

New generative technologies will make it increasingly hard to assess data integrity in this way. In fact, *JAMA Ophthalmology* recently published the first report of authors instructing GenAI to entirely fabricate a dataset in a way that would produce a statistically significant difference in outcomes between surgical interventions. The dataset produced by GenAI was "seemingly authentic" to expert analyses.⁶⁴ Authors concluded that the capabilities of GPT-4 "may pose a greater threat [than previous GPT models], being able to fabricate data sets specifically designed to quickly produce false scientific evidence...."⁶⁵

New Governance Proposals for Data Driven Research

There has been movement to start to update federal regulation in response to these technologies. For example, in October 2022, the Biden Administration released the *White House Blueprint for an AI Bill of Rights*, which offered a principlist approach to ensuring benefits of AI, while attempting to mitigate negative impact, including: (1) safe and effective systems, (2) algorithmic discrimination protections, (3) data privacy, (4) notice and explanation, and (5) human alternatives, consideration, and fallback.⁶⁶

Although initially well-received,⁶⁷ the AI Blueprint's premise — that industry would use it to facilitate "actualizing these principles in the technological design process"⁶⁸ — was upended the very next month with the release of GenAI and ChatGPT and resulting industry competition to get rival products to the market as quickly as possible. GenAI was moving *so* quickly, in fact, that many experts, including those in industry, called to "Pause Giant AI Experiments" and to "dramatically accelerate development of robust AI governance systems" such as "provenance and watermarking systems to help distinguish real from synthetic" and "liability for AI-caused harm."⁶⁹ That pause never happened.

The disruption of GenAI motivated the Biden Administration to put forth more concrete plans to guide and regulate the use of AI

and GenAI. The Biden AI Executive Order, the longest Executive Order in history, instructs HHS to "develop a strategic plan that includes policies and frameworks – possibly including regulatory action, as appropriate – on responsible deployment and use of AI and AI-enabled technologies in the health and human services sector" including "development, maintenance, and use of predictive and generative AI."⁷⁰ In November 2023, the Office of Management and Budget released a related request for comments on proposed guidance for agencies to "increase their capacity to successfully and responsibly adopt AI, including generative AI...." including "how errors from data entry [or] machine processing ... are adequately measured and limited, to include errors from relying on AI generated data as training data or model inputs."⁷¹ Trump revoked the Biden EO on his first day in office,⁷² and most recently announced a \$500 billion private investment in AI infrastructure,⁷³ making the future regulation of AI — if any — unclear.

In addition, as described above, the revised 2024 research integrity regulations argued that the changes were in response to "policy developments and technological changes applicable to research misconduct" including "the shift to saving data on the cloud" and "the ability to use artificial intelligence to detect image falsification...."⁷⁴ And yet, these issues are only raised in the preamble, and the parts of the updated regulations intended to apply to or alleviate the integrity challenges of AI technologies is unclear.

One thing that the 2024 regulations *did* do was specifically define "accepted practices of the relevant research community" to be "commonly accepted professional codes or norms within the overarching community of researchers and institutions that apply for and receive PHS awards."⁷⁵ While this shift is subtle, it potentially opens a new opportunity. To defer to "accepted practices of the relevant research community" requires establishing what those practices are, for example, such as what the average reasonable researcher would have done under similar circumstances (much like a "standard of care" for clinical medicine).

Unlike clinical medicine, however, the government and individual institutions should start seriously considering whether retrospective liability *if* research misconduct is established will be effective enough to protect the public in a GenAI-enabled research universe. I argue that this is likely not enough for two reasons: First, despite its inability to act with intentionality, GenAI's proclivity to copy others' work without attribution or hallucinate data means that *unintentional* research misconduct will likely become more common. Second, because bad actors will be able to easily fabricate datasets that support outcomes of interest, *intentional* research misconduct will become harder to establish.

Both arguments cut in favor of setting "accepted practices of the relevant research community" as *prospective requirements* that a researcher will be held accountable for following without a third party having to establish that an actual injury occurred. For example, if an accepted practice is to disclose that GenAI was used in the generation of a figure, and an investigator does not disclose that, then that lack of disclosure should warrant a response in and of itself — even if falsification, fabrication, or plagiarism is not (or cannot) be established. Many journals are beginning to set these standards as requirements for publication of research that uses GenAI, although what happens if they are not followed is not clear⁷⁶ and recent studies have found the practice to be already widespread.⁷⁷

Regulators should also consider switching to this prospective standard, like that of the human subjects research regulations, requiring compliance with a set of rules written to mitigate risk

and prevent fabrication, falsification, and plagiarism when using GenAI research tools before it can happen. This is because the risk of producing erroneous data unintentionally is so high — and the ability to establish intentionality so challenging — that researchers not following proscribed steps to protect against them should be strongly discouraged in and of itself.

Conclusions

The federal government has a long history of trying to find the right balance in supporting scientific and medical research while protecting the public and other researchers from potential harms. To date, this balance has been generally calibrated differently across contexts — including in clinical care, human subjects research, and research integrity. New challenges continue to face this disparate model of regulation, with GenAI being an excellent example of a new research tool that raises novel issues for data integrity.⁷⁸ Because of the likely increase of unintentional fabrication, falsification, and plagiarism while using GenAI tools, and challenges establishing both these errors and intentionality in retrospect, we should instead move toward a prospective regulatory system that sets community standards for the use of GenAI in research.

Acknowledgements. This work was funded by The Greenwall Foundation Faculty Scholars Program, the National Center for Translational Sciences (R01TR004244), and the National Institute on Aging (U54AG084520). The author would also like to thank Holly Fernandez Lynch, Nicolle Strand, Stephen Rosenfeld, and Kerry Ryan for their feedback on a previous draft. All errors are her own.

Kayte Spector-Bagdady is an Associate Professor at the University of Michigan Medical School and its Center for Bioethics & Social Sciences in Medicine in Ann Arbor, MI. She received her JD and MBe from the University of Pennsylvania in Philadelphia, PA.

References

1. K. Spector-Bagdady, “Generative-AI-Generated Challenges for Health Data Research,” *The American Journal of Bioethics* 23, no. 10 (2023): 1–5, <https://doi.org/10.1080/15265161.2023.2252311>.
2. J. Jones, *Bad Blood: The Tuskegee Syphilis Experiment*, Revised Edition (Free Press, 1993).
3. J.J. Peters et al., “Untreated Syphilis in the Male Negro,” *Journal of Chronic Diseases* 1, no. 2 (1955): 127–48.
4. J. Heller, “Syphilis Victims in U.S. Study Went Untreated for 40 Years,” *New York Times*, July 26, 1972, <https://www.nytimes.com/1972/07/26/archives/syphilis-victims-in-us-study-went-untreated-for-40-years-syphilis.html>.
5. C. Elliott, *The Occasional Human Sacrifice: Medical Experimentation and the Price of Saying No* (W. W. Norton & Company, 2024).
6. B.S. Bal, “An Introduction to Medical Malpractice in the United States,” *Clinical Orthopaedics and Related Research* 467, no. 2 (2008): 339, <https://doi.org/10.1007/s11999-008-0636-2>.
7. J.D. Lantos, “The Regulation of Clinical Research: What’s Love Got to Do with It?,” *Learning Health Systems* 2, no. 1 (2018): e10048, <https://doi.org/10.1002/lrh2.10048>.
8. C.F.R. § 46 (2018).
9. Department of Homeland Security et al., “Federal Policy for the Protection of Human Subjects,” 82 *Federal Register* 12 (2017), 7149–7274, at 7156, <https://www.federalregister.gov/documents/2017/01/19/2017-01058/federal-policy-for-the-protection-of-human-subjects#h-70>.
10. C.F.R. § 312 (1987).
11. C.F.R. § 50.1 (2019).
12. Federal Policy for the Protection of Human Subjects, *supra* note 9.
13. A.S. Relman, “Lessons from the Darsee Affair,” *New England Journal of Medicine* 308, no. 23 (1983): 1415–1417, at 1416, <https://doi.org/10.1056/NEJM198306093082311>.
14. Office of Research Integrity, *Historical Background*, <https://ori.hhs.gov/historical-background> (last visited August 31, 2024).
15. N.H. Steneck, “Research Universities and Scientific Misconduct: History, Policies, and the Future,” *Journal of Higher Education* 65, no. 3 (1994): 310–331, <https://www.jstor.org/stable/2943970>.
16. *The Maintenance of High Ethical Standards in the Conduct of Research* (Association of American Medical Colleges Ad Hoc Committee, 1982), <https://www.aamc.org/media/16196/download> (last visited February 27, 2025).
17. W.J. Broad, “Notorious Darsee Case Shakes Assumptions About Science,” *New York Times*, June 14, 1983, sec. C, <https://www.nytimes.com/1983/06/14/science/notorious-darsee-case-shakes-assumptions-about-science.html>.
18. See Relman, *supra* note 13, at 1415.
19. See Relman, *supra* note 13, at 1416.
20. See vRelman, *supra* note 13, at 1417.
21. D. Kevles, *The Baltimore Case: A Trial of Politics, Science, and Character* (W. W. Norton & Co Inc, 1998).
22. D. Weaver et al., “Altered Repertoire of Endogenous Immunoglobulin Gene Expression in Transgenic Mice Containing a Rearranged Mu Heavy Chain Gene,” *Cell* 45, no. 2 (April 25, 1986): 247–259, [https://doi.org/10.1016/0092-8674\(86\)90389-2](https://doi.org/10.1016/0092-8674(86)90389-2).
23. See Kevles, *supra* note 21, at 242; C.K. Gunsalus, Book Review: “The Baltimore Case: A Trial of Politics, Science, and Character,” *New England Journal of Medicine* 340, no. 242 (1999), <https://www.nejm.org/doi/abs/10.1056/NEJM199901213400320>.
24. See Kevles, *supra* note 21.
25. See Kevles, *supra* note 21, at 135.
26. See Kevles, *supra* note 21, at 155.
27. See Kevles, *supra* note 21, at 146; F. Hesselmann and M. Reinhart, “From Scandal to Reform: Approaches to Research Integrity at a Turning Point,” *Journal of Responsible Innovation* 11, no. 1 (2024): 2414491, <https://doi.org/10.1080/23299460.2024.2414491>.
28. Kevles, *supra* note 21, at 151.
29. PHS, “Responsibility of Awardee and Applicant Institutions for Dealing with and Reporting Possible Misconduct in Science; Final Rule,” *Federal Register* 54, no. 151 (August 8, 1989): 32446–51.
30. *Id.*
31. See PHS, *supra* note 29.
32. D. Weaver et al., “Retraction: Altered Repertoire of Endogenous Immunoglobulin Gene Expression in Transgenic Mice Containing a Rearranged Mu Heavy Chain Gene,” *Cell* 65, no. 4 (May 17, 1991): 536, [https://doi.org/10.1016/0092-8674\(91\)90085-D](https://doi.org/10.1016/0092-8674(91)90085-D).
33. Office of Science and Technology Policy, “Federal Policy on Research Misconduct,” *Federal Register* (Dec. 6, 2000), <https://www.govinfo.gov/content/pkg/FR-2000-12-06/pdf/00-30852.pdf>.
34. Department of Health and Human Services, “Public Health Service Policies on Research Misconduct,” *Federal Register*, May 17, 2005, <https://www.govinfo.gov/content/pkg/FR-2005-05-17/pdf/05-9643.pdf>.
35. C.F.R. § 93.408 (2005).
36. M.M. Caron et al., “Defining ‘Recklessness’ in Research Misconduct Proceedings,” *Accountability in Research* (2023): 1–23, <https://doi.org/10.1080/08989621.2023.2256650>.
37. “Negligence,” Legal Information Institute, Cornell Law School, <https://www.law.cornell.edu/wex/negligence> (last visited November 27, 2025).
38. *ORI v Kreipke*, Departmental Appeals Board, Civil Remedies Division, ORI (2018), https://ori.hhs.gov/sites/default/files/2018-07/282755_C-16-402%20RECOMMENDED%20DECISION_20180531104917.pdf (last visited November 3, 2024).
39. *Id.*
40. “Public Health Service Policies on Research Misconduct,” *Federal Register* 89, no. 180 (September 17, 2024), <https://www.federalregister.gov/documents/2024/09/17/2024-20814/public-health-service-policies-on-research-misconduct>.
41. *Id.*
42. C.F.R. § 93.200 (2024).

43. C.F.R. § 93.231 (2024).
44. See Bal, *supra* note 6.
45. See Lantos, *supra* note 7.
46. See Lantos, *supra* note 7, at 1.
47. See Kevles, *supra* note 21.
48. S. Eldridge, “negative externality,” *Encyclopedia Britannica*, December 23, 2024, <https://www.britannica.com/topic/negative-externality>, (last visited November 3, 2024).
49. National Research Council, *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease* (National Academies Press, 2011), <https://doi.org/10.17226/13284>.
50. *Final NIH Policy for Data Management and Sharing* (Office of The Director, National Institutes of Health, October 29, 2020), <https://grants.nih.gov/grants/guide/notice-files/NOT-OD-21-013.html>, (last visited September 1, 2024).
51. J. Lane and C. Schur. “Balancing access to health data and privacy: a review of the issues and approaches for the future,” *Health Services Research* 45, no. 5, Pt. 2 (2010): 1456–1467, <https://doi.org/10.1111/j.1475-6773.2010.01141.x>.
52. See Spector-Bagdady, *supra* note 1.
53. D. Milmo, “ChatGPT Reaches 100 Million Users Two Months after Launch,” *Guardian*, Feb. 2, 2023, <https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app> (last visited February 27, 2025).
54. P. Lee et al., *The AI Revolution in Medicine: GPT-4 and Beyond* (Pearson, 2023), <https://www.pearson.com/en-us/subject-catalog/p/the-ai-revolution-in-medicine-gpt-4-and-beyond/P200000011399/9780138279516>.
55. “Microsoft and Epic Expand Strategic Collaboration with Integration of Azure OpenAI Service,” *Microsoft News Center*, April 17, 2023, <https://news.microsoft.com/2023/04/17/microsoft-and-epic-expand-strategic-collaboration-with-integration-of-azure-openai-service/> (last visited February 27, 2025).
56. I.G. Cohen, “What Should ChatGPT Mean for Bioethics?,” *The American Journal of Bioethics* 23, no. 10 (2023): 8–16, <https://doi.org/10.1080/15265161.2023.2233357>; K. Spector-Bagdady and A.J. London, “Disclosure as Absolution in Medicine: Disentangling Autonomy from Beneficence and Justice in Artificial Intelligence,” *American Journal of Bioethics* 25, no. 3 (2025): 1–3, <https://doi.org/10.1080/15265161.2025.2458424>.
57. E.M. Bender et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜,” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT’21 (Association for Computing Machinery, 2021), 610–623, <https://doi.org/10.1145/3442188.3445922> (last visited Feb. 27, 2025); OpenAI et al., *GPT-4 Technical Report*, arXiv: 2303.08774v6 [cs.CL] (March 4, 2024), <https://doi.org/10.48550/arXiv.2303.08774> (last visited February 27, 2025).
58. G. Burke and H. Schellmann, “Researchers say an AI-powered transcription tool used in hospitals invents things no one ever said,” *AP* (October 26, 2024), <https://apnews.com/article/ai-artificial-intelligence-health-business-90020cdf5fa16c79ca2e5b6c4c9bbb14> (last visited February 27, 2025).
59. Y. Wu et al., “Pervasive Biases in Proxy Genome-Wide Association Studies Based on Parental History of Alzheimer’s Disease,” *Nature Genetics* 56, no. 12 (2024), <https://doi.org/10.1038/s41588-024-01963-9>.
60. See Lee, *supra* note 54.
61. J. Bohannon, “Science retracts gay marriage paper without agreement of lead author LaCour,” *Science*, May 28, 2015, <https://www.science.org/content/article/science-retracts-gay-marriage-paper-without-agreement-lead-author-lacour> (last visited February 27, 2025); M.R. Mehra et al., “Retraction—Hydroxychloroquine or Chloroquine with or without a Macrolide for Treatment of COVID-19: A Multinational Registry Analysis,” *Lancet* 395, no. 10240 (2020): 1820, [https://doi.org/10.1016/S0140-6736\(20\)31324-6](https://doi.org/10.1016/S0140-6736(20)31324-6).
62. Mehra, *supra* note 61.; M.R. Mehra et al., “Retraction: Cardiovascular Disease, Drug Therapy, and Mortality in Covid-19. N Engl J Med. DOI: 10.1056/NEJMoa2007621.” *New England Journal of Medicine* 382, no. 26 (2020): 2582–2582, <https://doi.org/10.1056/NEJMc2021225>.
63. Joseph, “Lancet, New England Journal Retract Covid-19 Studies, Including One That Raised Safety Concerns about Malaria Drugs,” *STAT*, June 4, 2020, <https://www.statnews.com/2020/06/04/lancet-retracts-major-covid-19-paper-that-raised-safety-concerns-about-malaria-drugs/> (last visited February 27, 2025).
64. Taloni et al., “Large Language Model Advanced Data Analysis Abuse to Create a Fake Data Set in Medical Research,” *JAMA Ophthalmology* 141, no. 12 (2023): 1174, <https://doi.org/10.1001/jamaophthalmol.2023.5162>.
65. *Id.*
66. The White House, “Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People,” archived version at the Internet Archive, <https://web.archive.org/web/20250120000129/https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf> (snapshot on January 20, 2025).
67. Atlantic Council experts, “Experts React: What Does Biden’s New Executive Order Mean for the Future of AI?,” *Atlantic Council*, Oct. 30, 2023, <https://www.atlanticcouncil.org/blogs/new-atlanticist/experts-react/experts-react-what-does-bidens-new-executive-order-mean-for-the-future-of-ai/> (last visited September 1, 2024).
68. See The White House, *supra* note 66.
69. “Pause Giant AI Experiments: An Open Letter,” *Future of Life Institute*, March 22, 2023, <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> (last visited September 1, 2024).
70. The White House, “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” October 30, 2023, archived version at the Internet Archive, <https://web.archive.org/web/20250120132537/https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/> (snapshot on January 20, 2025).
71. Executive Office of the President, “Proposed Memorandum for the Heads of Executive Departments and Agencies,” (November 2023), archived version at the Internet Archive, <https://web.archive.org/web/20250116091611/https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf> (snapshot on January 16, 2025).
72. D. Shepardson, “Trump revokes Biden executive order on addressing AI risks,” *Reuters* (Jan. 21, 2025), <https://www.reuters.com/technology/artificial-intelligence/trump-revokes-biden-executive-order-addressing-ai-risks-2025-01-21/> (last visited February 27, 2025).
73. Duffy, “Trump announces a \$500 billion AI infrastructure investment in the US,” *CNN*, January 21, 2025, <https://www.cnn.com/2025/01/21/tech/openai-oracle-softbank-trump-ai-investment/index.html>, (last visited February 27, 2025).
74. See Public Health Service Policies on Research Misconduct, *supra* note 40.
75. 42 C.F.R. § 93.200 (2024).
76. Elsevier, “Generative AI policies for journals,” <https://www.elsevier.com/about/policies-and-standards/generative-ai-policies-for-journals> (last visited September 1, 2024).
77. Glynn, “Suspected Undeclared Use of Artificial Intelligence in the Academic Literature: An Analysis of the Academ-AI Dataset,” arXiv:2411.15218 (Nov. 20, 2024), <https://arxiv.org/abs/2411.15218>.
78. See Spector-Bagdady, 2023, *supra* note 1.