Fixing Human Research Policy

By Molly Klote, MD

Executive Summary

For 50 years, the United States has employed a regulatory system to protect human research participants, which began with a unified purpose but has evolved into a complex, layered system. While this system has effectively prevented many past violations, it has also increased costs for scientific progress, reduced efficiency, and diminished public trust. The Common Rule oversees federally funded research, while the Food and Drug Administration manages product-related studies. Additionally, each federal agency, state, and institution creates its own interpretations and rules. This creates an ethical paradox: a system meant to protect participants now hinders innovation and obscures accountability.

This paper contends that the nation must enhance human research oversight by replacing unnecessary procedures with evidence-based coordination. Achieving this requires three interconnected reforms:

1. A National Center for Ethical Research Learning (NCERL). NCERL will focus on learning and collecting anonymized data from IRBs, institutions, and agencies to assess how protection systems function. It will oversee and coordinate evidence-based policy initiatives. It will establish national benchmarks for ethical quality, efficiency, and equity, enabling policy development based on evidence rather than assumptions. It will not regulate nor promulgate policy.
2. Statutory, Federal administration, and State legal modernization:
    - Rescind the Paperwork Reduction Act's application to federally funded research overseen by Human Research Protection Programs (HRPPs), ending redundant OIRA review that delays low-risk studies by months with no added protection.
    - Repeal or amend 10 U.S.C. § 980 to allow the Department of Defense to engage in minimal-risk and emergency research under the same standards as all other agencies.
    - Align implementation instructions across Common Rule Agencies to eliminate procedural conflicts that hinder interdepartmental collaboration.
    - Encourage interstate harmonization of State privacy and human-subjects protection laws through model state laws or federal incentives, to remove conflicts that slow national, multi-jurisdictional research.

3. A government-wide AI-driven harmonization tool. Before any federal agency releases new human research policies or guidance, it should run the draft through a federally maintained Harmonization Screening Tool that automatically compares language across existing statutes, regulations, and guidance. The system would flag contradictions or redundancies, allowing human reviewers to correct them before publication. This analytical safeguard ensures coherence without transferring authority to machines. (see Appendix). States would be encouraged to follow similar implementation of any new State law.

Together, these reforms will transform the current patchwork into a learning, adaptive, and transparent system that protects participants while accelerating the delivery of safe, fair innovation. The United States can regain its leadership in global ethics, not by creating new rules but by making existing ones work smarter and more consistently.

1. Introduction

Federal oversight of human subjects research operates across three distinct legal and policy layers. First, Congress enacts statutes, such as the National Research Act and the Public Health Service Act, that establish agencies, such as the Office for Human Research Protections (OHRP), and set broad requirements for research oversight. Second, federal agencies create regulations under statutory authority; the Common Rule is a binding regulation published in the Federal Register that details specific standards for research conduct and IRB operations. Third, OHRP issues non-binding guidance documents and interpretive materials intended to clarify, advise, or recommend practices for compliance with statutes and regulations. While statutes and regulations carry the force of law, guidance provides necessary interpretation for implementation support but does not create legally enforceable obligations on its own.

Human research protection in the United States has developed more through responses to events than through planning. Today's regulatory environment is complicated. Two main federal frameworks operate under the Department of Health and Human Services (HHS): the Common Rule (45 CFR part 46, which includes Subpart A and Subparts B, C, D, and E), which oversees federally funded research managed by the Office for Human Research Protections (OHRP), and the Food and Drug Administration's regulations that apply to products such as drugs, medical devices, and biologics used in human research.

Although these frameworks share common ethical goals, they differ in scope, definitions, and enforcement approaches. Twenty other federal agencies, along with HHS (for a total of 21 departments and agencies), have adopted the Common Rule, but each implements it with its own instructions and guidance. The Department of Defense (DoD), the Department of Veterans Affairs, the Department of Energy, the National Institute of Standards and Technology, and

others interpret these principles differently. Additionally, State laws and tribal sovereignty add further layers of complexity.

The result is expensive duplication. Institutions conducting multi-site, multi-State, or interagency studies must resolve conflicting requirements for consent language, data security, and reporting timelines. Review cycles extend into months; valuable research is delayed or abandoned. Every delayed study postpones potential benefits to patients, service members, and the public. Ethically, oversight structures meant to ensure fairness now amplify inequity: community hospitals and smaller universities cannot afford the compliance infrastructure of major academic centers, limiting who can participate in federally funded research. Meanwhile, participants face consent forms that are too long, obscuring understanding rather than promoting autonomy. Protection has become performative.

Legally, the issue is structural. The Administrative Procedure Act requires each agency to conduct its own notice-and-comment rulemaking. The Supreme Court's June 28, 2024, decision in *Loper Light Enterprises v. Raimondo* overturned the four-decade-old *Chevron* deference doctrine. Courts must now interpret statutes de novo rather than automatically deferring to agency interpretations. In human-research policy, this change is profound: if Congress uses terms like "minimal risk" or "informed consent" without clear definitions, judges, not scientists or regulators, will ultimately determine their meaning. Future legislation must therefore rely on demonstrable evidence, not interpretive habits.

This paper asserts that true modernization requires learning, accuracy, and coordination rather than increasing bureaucracy. The aim is to develop infrastructure that enables OHRP, FDA, and their federal partners to learn from outcomes, standardize language, and continuously update policies. NCERL will provide the empirical foundation; legislative reform will remove outdated restrictions; and AI-driven harmonization screening will ensure consistent guidance.

2. Historical Evolution of U.S. Human Research Protections

A. Origins

The modern system for protecting human subjects emerged from the aftermath of war. In 1947, the Nuremberg Code established the first worldwide declaration that voluntary consent is "absolutely essential." When the World Medical Association adopted the Declaration of Helsinki in 1964, the principle of informed consent developed into a formal doctrine: research must balance risks against benefits, and subject welfare against scientific goals.

For twenty years, these international statements lacked a domestic counterpart. U.S. investigators operated under local customs rather than national law. It was scandals, the 1963 Jewish Chronic Disease Hospital studies, and later the Tuskegee Syphilis Study (publicly exposed in 1972), that prompted congressional action. Public outrage led to the National Research Act of 1974, establishing the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. Its final report, The Belmont Report (1979), outlined the core principles of

Respect for Persons, Beneficence, and Justice. These three principles continue to serve as the fundamental moral foundation of every U.S. research regulation.

B. Institutionalization: The Common Rule

To implement the Belmont principles, originally 16 federal agencies, along with HHS (for a total of 17 departments and agencies), adopted the Federal Policy for the Protection of Human Subjects, known as the Common Rule, as a unified regulatory framework in 1991. The Common Rule represented a consolidation of existing departmental regulations into a single baseline. Codified at 45 CFR part 46, it required all federally supported institutions to obtain Institutional Review Board (IRB) review and obtain documented informed consent for most studies, and perform ongoing reviews of active human research.

The Common Rule succeeded in creating a national baseline of protection, but a uniform text did not lead to uniform practice. Each agency retained interpretive authority over its own implementation, leading to different understandings of key terms. Institutions conducting multi-agency studies must still reconcile conflicting expectations, usually defaulting to the most restrictive interpretation. What started as efforts to harmonize gradually became fragmented—precisely the situation the Common Rule aimed to prevent.

C. The FDA's Parallel Universe

As the Common Rule evolved, the Food and Drug Administration (FDA) maintained its separate, product-based system under 21 CFR parts 50 and 56. Its primary focus was on evidentiary integrity: ensuring that data supporting new drugs, devices, and biologics were produced under conditions that protected both participants and scientific validity. The FDA's system closely mirrored the Common Rule in requiring IRBs and informed consent, but subtle differences in documentation requirements, waiver restrictions, exemptions, and emergency-research rules created operational incompatibility.

As federal agencies, industry, and academia increasingly collaborated, sponsors were confronted with multiple regulatory oversight for identical protocols. Each discrepancy required reconciliation through legal opinion letters and duplicative review cycles. The two systems are ossified into parallel bureaucracies.

D. Post-Belmont Expansion

The 1970s and early 1980s introduced new protections under 45 CFR 46. Federal bioethics commissions recommended special protections for vulnerable groups, including pregnant women, prisoners, and children, and these were issued by the then Department of Health, Education, and Welfare (HEW), which is now known as HHS. When Subpart A was codified in 1991, those earlier issuances were then designated as the subparts.

In 1996, Congress passed the Health Insurance Portability and Accountability Act (HIPAA), establishing privacy and authorization rules that have become central for protecting all

identifiable health information. HIPAA's Privacy Rule was created to regulate electronic health transactions generally, not specifically for research, yet it now primarily guides IRB discussions whenever identifiable data are involved.

In 2008, Congress passed the Genetic Information Nondiscrimination Act (GINA). The Genetic Information Nondiscrimination Act (GINA) significantly benefits research participants by forbidding health insurers and employers from using genetic data for coverage or employment decisions, thereby alleviating a barrier to participation in genetic studies. However, GINA also makes informed consent forms more complex by requiring detailed explanations of these protections, their limits, and the confidentiality measures in place, resulting in longer and more technical documents that can be difficult for participants to navigate.

By the early 2000s, academic medical centers operated under at least five overlapping frameworks: their institutional policy, the Common Rule, FDA regulations, HIPAA, and various State laws. Compliance offices expanded; IRB chairs took on additional roles as risk managers alongside their ethical responsibilities.

E. The 2017 Revisions and Unfinished Harmonization

In 2017, HHS issued the first significant revision of the Common Rule since its promulgation in 1991. The final rule published on January 19, 2017, with a general compliance date of January 21, 2019, introduced new exempt research categories, required concise key-information summaries in consent forms, and formally required single-IRB review for multi-site research. However, these changes were based primarily on anecdotes and expert opinion rather than empirical evidence, and there was no plan to assess whether the new provisions achieved their aims, nor was there guidance on implementation in complex scenarios.

This experience highlights that policy changes alone are insufficient unless paired with mechanisms to measure their impact and adapt over time. Without infrastructure for ongoing evaluation and revision, well-intended reforms can quickly fall out of step with practical realities, demonstrating that ethical oversight depends on a learning system, not just periodic rule changes.

3. Legal Foundations and the Administrative Landscape

A. The Administrative Procedure Act

Every federal regulation, from drug approval to human-subject protections, must go through the Administrative Procedure Act (APA), enacted in 1946 as a compromise between congressional oversight and executive efficiency. The APA established a standard process for rulemaking: public notice, opportunity for comment, review and reconciliation of comments, and, finally, publication in the Federal Register.

For multi-agency frameworks like human-research policy, the APA's procedural design has become a brake on modernization. Each Common Rule agency must conduct its own notice-and-

comment rulemaking, even when proposing identical language. This multiplies administrative workload and elongates timelines. When HHS revised the Common Rule in 2017, the process took nearly eight years from notice to implementation.

Moreover, agencies cannot simply harmonize policy through informal agreement. Under the Administrative Procedure Act (5 U.S.C. § 553), any interagency agreement with binding regulatory effect must undergo public notice and comment. The result is a paradox: everyone agrees harmonization is desirable, but the APA requires that any joint rules affecting the public follow the same procedural safeguards as single-agency regulations. The United States remains locked in an administrative stasis where ethical principles evolve faster than the rules designed to uphold them can.

B. The End of Chevron Deference

For forty years, *Chevron U.S.A. v. Natural Resources Defense Council (1984)* shaped the balance of power between courts and agencies. Under that rule, when Congress passed semi-ambiguous laws and courts deferred to reasonable agency interpretations. This deference allowed technical agencies, such as OHRP and FDA, to adapt regulations flexibly through guidance as science advanced.

On June 28, 2024, the Supreme Court's decision in *Loper Light Enterprises v. Raimondo* overturned the four-decade-old *Chevron* deference doctrine. Courts must now interpret statutes de novo, exercising independent judgment rather than deferring automatically to agency interpretations. If Congress uses terms like "minimal risk" or "informed consent" without clear definitions, judges, not scientists, will ultimately determine their meaning. Although agencies' expertise may still receive Skidmore-style deference (*Skidmore v. Swift & Co.*, 323 U.S. 134 (1944)), it is no longer the controlling standard.

This legal shift is seismic. It limits agencies' ability to extend regulation beyond clear statutory authority. Conversely, it puts unprecedented pressure on Congress to craft legislation with technical accuracy. Future laws governing research oversight must be based on empirical evidence that can withstand judicial review. This requires a data infrastructure that can provide legislators and regulators with measurable evidence of ethical outcomes.

In the post-Chevron era, evidence becomes a form of authority. Agencies will require demonstrable results, not assumptions, to justify rulemaking and help Congress enact its intent.

C. Statutory Authority and Federalism

Human-subject protections hold a unique constitutional role. They are not explicitly mentioned in the Constitution but mainly stem from Congress's Spending Clause authority, which permits attaching ethical conditions to federal funds. The Common Rule applies to federally funded research because institutions voluntarily accept these conditions in exchange for federal grants. The FDA's oversight of clinical research stems from Congress's Commerce Clause power and the

Food, Drug, and Cosmetic Act, which governs the interstate sale of medical products. These different constitutional sources create overlapping but not identical protection systems.

States retain independent powers over health and safety, allowing them to enact their own privacy laws, age-of-majority laws, consent requirements, and research restrictions. Native American tribes, acting in their own sovereign capacity, may impose additional requirements. The result is a mosaic of jurisdictional variability. The federal government sets a baseline, but the practical complexity is determined locally. Without federal and interstate coordination mechanisms, harmonization is structurally constrained by design.

An effective analogy for understanding research oversight in the United States is the division of authority between federal and State levels in transportation law. Just as the federal government sets rules for interstate commerce, such as standards for trucking and moving goods across State lines, while States can establish their own speed limits and local traffic laws, research oversight in the U.S. functions at two levels. The federal Common Rule provides basic protections for federally funded human subjects research; however, individual States are free to add stricter or additional requirements for research within their borders. This results in a complex, sometimes overlapping regulatory environment similar to how drivers traveling interstate must comply with both national interstate rules and State speed limits.

D. Judicial Review and the Evolution of Ethical Policy

Judicial review has played an underappreciated role in shaping research oversight. Courts in the early 2000s affirmed IRBs' authority to waive documentation of consent in minimal-risk studies (*Robertson v. McGee*, No. 4:01-CV-60, 2002 WL 535045 (N.D. Okla. Jan. 28, 2002)), while later decisions upheld patient privacy and autonomy even when data use promised clear public benefit (*Sorrell v. IMS Health Inc.*, 564 U.S. 552 (2011)). The cumulative effect is a legal landscape that favors caution and procedural rigor.

4. The Fragmented Federal Oversight System

A. The Office for Human Research Protections and the Common Rule Agencies

OHRP, located within the Department of Health and Human Services, serves as the principal steward of the Common Rule. It enforces compliance with federal funding recipients and interprets the policy for the 21 agencies that have adopted it. OHRP's authority is advisory to other federal agencies rather than binding. Each adopting agency retains its own discretion in implementation.

B. The FDA's Parallel Oversight

The FDA oversees a separate but overlapping regime. Its mission, to protect the safety and efficacy of regulated products, drives its approach to protecting human subjects. While the FDA's IRB and informed consent rules appear similar to those of the Common Rule, critical distinctions persist.

The agencies have pursued harmonization, mandated by the 21st Century Cures Act of 2016, through joint guidance and proposed regulations. However, without a shared data infrastructure and given their fundamentally different missions, OHRP, which protects human subjects in federally funded research, and FDA, which ensures product safety and efficacy, complete alignment remains incomplete and, in some areas, may not be entirely achievable.

C. Divergent Implementing Instructions

Every federal agency that has adopted the Common Rule codifies it in its own section of the Code of Federal Regulations, often with agency-specific modifications to definitions, procedures, or documentation requirements. These regulations are published in the Federal Register and carry the force of law for research funded, conducted, or supported by that agency:

- Department of Veterans Affairs (VA): VA Directive 1200.05 adds procedural steps, including documentation of information security, privacy, and research and development committee review, in addition to the IRB.
- Department of Defense (DoD): Under DoDI 3216.02, research supported with DoD-appropriated funds that meets the Instruction's definition of a "human being as an experimental subject" (applying 10 U.S.C. § 980) generally requires advance informed consent.
- Department of Energy (DOE): Under Order 443.1C, it imposes human-subjects protections for any research it supports, even when the Common Rule would not automatically apply, thereby extending oversight but also adding institutional review layers.
- National Aeronautics and Space Administration (NASA): Under NPD 7100.8G, mandates compliance with the Common Rule and additional NASA-specific procedures, including for research aboard aircraft or involving astronauts.

Other agencies have similar requirements (e.g., USDA, EPA) and overlay the baseline federal framework, contributing to regulatory complexity and a multi-site operational burden. This duplicative system wastes scarce research resources for multi-agency work in understanding and complying with all requirements.

D. State, Tribal, and Institutional Layers

Beyond federal fragmentation, States and tribes hold independent authority to protect their residents. Many states have enacted privacy or genetic information laws that exceed federal requirements. Institutional policies add yet another layer. Most research universities and hospitals maintain their own "local policies" that elaborate on the Common Rule. These internal rules often add documentation requirements, expand categories requiring full board review, or impose additional consent elements.

5. Systemic Barriers to Reform

A. Process-Centric Metrics

The metrics used to judge oversight effectiveness measure process, not ethics. Agencies count the time to complete the review, but neither correlates with comprehension nor participant safety. There is no national data infrastructure linking regulatory burden to ethical outcomes.

The system confuses activity with accountability. Institutions that produce more paperwork seem compliant; those that innovate appear riskier. This counterproductive incentive structure discourages experimentation in ethical review models.

B. Administrative Inertia and Risk Aversion

Reform efforts within federal agencies face deep structural resistance. Staff turnover, divided jurisdiction, and political sensitivity make policy innovation a rare occurrence. No agency wants to appear to be "weakening protections."

Institutions reflect this inertia. Compliance officers worry that any departure from current practices could lead to audits or threaten funding. The system is designed to prevent criticism, not to enhance results.

C. Professional Silos

Multiple professionals now oversee ethics: regulators, ethicists, clinicians, data privacy officers, and compliance lawyers. Each operates with different motivations. Without a unifying data framework, disputes remain unresolved. Without a shared learning system, experience builds but never consolidates, and is not compared and justified. The absence of institutional memory causes each IRB to reinvent ethical reasoning on its own.

D. Operational Redundancy

The 2017 Common Rule revisions introduced a single-IRB requirement for multi-site research to simplify the approval process. However, because participating agencies and institutions often add extra "local context reviews," the requirement rarely results in a single, unified oversight. Sites still perform redundant reviews to meet agency-specific or institutional rules.

Without unified definitions, digital infrastructure, and shared trust, procedural reform cannot fix cultural fragmentation. The system's inefficiencies stem from incompatible architecture.

6. Statutory and Procedural Modernization

Modernizing U.S. human research oversight doesn't require new ethical principles. It requires the courage to eliminate outdated barriers and the discipline to align existing ones. The following reforms, all of which are doable within current institutional frameworks, would be a start at turning stagnation into progress.

A. Reforming the Paperwork Reduction Act

The Paperwork Reduction Act of 1980 (PRA) was conceived to protect the public from unnecessary federal information collection. In the research context, it now obstructs ethically approved, low-risk studies that have already undergone regulatory review.

Under the PRA, any federally sponsored survey or data collection involving ten or more individuals requires approval from the Office of Management and Budget's (OMB) Office of Information and Regulatory Affairs (OIRA). This secondary review duplicates IRB oversight and routinely adds six to twelve months to project timelines.

When the Uniformed Services University (USU) piloted a legislative exemption from PRA review, the results were unambiguous: faster study initiation and no decrease in data integrity. In a 2023 report to Congress, the USU exemption (which became effective January 31, 2021) saved 25,506 administrative days and $318,198 in costs during 24 months without a single compromise to participant protection.

The conclusion is clear: the PRA, as applied to human research that has undergone regulatory review, is neither protective nor permissive; it is prohibitive.

Policy Recommendation: Congress should amend 44 U.S.C. § 3502(3) to exclude all human-subject research subject to the requirements of 45 CFR 46 or 21 CFR 50/56a from the definition of "information collection." To functionally implement while Congress acts, the OMB should immediately issue guidance clarifying that studies subject to 45 CFR 46 or 21 CFR 50/56 are exempt from PRA clearance. This simple revision would restore months of lost time to many federally sponsored research studies.

B. Repealing or Modernizing 10 U.S.C. § 980

Enacted in 1984 and became effective in 1985, 10 U.S.C. § 980 prohibits the Department of Defense from funding research involving human subjects without prospective consent. It was a well-intentioned reaction to Cold War-era fears of military service-member abuse. But four decades later, it is a relic, inconsistent with the Common Rule, and counterproductive to ethical science.

The statute bars the DoD from participating in studies where informed consent is waived under 45 CFR 46.116(f), even when risk is minimal and the IRB has formally approved the waiver. It also blocks emergency and casualty-care trials that could save lives.

Policy Recommendation: Congress should repeal 10 U.S.C. § 980 in its entirety, or amend it to mirror the Common Rule and FDA waiver provisions for minimal-risk and emergency research. This reform would not weaken protections; it would unify them, allowing DoD research to proceed under the same ethical architecture as every other federal entity.

C. Harmonizing Implementing Instructions

Even if statutes align, fragmentation will persist unless agencies standardize their implementing instructions. Today, every department maintains its own manual.

Policy Recommendation: The Office of Science and Technology Policy (OSTP) should convene a Human Research Harmonization Task Force, co-chaired by OHRP and FDA, to standardize implementing language across all Common Rule agencies. Using model templates for definitions, documentation, and IRB procedures, the Task Force could achieve practical harmonization without the need for new legislation. NCERL's analytic outputs would ultimately inform these alignments. This initiative would finally make the "single IRB" concept a reality, not just a rhetorical concept, at least within interagency studies.

D. AI-Enabled Policy Harmonization Screening

Regulatory consistency must be sustained, not just restored. To prevent re-fragmentation, the federal government should deploy a Retrieval-Augmented Generation (RAG) Harmonization Screening Tool to review new policy drafts for conflicts or redundancies.

Under this system, when an agency drafts new human-research policy or guidance, it would submit the text to a secure, government-operated AI platform, maintained by the NIST or the OSTP's AI Policy Lab (or another designated agency). The system would:

1. Retrieve all relevant statutes, regulations, and guidance from the federal corpus;
2. Compare the draft text semantically against existing language;
3. Flag inconsistencies, missing cross-references, or definitional divergences; and
4. Generate a harmonization report for human review before public notice.

The AI tool would serve as a pre-publication quality control layer, a form of "ethical spell-check" for policy coherence. It would not interpret or decide policy, but ensure that new language does not reintroduce avoidable contradictions. States would be encouraged to develop a similar tool and apply to have new State research-related laws evaluated by the Federal Harmonization Screening Tool. (See Appendix)

7. Establishing a National Center for Ethical Research Learning (NCERL)

A. Purpose and Mission

The United States needs a mechanism not for more regulation but for better evidence about how regulation performs. The National Center for Ethical Research Learning (NCERL) would serve as that mechanism, a non-regulatory, data-driven entity that studies the performance of human-subject protections nationwide.

Its mandate would be transformative but straightforward: collect, analyze, and publish metrics on how oversight functions, what delays it creates, how well it protects participants, and how

equitably research opportunities are distributed. NCERL would convert ethics from an assertion into an empirical discipline.

Unlike OHRP or FDA, NCERL would have no enforcement power. Its credibility would rest on neutrality, transparency, and analytic rigor.

B. Core Functions

1. Ethical Performance Analytics: Collect de-identified national data on IRB processing times, approval rates, adverse events, and participant comprehension.
2. Regulatory Experimentation: Be funded to support pilot projects testing proportional-review models, simplified consent, or risk-based oversight structures.
3. Annual Learning Report: Publish an evidence-based report assessing the ethical and operational impact of current oversight systems, providing Congress, OHRP, and FDA with actionable recommendations.
4. Policy Evaluation Collaboration: Partner with agencies to evaluate the real-world outcomes of new regulations or guidance using post-implementation data.

C. The Decision Bank

A cornerstone of NCERL's mission would be a National Decision Bank, a secure, de-identified national repository of IRB determinations, exemption categorizations, HIPAA applications, research vs not-research determinations, and the supporting rationales. Each record would summarize key decision variables: study type, risk level, regulatory provisions cited, waiver decisions, and the reasoning used.

The Decision Bank would allow IRBs, agencies, and policymakers to learn from non-binding precedent. It would reveal patterns in ethical reasoning and identify sources of inconsistent interpretation. Participating institutions would submit structured decision data through a standardized interface. NCERL analysts could then detect trends, such as disproportionate risk determinations for certain studies, that indicate where definitions or rules need clarification and pass that information to OHRP and FDA.

By aggregating thousands of decisions nationwide, the Decision Bank would create the first empirical map of ethical judgment in the United States.

D. Governance and Transparency

To preserve trust, NCERL should operate as a federally funded research and development center (FFRDC) or within the Agency for Healthcare Research and Quality (AHRQ). Governance would include a board representing OHRP, FDA, VA, DoD, NIH, academic institutions, and patient advocates.

All data would be anonymized and publicly reported through dashboards summarizing national metrics. NCERL's methodology, data sources, and algorithms would be open to independent audit to ensure neutrality.

8. Global Harmonization and International Leadership

A. The Global Context

Human-subject research is now an international enterprise. Clinical trials often span multiple continents, and multinational consortia share biospecimens and genomic information in real-time. Yet while science has globalized, ethics remains local.

The United States, long a moral leader in bioethics, is increasingly perceived as an administratively complex entity. The European Union Clinical Trials Regulation (EU CTR) of 2022 created a unified digital portal for approvals across 27 member states. The Organisation for Economic Co-operation and Development (OECD) and the Council for International Organizations of Medical Sciences (CIOMS) have developed harmonized ethical frameworks adopted by most developed nations. Meanwhile, U.S. institutions navigate a complex web of inconsistent federal and state regulations.

Without harmonization at home, the U.S. cannot credibly advocate for it abroad.

B. Alignment with Global Standards

Modernization must include international interoperability. U.S. definitions of "minimal risk," "consent," "vulnerability," and "secondary use of data" should align conceptually with those in the EU CTR, OECD Good Clinical Practice guidelines, and CIOMS 2023 Principles. Doing so would facilitate reciprocal recognition of ethics review, enhance scientific agility during public-health emergencies, and reinforce U.S. leadership as a trusted guardian of ethical research.

OHRP and FDA already participate in the International Council for Harmonisation (ICH). A unified domestic system would strengthen their negotiating position.

C. NCERL's International Role

NCERL can serve as the analytic bridge between domestic reform and global alignment. By benchmarking U.S. oversight performance against international comparators, consent comprehension rates, and definitional risk application, it can identify both strengths and areas for improvement.

Over time, NCERL could host a Global Ethics Metrics Forum, convening regulators, researchers, and patient representatives to exchange lessons on consent innovation, risk stratification, and data-sharing safeguards.

9. Implementation Pathways

A. Phased Approach

Phase 1: Foundation (Legislative and Administrative Setup)

- Congress appropriates funding to authorize NCERL as a federally funded research and development center (FFRDC) or AHRQ agency..
- OSTP establishes the Human Research Harmonization Task Force to coordinate federal agency alignment efforts.
- NIST or OSTP's AI Policy Lab begins developing the harmonization-screening tool prototype.
- NCERL is officially launched with initial staffing, governance structure, and pilot infrastructure.

Phase 2: Operationalization (Data Infrastructure and Pilot Programs)

- OHRP, FDA, and participating agencies formalize data-sharing agreements with NCERL.
- Federal funding agencies amend grant language to require reporting to the NCERL and include additional funding support for reporting.
- NCERL launches pilot projects to test data collection, analysis, and reporting systems.
- The Initial Decision Bank dataset is established with mandatory federally sponsored research reporting and voluntary institutional participation.
- Harmonization-screening tool is piloted for internal agency policy review.
- A State-of-the-Policy Report is created as the starting point for federal alignment.
- OIRA grants waiver for research approved under 45 CFR 46 or 21 CFR 50/56 from OIRA PRA review.

Phase 3: Statutory Reform and Scale-Up

- Congress amends the Paperwork Reduction Act to exempt research approved under 45 CFR 46 or 21 CFR 50/56 from OIRA PRA review.
- Congress repeals or modernizes 10 U.S.C. § 980 to align DoD research with Common Rule waiver provisions.
- DOD revises DODI 3216.02 to reflect the change of 10 USC 980.
- Agencies adopt harmonized implementing instructions based on Task Force recommendations and NCERL data.
- Harmonization-screening tool becomes mandatory for all new human-research policy drafts.

Phase 4: Cultural Integration and Continuous Improvement

- Institutions receive incentives (streamlined review, reduced reporting burden) for voluntary data submission to NCERL.
- NCERL publishes annual "Ethics Learning Reports" with national benchmarks and policy recommendations.
- Ongoing evaluation and refinement of harmonization tools, data systems, and regulatory alignment.
- NCERL coordinates with international partners to streamline international research interoperability through metrics.

B. Managing Institutional Resistance

Reform will meet resistance from entrenched bureaucratic habits and institutional risk aversion. The antidote is evidence. NCERL's early pilots should demonstrate quantifiable benefits: shorter review times, consistent consent quality, participant understanding, without any decline in participant safety. Once success is empirical rather than speculative, resistance will erode.

Transparency also defuses opposition. Public dashboards summarizing NCERL metrics can show taxpayers and Congress that reform saves both time and money while maintaining protection.

C. Building Coalitions

Lasting change requires broad alignment. Reformers should cultivate a coalition that includes federal agencies, academic and industry sponsors, patient and veteran organizations, and philanthropic partners. By framing modernization as both an ethical and economic imperative, the coalition can transcend partisanship.

10. Conclusion and Policy Recommendations

The United States stands at a turning point. Its human-research oversight system, once a global leader, now lags behind the pace of scientific innovation and international cooperation. This paper demonstrates that reform is not only possible but urgently necessary.

The actions described in Section 9 would replace procedural stagnation with empirical agility. They would transform oversight from a defensive bureaucracy into a proactive learning system that protects participants, empowers investigators, and restores U.S. leadership in global bioethics.

By reforming its own system, the United States can once again establish the global standard for what constitutes ethical research.

Works Cited

- The Nuremberg Code. 1947.
- World Medical Association. *Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects.* Adopted 1964; most recent revision 2013.
- National Research Act. Pub. L. No. 93-348, 88 Stat. 342 (1974).
- The Belmont Report. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1979).
- Administrative Procedure Act. 5 U.S.C. §§ 551–559, 701–706.
- Federal Policy for the Protection of Human Subjects ("Common Rule"). 45 C.F.R. part 46 (1991; revised 2017; general compliance 2019).
- FDA Human Subject Protections. 21 C.F.R. parts 50, 56; and related IND/IDE regulations, 21 C.F.R. parts 312, 812.
- Health Insurance Portability and Accountability Act (HIPAA). Pub. L. No. 104-191 (1996); Privacy Rule, 45 C.F.R. parts 160, 164.
- Genetic Information Nondiscrimination Act (GINA). Pub. L. No. 110-233 (2008).
- 21st Century Cures Act. Pub. L. No. 114-255 (2016).
- Paperwork Reduction Act. 44 U.S.C. §§ 3501–3521.
- 10 U.S.C. § 980. Human subjects; limitation on research without consent.
- Chevron U.S.A. Inc. v. Natural Resources Defense Council, Inc., 467 U.S. 837 (1984).
- Loper Light Enterprises v. Raimondo, 603 U.S. ___ (2024).
- International Council for Harmonisation (ICH). *E6 Good Clinical Practice* (E6(R2) 2016; E6(R3) draft).
- Council for International Organizations of Medical Sciences (CIOMS). *International Ethical Guidelines for Health-related Research Involving Humans* (2016; 2023 update).
- European Union Clinical Trials Regulation. Regulation (EU) No 536/2014 (in force 2022).

Appendix

The Harmonization Screening Tool Powered by AI-Driven RAG Architecture

Purpose and Context

As part of a comprehensive modernization of U.S. human-research oversight, this paper proposes a Harmonization Screening Tool. This federal, AI-assisted system automatically reviews proposed policy and guidance documents for consistency across agencies before they are published for public comment.

The purpose of the tool is to detect and prevent policy conflicts before they occur, ensuring that new regulations, implementing instructions, and guidance align semantically and substantively with existing federal human-research frameworks.

This screening process will be powered by an AI-driven Retrieval-Augmented Generation (RAG) architecture, an advanced information-processing system that combines precise document retrieval with natural-language reasoning. The RAG engine enables the tool to compare new policy language against the entire body of existing regulations and automatically flag discrepancies.

Importantly, the Harmonization Screening Tool is not part of NCERL. It will be operated independently under the technical stewardship of the National Institute of Standards and Technology (NIST) or another designated organization and coordinated through the Office of Science and Technology Policy (OSTP). NCERL remains a separate, non-regulatory learning infrastructure that examines the effectiveness of policies.

Operational Framework

1. Repository Construction

The system will maintain a secure, continuously updated repository of all relevant statutes, regulations, and guidance, including:

- The Common Rule (45 CFR 46) and its agency-specific implementations;
- FDA human-subject regulations (21 CFR 50, 56, 312, 812);
- Foundational statutes such as the Food, Drug, and Cosmetic Act and Public Health Service Act;
- Departmental manuals and implementing instructions (e.g., VA Handbook 1200.05, DoDI 3216.02);
- Historical Federal Register notices, OHRP and FDA guidance, and interpretive memos.

All documents are embedded in a vector database, allowing the AI to search and compare based on semantic meaning rather than exact wording. This enables the detection of definitional drift and unintended conflicts across agency texts.

2. Draft Policy Submission

When an agency (for example, OHRP, FDA, VA, or DoD) develops new or revised human-research guidance, it submits the draft to the secure screening portal. The system operates within a FedRAMP-authorized government cloud environment, ensuring confidentiality and compliance with federal information security standards.

3. AI-Powered Retrieval and Comparison

Using its RAG architecture, the system:

- Retrieves relevant passages from existing regulations and guidance related to each section of the draft;
- Compares definitions, requirements, and procedural language across agencies;
- Flags inconsistencies or gaps (e.g., diverging consent definitions, waiver criteria, or reporting timelines); and
- Generates a structured summary identifying where harmonization may be strengthened.

This process is fully automated and produces output within minutes to hours, reducing the time currently required for months of interagency review.

4. Harmonization Report Generation

The system produces a Harmonization Analysis Report detailing:

- The relevant sources and citations for each flagged issue;
- The nature of the discrepancy (minor variation, definitional conflict, procedural contradiction);
- A severity or priority level; and
- Plain-language summaries for agency reviewers.

The report is sent to both the issuing agency and the Human Research Harmonization Task Force (co-chaired by OHRP and FDA under OSTP coordination). The report's purpose is advisory; final policy decisions are made solely by humans and are fully transparent.

The first use of the system should be a state-of-the-policy report detailing where all current regulations are inconsistent or unnecessarily duplicative as a starting point for alignment.

Governance and Oversight

To maintain neutrality and credibility, the Harmonization Screening Tool will be managed outside of policy agencies.

- NIST (or designated organization) will oversee the technical infrastructure, algorithmic transparency, and cybersecurity.
- OSTP will coordinate interagency use, ensuring every Common Rule agency applies the tool before publication.
- The Human Research Harmonization Task Force (OSTP/OHRP/FDA co-chaired) will review output reports but will not alter the AI's analytic process.

Annual audits will evaluate accuracy, false-positive rates, and interagency responses to flagged issues. Summary statistics (not policy drafts) should be publicly released to show accountability and ongoing improvement.

Technical Feasibility

The proposed architecture relies on mature, commercially proven technologies:

- Vector Databases: (e.g., FAISS, Milvus, Pinecone) to store and search millions of regulatory text fragments.
- Embedding Models: such as OpenAI's text-embedding-3 series or government-trained equivalents, to encode text semantically.
- Large Language Models (LLMs): specialized for legal and regulatory text summarization, hosted securely in government cloud environments.
- APIs and Audit Logs: to document retrieval, reasoning steps, and final outputs for transparency and reproducibility.

The necessary technical infrastructure is already in use in other government applications, such as the AI Bill of Rights prototype systems at NIST. Deployment would focus on integration rather than creation.

Safeguards and Limitations

To maintain trust and accuracy, several safeguards are essential:

1. Human Accountability: The tool identifies conflicts but never modifies policy text or issues interpretations. Agencies retain full decision-making authority.

2. Corpus Integrity: The accuracy of analysis depends on keeping the document repository complete and current; a dedicated interagency data maintenance team will manage updates.
3. Explainability and Traceability: Every flagged item will include explicit citations so human reviewers can verify the AI's reasoning.
4. Privacy and Security: Drafts remain within government-secure systems; no external or public LLM interfaces are used.
5. Continuous Evaluation: OSTP and NIST will monitor performance and bias, publishing audit results annually.

The system's strength is in prevention: by identifying inconsistencies before publication, it prevents years of interpretive confusion.

Relationship to NCERL

The Harmonization Screening Tool is independent of NCERL.

- NCERL's mandate is to collect and analyze data on how oversight performs, timelines, equity, and ethical outcomes, after policies are implemented.
- The Harmonization Tool's mandate is to analyze the text of policies before they are published.

Occasional data sharing might happen: if NCERL detects patterns of confusion or bottlenecks in implementation, it could provide anonymized insights to NIST or OSTP to improve the tool's document collection. However, NCERL does not have the authority to review, approve, or vet policy language.

This clear separation ensures that NCERL remains an analytical institution, not a policy office, preserving its credibility as the nation's evidence hub for ethical oversight and performance.

Integration with International Standards

The architecture can later be expanded to include global ethical frameworks, such as the EU Clinical Trials Regulation, OECD Good Clinical Practice, and CIOMS Guidelines, enabling automatic crosswalk analysis between U.S. and international terminology. This would position the United States as a leader in transparent, evidence-driven harmonization of global research ethics.

Conclusion

The Harmonization Screening Tool, powered by AI-driven RAG architecture, embeds intelligence into the machinery of government. By integrating this system into the rulemaking

workflow, the federal government can maintain consistency, eliminate redundant regulation, and restore agility to human-research oversight.

Together with NCERL's evidence-based learning mission, this tool ensures that future policy is both ethically principled and technically coherent, a model of how artificial intelligence can strengthen, rather than supplant, human judgment in public governance.

Disclaimer

This manuscript was prepared with the assistance of ChatGPT and Perplexity, AI language models, which contributed to editing and technical synthesis under the author's direction. The author takes responsibility for the content.