AI Accountability Requires Competition and Evidence

Robert P Strauss, Professor of Economics & Public Policy, Carnegie-Mellon; rpstrauss@gmail.com; www.andrew.cmu.edu/user/rs9f October 14,2026

Public concern about artificial intelligence has concentrated on whether the technology will become too powerful, eliminate jobs, spread misinformation, or escape meaningful human control. Those concerns are legitimate. But the debate too often treats AI policy as a choice between rapid innovation and restrictive regulation. A sound policy must do something more practical: preserve competition among AI providers while making those providers answerable for the harms they cause.

The benefits of AI are already tangible. Combined with instructional video, an AI service can help a homeowner diagnose a repair, judge whether it is feasible, identify the right tools and parts, and compare prices within minutes. This reduces the informational advantage of contractors and sellers, expands do-it-yourself activity, and directs purchases toward retailers offering the best combination of price and availability. Similar gains are possible in education, professional services, and small business. These benefits depend, however, on users having a choice among trustworthy, and ethical services. A market controlled by a few vertically integrated companies could convert a technology that lowers information costs into another source of dependency and economic power.

The first requirement is vigorous antitrust enforcement. Artificial intelligence has strong tendencies toward concentration; it is a classic example of economies of scale. Training advanced models requires enormous computing capacity, specialized chips, immense stores of data legitimately obtained, and access to huge, cloud platforms. A few firms can use control over those inputs, as well as exclusive distribution arrangements, to foreclose competitors. If policymakers allow those firms to dominate the market, consumers, businesses, universities, and governments may become dependent on a handful of private systems whose prices, operating rules, and political influence they cannot effectively challenge.

Regulation itself can worsen this danger through regulatory capture. Dominant companies can absorb elaborate compliance expenses that smaller entrants cannot. They can also help write technical standards that favor their own systems. A regulatory structure advertised as protecting the public may

therefore protect established firms from competition. Antitrust scrutiny of acquisitions, exclusive cloud agreements, preferential access to chips and data, and self-preferencing by digital platforms must be part of AI policy from the beginning. Interoperability and the ability to move data and work among competing models would further reduce dependence on any one supplier.

Competition, however, cannot substitute for rules governing safety and conduct. Ethical principles are useful only when they independently define concrete duties. AI providers should test their systems for foreseeable hazards, disclose important limitations, protect confidential information, respond to discovered defects, and maintain effective human oversight in high-risk uses. Requirements should reflect the application. A system offering restaurant suggestions should not face the same obligations as one used to make medical, employment, lending, insurance, or public-safety decisions.

Those duties also need legal consequences. Voluntary promises invite firms to publicize safety while quietly accepting risks that increase market share. Conversely, imposing liability for every mistaken answer could make entry impossible and leave only the largest companies able to insure or litigate. The better approach is to assign responsibility according to control. Model developers should be responsible for foreseeable defects in design, training, testing, and warnings. Businesses deploying models should be responsible for unsafe applications and inadequate supervision. Professional users should retain their established duties of competence, while end users should remain responsible for deliberate misuse.

Yet even well-designed liability rules will fail without evidence. An injured person may need to establish which model and version produced an output, what system instructions governed it, what information the model retrieved, which safeguards operated, what warnings were given, and whether the provider already knew of similar failures. If the relevant, nonalterable records were never created, were routinely erased, or remain inaccessible behind trade-secret claims, a right to sue exists largely on paper.

This makes recordkeeping a central element of AI regulation rather than an administrative afterthought. Providers and deployers of higher-risk systems should preserve proportionate records of model versions, material system settings, safety tests, significant incidents, inputs and outputs needed

to reconstruct disputed decisions, and the identity of human reviewers. Once litigation or a regulatory investigation is reasonably foreseeable, ordinary preservation duties should prevent deletion. Courts should be able to protect legitimate trade secrets through confidentiality orders without denying injured parties the evidence required to prove fault and causation.

Recordkeeping must also respect privacy. Requiring indefinite retention of every personal conversation would create its own danger. Congress and regulators should establish risk-based retention periods, strict access controls, secure segregation of sensitive material, and prompt deletion when neither safety nor legal accountability requires continued storage. The objective is not mass surveillance. It is preservation of the minimum reliable evidence necessary to investigate consequential failures.

These policies as a system support one another. Competition limits the power of individual AI companies and gives users alternatives. Enforceable standards specify the care those companies owe. Liability gives them a financial reason to comply. Record keeping permits regulators, courts, and injured parties to determine what actually happened. Remove any one element and the system weakens: standards without liability become aspirations; liability without records becomes largely unenforceable; regulation without competition can become a barrier protecting the firms already in control.

AI can lower the cost of knowledge, expand individual capability, and make markets more efficient. It can also concentrate economic power and distribute harm on an unprecedented scale. The appropriate response is neither fear nor exemption. It is a competitive market governed by enforceable duties and supported by evidence. That framework would allow innovation to continue while ensuring that those who design and sell AI services remain accountable for the consequences.