ACADEMIA WITHOUT SCHOLARS: ARTIFICIAL INTELLIGENCE AND THE EROSION OF INDEPENDENT JUDGMENT

Kildeev Adel
Leningrad State University named after A.A. Zhdanov (1990)
PhD in Law
ORCID: 0009-0001-2211-4944

Crossref DOI: Pending Registration
SSRN DOI: https://doi.org/10.2139/ssrn.7620899

Abstract: This article examines the risk that contemporary academic institutions may preserve the visible products of scholarship while progressively weakening the human capacities required to produce, verify, and transmit knowledge independently. Rather than attributing academic mediocrity to generative artificial intelligence or declaring the collapse of modern universities, the study situates LLM-assisted scholarly production within longstanding problems of publication incentives, research integrity, institutional authority, and intellectual continuity.

Drawing upon epistemology, philosophy of language, legal reasoning, research on cognitive offloading and automation, and empirical evidence from the United States, Canada, and England, the article investigates the growing asymmetry between the speed of generating academically convincing texts and the human capacity to understand and independently verify their contents. Its empirical component examines changes in higher-education participation, measured educational skills, academic employment, doctoral career pathways, university financing, and research expenditure. These indicators reveal significant institutional transformations and contrasting trends, but do not directly measure independent scholarly judgment or establish causal cognitive deterioration attributable to artificial intelligence.

The analysis distinguishes legitimate technological assistance from the substantive substitution of human intellectual work. It argues that linguistic sophistication, AI-use disclosure, formal peer review, and nominal human oversight cannot independently establish scholarly competence or meaningful verification. Preserving the integrity of academic knowledge requires competent evaluation, accountable authorship, institutional independence, and sustained investment in research mentorship and intellectual continuity.

The study concludes that *academia without scholars* describes a possible institutional condition rather than an accomplished historical fact: universities may expand their measurable academic output while weakening the processes through which independent scientific judgment is cultivated and reproduced.

Keywords: Artificial Intelligence; Large Language Models; Independent Academic Judgment; Cognitive Offloading; Academic Verification; Research Integrity; Peer Review; Institutional Amnesia; Higher Education; Intellectual Continuity.

The use of the generative models ChatGPT (OpenAI), Gemini (Google), Grok (xAI) in drafting this article illustrates, rather than undermines, its central thesis: LLMs function as instrumenta vocalia—tools assisting in linguistic production without possessing understanding, intention, or authorship. Their involvement does not alter the locus of legal or intellectual responsibility, which remains exclusively with the human author, who makes all substantive, structural, and evaluative decisions. LLM outputs cannot be accorded independent legal or intellectual status apart from the human author.

Scope of Interdisciplinary Evidence

This article is a work of legal philosophy and institutional analysis. It does not undertake independent psychological, medical, or educational research, nor does it offer clinical assessments of cognitive impairment. Findings from these disciplines are examined as reported in published specialist studies, subject to their respective methodological limitations. Their inclusion serves exclusively to assess the possible implications of empirical evidence for independent scholarly judgment, epistemic responsibility, and academic institutions. No generalized cognitive deterioration or causal relationship is inferred beyond what the cited studies establish.

The question is," said Humpty Dumpty,
"which is to be master — that's all."

Lewis Carroll, *Through the Looking-Glass*
(1871), Chapter VI,
quoted by Lord Atkin in his dissent in
Liversidge v Anderson [1942] AC 206.

Chapter 1. Introduction: The Crisis of Independent Scholarship

The contemporary university occupies a paradoxical position. Never before have academic institutions possessed such extensive technological capabilities for the production, storage, distribution, and retrieval of information. Scholarly publications circulate across international repositories, educational programmes reach increasingly diverse populations, and digital systems enable access to research materials on a scale unimaginable to previous generations. Yet the expansion of access to information raises a question that cannot be answered by reference to technological progress or institutional output alone: to what extent do contemporary educational and academic institutions continue to cultivate the independent intellectual judgment upon which the production, verification, and transmission of knowledge ultimately depend?

This question becomes particularly important in the context of the widespread adoption and integration of artificial intelligence into education, scientific research, academic publishing, and professional intellectual activity. Generative large language models can produce sophisticated academic prose, summarize literature, organize arguments, and imitate many of the linguistic conventions traditionally associated with scholarly work.

Their capabilities offer considerable opportunities for research and education. At the same time, they create new possibilities for separating the production of academically recognizable texts from the independent intellectual activity historically required to produce, evaluate, and defend them. The existence of such possibilities does not establish that intellectual degradation is universal or inevitable. It does, however, require investigation into the institutional conditions under which technological assistance becomes a substitute for independent judgment.

The crisis of contemporary academia cannot be adequately explained by the widespread adoption and integration of artificial intelligence alone. Any serious investigation must place these technologies within the longer history of institutional transformation, changing educational economics, the expansion and restructuring of higher education, academic commercialization, and the growing importance of measurable institutional performance. Artificial intelligence enters an academic environment already shaped by decades of political decisions, economic pressures, administrative reforms, and changing expectations concerning the purposes of universities and scientific research.

1.1. Historical and Geographical Scope

The present study examines the period from 1990 to 2026. The selection of 1990 is deliberate, but it does not imply that the problems of contemporary education originated in that year. Earlier developments in university governance, educational financing, academic professionalization, and scientific organization remain historically relevant. A comprehensive account of those developments, however, would exceed the purposes and reasonable scope of the present article.

The beginning of the 1990s provides a meaningful historical reference point because it coincided with a major transformation of the European political and institutional order. The dissolution of the socialist political systems of Central and Eastern Europe, followed by the collapse of the Soviet Union in 1991, initiated substantial changes in national economic structures, public administration, institutional financing, and educational governance. These transformations did not occur uniformly, nor did they produce identical consequences across different countries. Nevertheless, they provide an important context for examining the restructuring of educational institutions during the subsequent decades.

The same period witnessed broader changes in international economic integration, university financing, educational mobility, and the organization of higher education. In Europe, the Bologna Declaration of 1999 marked a significant stage in efforts to create more comparable higher education qualifications and facilitate academic mobility. Subsequent reforms modified degree structures, accreditation arrangements, and institutional expectations across participating countries. Whether these changes strengthened educational quality, weakened particular forms of disciplinary preparation, or produced different effects across national systems remains a question for empirical examination rather than an assumption adopted in advance.

Although European developments provide an important historical reference point, the systematic empirical comparison is confined to **the United States**, **Canada**, and **England**. Sources concerning other countries may be used for historical, philosophical, or methodological context, but they do not enlarge the comparative sample. The article does

not claim to provide an exhaustive history of education in every country, nor does it presume that national systems experienced a common trajectory of institutional decline.

The terminal date of 2026 identifies the contemporary context of the investigation. Where statistical series or institutional studies do not extend to that year, the analysis will employ the most recent available evidence and identify the actual observation period.

1.2. Subject, Research Questions, and Central Hypothesis

The subject of this article is the transformation of the institutional conditions under which independent scholarly judgment is developed, exercised, verified, and transmitted between generations.

For the purposes of this study, independent scholarly judgment denotes the capacity to formulate research questions, critically examine sources, assess competing explanations, distinguish evidence from assertion, and defend conclusions through reasoned argument. Such independence does not require intellectual isolation or freedom from technological assistance. Scientific knowledge has always depended upon instruments, collaboration, institutional resources, and inherited intellectual traditions. The decisive distinction concerns whether the researcher retains the competence and responsibility necessary to evaluate the work performed with those instruments.

The investigation addresses three interconnected questions. First, how have changes in educational financing, institutional governance, degree structures, and academic incentives affected the conditions of scholarly preparation and research since 1990? Second, what do longitudinal educational assessments, sociological investigations, and studies of scientific activity reveal about the development of relevant intellectual competencies and the institutional production of knowledge? Third, how might the widespread adoption of generative AI interact with these existing conditions, either by supporting independent inquiry or by encouraging the delegation of intellectual tasks without corresponding verification?

The central hypothesis is that the expansion of formal educational provision and scholarly output does not necessarily entail an equivalent strengthening of independent intellectual competence. Where institutional incentives increasingly privilege measurable production over the demonstrable quality of research and judgment, educational systems may become vulnerable to a divergence between the appearance of academic productivity and the substantive reproduction of scholarly knowledge. Generative AI may intensify this divergence by reducing the effort required to produce plausible academic language without necessarily ensuring the understanding, originality, or evidentiary reliability of the resulting work.

This proposition is advanced as a hypothesis requiring critical examination, not as a predetermined conclusion concerning the condition of contemporary universities.

1.3. Research Methodology and Limitations

The article employs an interdisciplinary methodology combining comparative historical-institutional analysis, secondary analysis of educational and sociological evidence, critical examination of scholarship on academic organization and scientific

productivity, and philosophical and jurisprudential analysis of independent judgment and institutional authority.

The historical-economic component compares selected indicators of educational financing, participation, institutional organization, and academic employment around the beginning of the 1990s with the most recent available data. Intermediate developments are considered where necessary to explain significant reforms or changes in institutional conditions. These comparisons will distinguish total expenditure from expenditure per student, public financing from private contributions, and quantitative expansion from qualitative educational outcomes.

The sociological component draws upon established international assessments and research concerning educational competencies, university practices, academic employment, and scientific production. Differences in national coverage, sampling, measurement, and historical comparability will be acknowledged. Educational test results, publication statistics, and institutional performance indicators will not be treated as interchangeable measures of intellectual capacity or scholarly quality.

The philosophical and jurisprudential component examines the relationship between intellectual judgment, the production of meaning, and the authority claimed by institutions that certify knowledge. Particular attention is given to the distinction between the formal recognition of academic achievement and the substantive intellectual competence that such recognition is intended to represent.

The legal-philosophical foundations of the present study were developed in the author's earlier works. In *Artificial Intelligence Is Slave of the Lamp*, the doctrine of *instrumentum vocale* was employed to distinguish technologically sophisticated instruments from legal subjects capable of independent will and responsibility.[1] *LLM Is Not Artificial Intelligence* extended this distinction to the epistemological status of large language models, examining the difference between linguistic performance, understanding, and independent judgment.[2] The subsequent joint study, *Instrumentum Vocale and the Architecture of Responsibility: From Liability to Embedded Accountability*, developed an institutional framework of verification, attribution, and non-delegable human responsibility.[3] The present article applies these foundations to academic knowledge production, scholarly authorship, independent verification, and the institutional certification of intellectual competence.

The analysis does not undertake original survey research, claim to establish a universal causal relationship between artificial intelligence and cognitive decline, or seek to identify individual political or institutional actors as personally responsible for a presumed global collapse of education. Its purpose is to evaluate documented institutional developments,

[1] Kildeev, A. (2025a). Artificial Intelligence Is Slave of the Lamp. Scientific Platform XXI Century, No. 6 (September 2025), 3–31. Published journal version. https://episs.ru/

[2] Kildeev, A. (2025b). LLM Is Not Artificial Intelligence. Scientific Platform XXI Century, No. 7 (October 2025), 3–44. Published journal version. https://episs.ru/

[3] Liebig, T., & Kildeev, A. (2026). Instrumentum Vocale and the Architecture of Responsibility: From Liability to Embedded Accountability. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(1), 11–84. Published journal version. https://crail.ca/

identify plausible mechanisms affecting the reproduction of knowledge, and distinguish demonstrable findings from interpretive conclusions.

The broader question is therefore not whether modern universities produce more publications, award more qualifications, or employ more advanced technologies than their predecessors. It is whether the institutions responsible for certifying knowledge continue to preserve the human intellectual capabilities necessary to justify that certification.

The survival of academia as an institutional system should not be confused, without further examination, with the preservation of scholarship as an independent intellectual practice.

Chapter 2. The University, Scholarly Apprenticeship, and the Continuity of Knowledge

2.1. The Historical Foundations of the University

The contemporary university did not emerge as an institution primarily designed to issue academic qualifications, generate publications, or satisfy the immediate demands of the labour market. Its historical origins lie in the organization of advanced learning, the preservation and interpretation of accumulated knowledge, and the preparation of individuals capable of exercising intellectual and professional judgment.

The European university developed during the Middle Ages through communities of teachers and scholars whose activities gradually acquired institutional recognition and legal privileges. In England, teaching at Oxford can be traced to 1096, while Cambridge emerged as a scholarly community in 1209. These institutions inherited and developed traditions of advanced learning already established in continental Europe. Their educational activities included the liberal arts, theology, law, and medicine.[4]

The medieval university differed fundamentally from its modern counterpart. Its intellectual organization depended heavily upon authoritative texts, lectures, commentaries, and formal disputations. Nevertheless, an important feature of this educational tradition deserves emphasis: academic preparation was not confined to the passive reception of information.

At Cambridge, for example, medieval examinations involved oral disputations in which candidates defended propositions against objections raised by other scholars and their teachers. Academic qualifications were connected to participation in recognized intellectual procedures rather than merely to the completion of a prescribed period of attendance.

Such practices did not guarantee intellectual independence. Medieval universities operated within religious, political, and social hierarchies that frequently constrained inquiry. They were neither universally accessible nor institutions of unrestricted intellectual freedom.

[4] University of Oxford. Our History. https://www.ox.ac.uk/about/the-university/history ; University of Cambridge. The Medieval University. https://www.cam.ac.uk/about-the-university/history/the-medieval-university

Their historical significance lies elsewhere: they established durable organizational mechanisms through which knowledge, disciplinary methods, standards of argumentation, and scholarly qualifications could be transmitted between generations.

During the early modern period, the development of experimental science introduced additional principles of inquiry. The establishment of the Royal Society in England in 1660, followed by the adoption of its motto *Nullius in verba* in 1662, expressed a commitment to examining scientific propositions through observation and experiment rather than accepting them solely upon the authority of established figures.[5]

The Royal Society was not itself a university. Its significance for the present discussion lies in demonstrating that the development of modern scientific inquiry required more than the institutional preservation of received knowledge. It demanded procedures through which established propositions could be questioned, tested, and potentially rejected.

This distinction between the possession of knowledge and the capacity to examine its foundations would become essential to the subsequent development of the research university.

2.2. The University Tradition in North America

The transplantation of European educational institutions to North America produced different national trajectories, but the intellectual influence of English university traditions remained substantial.

Harvard College, founded in Massachusetts in 1636, emerged within the institutional and religious environment of English colonial society. Its early educational purposes reflected the preparation of clergy and educated members of colonial communities. The institution did not begin as a research university in the modern sense. Its later transformation formed part of a much broader evolution of American higher education.[6]

In Canada, British institutional traditions were reflected in the establishment of King's College in Upper Canada by royal charter in 1827. The institution subsequently became the University of Toronto, with its development shaped by disputes concerning religious authority, public education, and the secularization of higher learning.

Canadian university history, however, cannot be understood exclusively through English institutional inheritance. The French-language educational tradition developed through institutions such as the Séminaire de Québec, founded in 1663, from which Université Laval emerged in 1852.[7]

These historical differences are relevant because they demonstrate that the universities of England, the United States, and Canada did not develop according to a single institutional model.

---

[5] The Royal Society. History of the Royal Society. https://royalsociety.org/about-us/who-we-are/history/

[6] Harvard University. The History of Harvard University. https://www.harvard.edu/about/history/

[7] University of Toronto Archives & Records Management Services. The University's Original Charter. https://utarms.library.utoronto.ca/history/university-original-charter ; Université Laval. Origin and History. https://www.ulaval.ca/en/about-us/origin-and-history

The modern research university itself emerged through a combination of intellectual traditions and institutional innovations. During the nineteenth century, the relationship between university instruction and original scientific investigation became increasingly important. Research was no longer understood merely as an activity conducted by individual scholars alongside teaching; it progressively became an institutional responsibility.

The foundation of Johns Hopkins University in Baltimore in 1876 represented an important stage in this development. Its educational philosophy emphasized original investigation, advanced scholarly preparation, and the development of researchers capable of extending existing knowledge.[8]

The significance of this transformation cannot be reduced to the creation of specialized laboratories or graduate programmes. It concerned the intellectual identity of the university. The institution increasingly assumed responsibility not only for transmitting established knowledge but also for sustaining the human and organizational conditions required to produce new knowledge.

By the late twentieth century, universities in England, the United States, and Canada had inherited institutions shaped by centuries of teaching, scholarly argumentation, professional preparation, research, and intellectual exchange. This inheritance was neither homogeneous nor immune to criticism. Nevertheless, it provides the historical context for examining the transformations that occurred during the period addressed by the present article.

2.3. Scientific Schools, Apprenticeship, and the Formation of Scholars

The reproduction of scientific knowledge involves more than the preservation of books, archives, experimental results, and published research.

It also requires the transmission of the intellectual capacities necessary to understand, evaluate, and extend that knowledge. This distinction is particularly important in considering scientific schools and academic apprenticeship.

For the purposes of the present study, a *scientific school* is not understood as a formally designated institution or a group united by unquestioning adherence to the authority of a particular professor. Rather, it denotes a continuing intellectual community in which research problems, methodological standards, disciplinary practices, and forms of scholarly judgment are developed and transmitted across generations.

Such communities may exist within laboratories, departments, research institutes, doctoral programmes, and networks of scholarly collaboration. Their defining feature is intellectual continuity. The preparation of a researcher therefore involves a process fundamentally different from the simple acquisition of information. An aspiring scholar must learn to formulate research questions, distinguish significant problems from trivial ones, identify relevant sources, evaluate competing explanations, and recognize the limitations of available evidence.

[8] Johns Hopkins University. History & Mission. https://www.jhu.edu/about/history/

These capacities cannot be acquired solely through the memorization of established propositions.

Michael Polanyi's conception of *tacit knowledge* provides a philosophical foundation for understanding this process. In *The Tacit Dimension*, Polanyi emphasized that human knowledge includes practical and intellectual capacities that cannot be completely reduced to explicit verbal formulations or codified instructions.[9]

Applied to scholarship, this distinction suggests that familiarity with methodological rules does not necessarily confer the capacity to exercise sound methodological judgment. A researcher may know the formal requirements for scientific citation without understanding the evidentiary significance of a primary source. A doctoral candidate may reproduce accepted theoretical propositions without possessing the competence to challenge their foundations. Similarly, the ability to describe an experimental procedure does not necessarily establish the capacity to evaluate an unexpected result.

Jean Lave and Etienne Wenger's theory of *situated learning* provides a complementary perspective. Their concept of legitimate peripheral participation explains how newcomers acquire knowledge and competence through progressively fuller participation in communities of practice.[10]

Within academic research, such participation may involve seminars, laboratory work, supervised investigations, scholarly correspondence, conferences, critical discussion, and independent publication. The significance of these activities lies not merely in their formal completion, but in the gradual development of intellectual responsibility.

Research supervision, when properly exercised, is therefore not simply an administrative procedure for ensuring that a thesis satisfies institutional requirements. It is a form of intellectual apprenticeship. Its purpose is to develop the researcher's capacity to reach and defend conclusions through independent examination of evidence.

This understanding is supported by the United States National Academies' report *The Science of Effective Mentorship in STEMM* (2019), which identifies mentoring as an important component of research training and professional development. The report also emphasizes that effective mentorship requires institutional attention and cannot be assumed to arise automatically from the existence of academic hierarchies.[11]

It would nevertheless be mistaken to idealize traditional scholarly apprenticeship. The dependence of junior researchers upon established professors may reproduce intellectual conformity, discourage criticism, or reinforce institutional inequalities. The existence of a scientific school does not, by itself, guarantee the quality or independence of its members.

---

[9] Polanyi M. The Tacit Dimension. Garden City, NY: Doubleday, 1966. Reissued by the University of Chicago Press, 2009.

[10] Lave J., Wenger E. Situated Learning: Legitimate Peripheral Participation. Cambridge: Cambridge University Press, 1991. DOI: 10.1017/CBO9780511815355.

[11] National Academies of Sciences, Engineering, and Medicine. The Science of Effective Mentorship in STEMM. Washington, DC: National Academies Press, 2019. DOI: 10.17226/25568.

The decisive criterion is whether the educational relationship develops the intellectual autonomy of the succeeding generation.

A scientific school fulfils its purpose not when its students faithfully reproduce the conclusions of their teachers, but when they acquire the competence to question, improve, and, where necessary, refute them.

2.4. Fundamental Research and the Institutional Reproduction of Knowledge

The relationship between fundamental research and scholarly education constitutes another essential dimension of the university's intellectual function.

Fundamental research is characterized by inquiry directed toward the development of knowledge whose ultimate significance, applications, or practical consequences may not be immediately predictable.

Its value cannot always be measured through short-term economic returns, commercial applications, or predetermined institutional performance targets.

This does not mean that fundamental research should be exempt from evaluation or public accountability. It means that the criteria appropriate to its assessment must reflect the uncertainty, complexity, and potentially long-term significance of scientific discovery.

Universities occupy a distinctive position in this process because they combine research activity with the preparation of future researchers.

The relationship between these functions has repeatedly been emphasized in educational policy. In the United States, the Boyer Commission's 1998 report, *Reinventing Undergraduate Education: A Blueprint for America's Research Universities*, argued that research universities should integrate the process of discovery more directly into undergraduate education. Its recommendations addressed inquiry-based learning, research participation, and the development of students as active participants in intellectual investigation.[12]

Importantly, the report identified institutional shortcomings that already existed during the 1990s. It therefore provides historical evidence that concerns about the relationship between university research and the quality of education long preceded the widespread deployment of generative artificial intelligence.

In Canada, the Natural Sciences and Engineering Research Council's Discovery Grants programme reflects an institutional model in which sustained research activity and the training of highly qualified personnel are treated as interconnected objectives. The programme supports continuing research agendas and recognizes research training as a relevant element in assessing scholarly activity.[13]

---

[12] Boyer Commission on Educating Undergraduates in the Research University. Reinventing Undergraduate Education: A Blueprint for America's Research Universities. Stony Brook, NY: State University of New York at Stony Brook, 1998. https://repo.library.stonybrook.edu/xmlui/handle/11401/57972

[13] Natural Sciences and Engineering Research Council of Canada. Discovery Grants Program (Individual). Official programme documentation. https://nserc.canada.ca/en/funding-opportunity/discovery-grants-program-individual

In England, Research England's quality-related research funding provides institutional support that universities may allocate toward research priorities, infrastructure, and postgraduate research supervision. This arrangement forms part of a broader research-funding structure combining institutional resources with competitively awarded project funding.[14]

These examples do not establish that the three systems provide equivalent levels of financial support or achieve comparable educational outcomes.

Nor do they demonstrate that contemporary research funding necessarily protects independent intellectual inquiry. Their significance is more limited and more precise: they confirm that the institutional relationship between the production of knowledge and the preparation of future researchers remains an explicit objective of university research policy.

The effectiveness of that relationship must be evaluated separately. A university may expand its research budget while reducing the time available for meaningful supervision. It may develop commercially successful research partnerships while placing increasing pressure upon scholars to pursue projects capable of attracting external funding.

Conversely, institutional reforms may strengthen research infrastructure, broaden participation, and create new opportunities for scholarly collaboration. These possibilities require empirical examination rather than assumptions concerning inevitable improvement or decline.

2.5. Intellectual Continuity as an Institutional Responsibility

The historical development of the university reveals a persistent distinction between the accumulation of knowledge and the reproduction of the human competence necessary to sustain it.

Libraries may preserve texts without ensuring that succeeding generations understand their significance. Research repositories may accumulate publications without establishing the originality or reliability of their contents. Universities may confer qualifications without necessarily demonstrating that every recipient has developed the capacity for independent scholarly judgment.

These observations do not constitute evidence of universal academic deterioration.

They identify a structural problem requiring investigation. The survival of scholarly knowledge depends upon the continuity of intellectual practices through which propositions are examined, sources are verified, methodological decisions are justified, and conclusions remain open to criticism.

This continuity cannot be established solely by the number of academic institutions, students, degrees, or published articles. Such indicators measure important dimensions of educational activity. They do not independently measure the quality of scholarly judgment.

[14] UK Research and Innovation / Research England. How We Fund Higher Education Providers. https://www.ukri.org/publications/research-england-how-we-fund-higher-education-providers/how-we-fund-higher-education-providers/

Accordingly, the present investigation distinguishes between three related but analytically separate processes: the expansion of educational institutions, the production of formal academic outputs, and the reproduction of independently competent scholars.

The historical discussion in this chapter establishes the intellectual and institutional background of these processes. It does not extend the empirical research beyond the period 1990–2026 or beyond the three selected educational systems: England, the United States, and Canada.

The quantitative evidence concerning their transformation is presented in Exhibits A and B, while the methodological limitations and boundaries of inference are specified in Exhibit C.

The central question is not whether universities at the beginning of the 1990s were necessarily superior to their contemporary counterparts. It is whether subsequent changes in university organization, financing, research incentives, and educational practice have strengthened or weakened the conditions required for independent scholarly judgment.

The capacity of an institution to produce academic qualifications and publications must not be confused with its capacity to reproduce independent scholars.

This distinction provides the conceptual foundation for the following examination of educational expansion, institutional commercialization, changing research incentives, and the political economy of contemporary academia.

Chapter 3. From Knowledge to Academic Production

3.1. The Transformation of the University: From Intellectual Institution to Measurable Enterprise

The historical development examined in the preceding chapter established the university as an institution whose responsibilities extend beyond the transmission of established information. The preservation of scholarly traditions, the cultivation of independent judgment, the preparation of successive generations of researchers, and the pursuit of fundamental knowledge constitute interdependent elements of its intellectual mission.

During the period 1990–2026, universities in the United States, Canada, and England experienced substantial changes in their organization, financing, governance, and systems of evaluation.

These transformations occurred under different national conditions. They cannot be reduced to a single policy programme or explained exclusively through the influence of market economics. Nevertheless, they raise a common question: how does an institution devoted to producing and transmitting knowledge preserve its intellectual purpose when its performance is increasingly expressed through measurable administrative and economic indicators?

The distinction is not between universities that operate with financial resources and those that do not. Universities have always required funding, institutional administration, and mechanisms for evaluating educational performance.

Rather, the issue concerns the relationship between institutional means and intellectual ends. Enrolment numbers, degrees awarded, research grants obtained, publications produced, citation indicators, patents, and external partnerships may all provide valuable information about university activity. Yet they describe different aspects of institutional performance and cannot be treated as interchangeable measures of educational or scientific quality.

A university's capacity to attract students does not independently demonstrate the quality of its teaching. Its research income does not establish the originality of its discoveries. Its publication volume does not necessarily reveal whether its scholars have advanced understanding of a particular problem.

The danger arises when indicators intended to inform institutional judgment begin to replace that judgment. This distinction provides the analytical foundation for examining the contemporary movement from knowledge as an intellectual responsibility toward academic production as a system of measurable outputs.

3.2. The Changing Political Economy of Higher Education, 1990–2026

The transformation of university financing represents one of the most significant institutional developments of the period under examination.

In the United States, higher education operates through a complex combination of state universities, private nonprofit institutions, community colleges, and other educational providers. Their revenue structures vary substantially, as do their research activities, governance arrangements, and dependence upon tuition.

Consequently, American higher education cannot be described accurately through a single national funding model.

Nevertheless, the State Higher Education Executive Officers Association's *State Higher Education Finance* reports document the long-term relationship between state appropriations, public university enrolment, and net tuition revenue. The fiscal year 2025 report provides historical financial series extending to 1980, allowing developments within the present study's 1990–2026 period to be examined against a consistent institutional framework.[15]

These records demonstrate the importance of distinguishing aggregate public expenditure from funding available per student, and gross tuition charges from net tuition revenue. Changes in enrolment, financial aid, institutional costs, and state-level policy may alter the interpretation of apparently straightforward funding trends.

The American case therefore illustrates a broader methodological problem: the growth of educational expenditure does not, without additional evidence, establish an equivalent expansion of resources available for independent academic work.

Canada presents a different institutional structure. Universities operate predominantly within systems of provincial responsibility, supplemented by federal research funding,

[15] State Higher Education Executive Officers Association. State Higher Education Finance: FY 2025. Published 2026. https://shef.sheeo.org/report/

institutional revenues, philanthropy, and other sources. Statistics Canada recorded that student fees represented approximately 12% of Canadian university revenues in 1990/91, while total government sources accounted for approximately 69%. By 2002/03, student fees had reached 20.5% of revenues. More recent administrative data reported tuition revenues amounting to 30.1% of total university revenues in 2023/24.[16]

These figures provide evidence of an important long-term change in the composition of university financing. They should not, however, be interpreted as a completely harmonized time series without further examination of historical accounting definitions, institutional coverage, and revenue classifications.

The Canadian experience also reveals substantial provincial variation. Quebec, Ontario, British Columbia, and the Atlantic provinces have followed different financing trajectories, making national averages informative but insufficient for explaining the circumstances of particular universities.

England experienced especially visible legislative and financial restructuring.

The Further and Higher Education Act 1992 reorganized important elements of higher education governance and enabled former polytechnics to obtain university status. This change must be understood as an institutional reclassification and transformation, rather than automatically treated as evidence of the creation of equivalent numbers of new research universities.[17]

The introduction of undergraduate tuition fees in 1998, subsequent changes in 2006, and the increase of the standard upper fee limit to £9,000 in 2012 significantly altered the relationship between students, universities, and public financing.

Later reforms established the Office for Students as the principal regulator of higher education providers in England, reflecting the changed regulatory environment associated with tuition finance, institutional competition, and reduced dependence upon direct teaching grants.[18] The financial consequences remain visible in contemporary data.

A September 2026 House of Commons Library briefing recorded that tuition fees had become a substantially larger component of higher education providers' income, while the relative contribution from direct funding-body grants had declined. The same report documented financial deficits among a significant proportion of publicly funded English universities in 2024/25.[19]

The broader financial series published with the parliamentary briefing include United Kingdom aggregates. Such series must not be presented as exclusively English observations.

---

[16] Statistics Canada. University Finances, 2002/03. The Daily, 19 August 2004; University Revenues Increase More Than Expenditures in 2023/2024. The Daily, 29 July 2025.

[17] Further and Higher Education Act 1992, c. 13. https://www.legislation.gov.uk/ukpga/1992/13/contents

[18] House of Commons Library. Tuition Fees in England: History, Debates, and International Comparisons. Research Briefing CBP-10155, 23 February 2026; Higher Education and Research Act 2017, c. 29.

[19] Bolton, P., and Lewis, J. Higher Education Finances and Funding in England. House of Commons Library, Research Briefing CBP-10037, 4 September 2026.

The present study distinguishes England-specific evidence from UK-wide contextual statistics.

Across the three countries, the resulting institutional arrangements differ considerably. What connects them is not an identical process of privatization or the disappearance of public funding, but the increasing importance of complex revenue structures, financial sustainability, external accountability, and institutional competition. Whether these changes improved or weakened scholarly competence must be investigated rather than presumed.

3.3. Academic Capitalism and the Commercialization of Knowledge

The transformation of university financing has been accompanied by changes in the relationship between academic institutions and the wider economy.

Sheila Slaughter and Gary Rhoades developed the concept of *academic capitalism* to examine how universities participate in markets for educational services, research products, intellectual property, and commercially valuable knowledge. Their work emphasizes that market-oriented academic activity is not necessarily external to public institutions; it can develop through relationships involving universities, governments, commercial enterprises, and research organizations.[20]

Derek Bok, in *Universities in the Marketplace* (2003), approached the commercialization of higher education through its implications for academic values, research independence, educational priorities, and institutional integrity. His argument was not that universities should be isolated from economic activity, but that financial opportunities create conflicts requiring explicit safeguards.[21]

These contributions identify a fundamental distinction between the economic value of scientific knowledge and the intellectual responsibilities of the institutions producing it.

The commercial application of discoveries may provide substantial public benefits. University research can contribute to medical innovation, technological development, industrial productivity, and economic growth. External partnerships may provide facilities, financial resources, and opportunities that would otherwise be unavailable.

Consequently, commercialization cannot be treated as inherently incompatible with academic freedom or scientific progress. The relevant question is whether commercial objectives remain subordinate to the standards governing independent inquiry.

A research institution may legitimately cooperate with private industry while protecting investigators' freedom to publish inconvenient findings. It may develop intellectual property without allowing proprietary interests to determine the validity of scientific conclusions. Conversely, difficulties arise when the terms of financing, contractual restrictions, reputational interests, or anticipated commercial returns influence the selection, interpretation, or disclosure of research results.

[20] Slaughter, S., and Rhoades, G. Academic Capitalism and the New Economy: Markets, State, and Higher Education. Johns Hopkins University Press, 2004.

[21] Bok, D. Universities in the Marketplace: The Commercialization of Higher Education. Princeton University Press, 2003.

These are questions of governance and institutional responsibility, not merely questions of public versus private ownership. The university's scientific authority derives from the quality and independence of its research, not from the commercial value of its discoveries. The distinction becomes particularly important when universities increasingly evaluate departments and researchers through indicators associated with external income, intellectual property, publication activity, and institutional reputation.

3.4. The Academic Production System: Publications, Grants, and Performance Indicators

Scientific publications constitute an indispensable means of communicating research.

They permit findings to be examined, criticized, replicated, extended, and incorporated into continuing scholarly discussion. Without publication, scientific knowledge would lose one of its principal mechanisms of intellectual continuity.

The problem is therefore not publication itself, nor the legitimate expectation that publicly supported researchers should account for their activities. It arises when the numerical production of academic outputs becomes a substitute for the substantive evaluation of research.

A publication count establishes that a certain number of documents have appeared. It does not independently establish their originality, methodological reliability, or importance. Citation counts provide information about patterns of attention and scholarly reference. They may be influenced by disciplinary practices, research-community size, the age of a publication, methodological controversies, or the visibility of particular topics.

Journal reputation may provide contextual information about editorial processes, but it cannot substitute for examination of an individual article. These distinctions have been explicitly recognized within the scientific community.

The *San Francisco Declaration on Research Assessment* (DORA), initiated in 2012, challenged the use of journal-based indicators as proxies for the quality of individual articles or researchers. Its recommendations address hiring, promotion, funding, and institutional research evaluation.[22]

The *Leiden Manifesto for Research Metrics*, published in *Nature* in 2015, similarly emphasized the need for quantitative indicators to support expert assessment rather than replace it. It called attention to disciplinary differences, transparency, verification, and the risks associated with poorly designed evaluation practices.[23]

These initiatives are important for the historical argument of the present article. They demonstrate that concerns about excessive reliance upon academic performance indicators were articulated well before generative artificial intelligence became widely available.

[22] San Francisco Declaration on Research Assessment (DORA), 2012/2013. https://sfdora.org/read/

[23] Hicks, D., Wouters, P., Waltman, L., de Rijcke, S., and Rafols, I. "Bibliometrics: The Leiden Manifesto for Research Metrics." Nature, 520 (2015): 429–431. https://doi.org/10.1038/520429a

The problems identified were therefore not created by language models. They arose within institutional systems already attempting to reconcile research quality, accountability, financial competition, and measurable performance.

England offers a particularly relevant example of formalized research evaluation. The Research Excellence Framework, first conducted in 2014, replaced the earlier Research Assessment Exercise and became a major mechanism for evaluating university research. Its results inform the allocation of public research funding. The framework covers the United Kingdom, although Research England administers the relevant funding arrangements for English institutions.[24]

Such systems serve legitimate public purposes, including accountability, the identification of research strengths, and the distribution of limited resources. Their existence does not demonstrate that British universities have abandoned scholarly judgment.

Nevertheless, any assessment framework may influence institutional priorities, career incentives, the organization of research, and decisions about which activities receive financial or administrative support. The relevant question is whether research assessment preserves the distinction between evidence of academic activity and judgment concerning academic quality.

An institution may be highly successful in generating measurable research outputs without necessarily demonstrating an equivalent capacity for original intellectual discovery. The present investigation therefore treats publication volume, research income, and evaluation indicators as evidence of institutional activity rather than direct measures of scientific competence.

3.5. Academic Labour, Mentorship, and Institutional Time

The transformation of university governance and financing must also be considered in relation to academic employment. Research and teaching depend upon human labour. The preparation of new scholars requires sustained supervision, intellectual discussion, feedback, and opportunities for participation in research.

These activities consume resources and, importantly, time. Changes in academic employment arrangements may therefore have consequences extending beyond employment policy.

In the United States, a 2023 analysis published by the American Association of University Professors, drawing upon National Center for Education Statistics data, documented the extensive use of contingent academic appointments. Its findings showed that approximately 68% of faculty members in American colleges and universities held contingent appointments in 2021.[25]

[24] UK Research and Innovation. Research Excellence Framework. Official documentation, updated 2026. https://www.ukri.org/who-we-are/research-england/research-excellence/research-excellence-framework/

[25] Colby, G. “Data Snapshot: Tenure and Contingency in US Higher Education.” Academe, American Association of University Professors, Spring 2023.

The figure refers to a broad higher education workforce and does not describe the composition of research-university faculty alone. It cannot establish the quality of supervision or research in particular institutions.

Nevertheless, it identifies an institutional development relevant to the present inquiry: a substantial proportion of teaching is performed under employment arrangements differing from the traditional tenured academic appointment.

In Canada, historical data published by Statistics Canada likewise recorded changes in the composition of the university teaching workforce during the 1990s, including an increase in part-time faculty alongside a reduction in full-time faculty between 1990/91 and 1997/98.[26]

These observations do not justify assuming that temporary or part-time academic personnel are less competent than tenured professors.

Employment status is not a measure of intellectual merit. The relevant concern is institutional continuity. Where teaching and research responsibilities are divided among personnel facing different contractual conditions, the university must ensure that supervision, research training, and scholarly mentorship remain adequately supported.

The distinction is particularly significant for scientific schools. Such communities depend upon more than individual intellectual achievement. They require continuity of research programmes, the progressive formation of junior scholars, and institutional opportunities for sustained inquiry.

A university may employ accomplished researchers while failing to provide adequate conditions for intellectual succession. Conversely, stable academic employment does not by itself guarantee effective mentorship. The appropriate question is whether the university's organizational arrangements enable its faculty to perform the intellectual functions for which the institution exists.

3.6. When Institutional Indicators Become Intellectual Objectives

The preceding evidence suggests a distinction between three forms of institutional accountability.

Financial accountability concerns the legitimate management and use of resources.

Administrative accountability concerns compliance with institutional and regulatory obligations.

Epistemic accountability concerns the justification of scholarly claims through appropriate methods, evidence, criticism, and verification.

All three are necessary.

They are not interchangeable.

An institution can satisfy accounting requirements while producing poor research. A publication may meet formal editorial requirements while containing unsupported

[26] Statistics Canada. "Part-time University Faculty, 1990/91 to 1997/98." The Daily, 8 May 2002.

conclusions. A research project may receive substantial external funding without producing significant intellectual advances.

Equally, important scientific findings may emerge from relatively modest projects whose value was not anticipated by their original financial or administrative assessments. The danger is not measurement itself.

It is the transformation of an indicator of scholarly activity into the governing objective of scholarship.

When universities are evaluated principally through what can be counted, their internal priorities may increasingly favour activities that generate measurable evidence of institutional success.

Such an adjustment may be rational from an administrative perspective without necessarily being optimal for the long-term development of scientific knowledge. This is an institutional hypothesis, not a demonstrated universal outcome. Different disciplines, universities, and funding systems may respond differently to similar incentives. Quantitative evaluation can expose previously hidden inequalities, identify ineffective programmes, and improve the accountability of public institutions.

The analytical task is therefore to determine the conditions under which measurement supports scholarly judgment and those under which it displaces it. That distinction becomes particularly important in examining academic publishing. The increasing capacity to produce documents, references, summaries, and apparently sophisticated arguments does not necessarily correspond to an increasing capacity to establish the truth of their contents.

The relevant problem existed before contemporary generative artificial intelligence. Technological developments may nevertheless alter its scale, speed, and institutional consequences. The present chapter does not yet attempt to establish those consequences. It identifies the academic environment into which these technologies were introduced.

3.7. From the Production of Knowledge to the Production of Academic Outputs

Between 1990 and 2026, the universities of the United States, Canada, and England underwent substantial institutional transformations.

The available evidence documents changes in educational financing, student contributions, research evaluation, academic employment, and relationships with commercial and governmental organizations. These changes were not uniform. Nor do they establish that contemporary universities have collectively abandoned their intellectual responsibilities. Some developments expanded educational access, strengthened research infrastructure, promoted public accountability, and created new opportunities for scientific collaboration. Others introduced or intensified questions concerning financial dependence, employment stability, institutional competition, and the relationship between measurable performance and scholarly competence.

The distinction established in Chapter II therefore acquires a new significance. A university may preserve the institutional appearance of scholarship while reorganizing the conditions under which scholarship is performed. It may increase the number of

qualifications awarded, publications produced, or research projects administered without thereby establishing that it has improved its capacity to develop independent scholars.

The central question is whether the university remains primarily an institution in which intellectual judgment is cultivated, tested, and transmitted, or whether those functions increasingly compete with objectives associated with measurable academic production.

The production of academic outputs is not identical to the production of knowledge; and the production of knowledge is not identical to the reproduction of independent scholars.

This threefold distinction will become essential in examining the relationship between academic publishing, institutional responsibility, and generative artificial intelligence. The challenge posed by contemporary technologies is not simply that academic documents can now be produced more rapidly. It is whether universities already organized around measurable outputs possess the institutional capacity to distinguish intellectual production from its simulation. That question forms the basis of the next stage of the investigation.

Chapter 4. The Abdication of Academic Judgment

The transformation of universities into increasingly sophisticated systems of academic production raises a question that cannot be answered by examining publication statistics, research budgets, or institutional rankings alone. What happens to scholarship when the production of its visible results becomes progressively detached from the intellectual processes that once gave those results their meaning? A university may publish more articles, award more degrees, attract greater investment, and expand its digital infrastructure without necessarily strengthening the capacity of its students and researchers to think independently. The difficulty lies in distinguishing the measurable products of academic activity from the less readily measurable qualities that make such activity intellectually valuable.

This distinction has acquired particular importance with the arrival of generative artificial intelligence. Large language models can now produce texts that resemble scholarly arguments, organize extensive bibliographies, formulate hypotheses, compare theoretical positions, and present their conclusions in language sufficiently polished to satisfy many of the formal conventions of academic writing. These capabilities are useful and, in competent hands, can substantially assist research. Yet the sophistication of the resulting document tells us remarkably little about the intellectual contribution of the person who submits it. The outward appearance of scholarship has become easier to reproduce than the independent judgment upon which scholarship ultimately depends.

There is no reason to idealize the academic past. Intellectual mediocrity, professional conformity, questionable research practices, and the pursuit of recognition without corresponding achievement are hardly inventions of the digital age. Universities have never been inhabited exclusively by brilliant scholars, nor have academic titles invariably reflected exceptional intellectual merit. What has changed is the technological environment in which scholarly work is produced and circulated. Generative systems can reduce the time required to assemble a plausible academic manuscript, potentially allowing the weaknesses of an already imperfect publication culture to manifest themselves on a substantially greater scale.

The central concern, therefore, is not whether scholars should use artificial intelligence. The author of this study openly acknowledges the use of such instruments in scholarly preparation, textual analysis, and editorial work. The relevant distinction concerns the intellectual functions that remain genuinely human. A researcher who employs an LLM to improve the clarity of an argument already developed and independently verified is engaged in a very different activity from someone who accepts a generated argument without understanding its assumptions, testing its evidence, or examining its conclusions. Both may submit polished manuscripts. Only the actual process of scholarly work can establish what intellectual contribution stands behind them.

This is where the problem of academic production becomes a problem of academic judgment. Once the appearance of intellectual accomplishment can be produced independently of the intellectual accomplishment itself, the university faces a challenge that extends beyond plagiarism, inaccurate citations, or violations of publication ethics. It must determine whether its established procedures can still distinguish genuine scholarship from its increasingly convincing imitation.

4.1. Cognitive Offloading and the Erosion of Independent Inquiry

Human beings have never developed knowledge without instruments. Writing extended the capacities of memory, libraries preserved the intellectual achievements of previous generations, and printing made knowledge available to audiences unimaginable in antiquity. Calculators relieved mathematicians and scientists of repetitive operations, while computers transformed the possibilities of scientific modelling, information retrieval, and experimental analysis. Each innovation changed the relationship between human thought and the instruments through which that thought was expressed. Yet none of these developments made the intellectual activity of the researcher irrelevant. The value of an instrument depended upon the questions being asked, the methods employed, and the capacity of a human being to interpret and evaluate the resulting information.

The use of external resources to reduce mental effort is commonly described as *cognitive offloading*. In their influential review, Risko and Gilbert examined this phenomenon as a normal feature of human cognition, emphasizing that individuals regularly employ external actions and resources to reduce the cognitive demands of particular tasks. Their analysis also demonstrates the importance of metacognition: people make judgments about their own abilities when deciding whether to perform an operation themselves or entrust it to an external resource. Such decisions may be entirely rational, although inaccurate assessments of one's own competence can also produce inappropriate reliance on external assistance.[27]

There is consequently little justification for treating cognitive offloading as an intellectual deficiency in itself. A scientist who uses statistical software does not thereby abandon statistical reasoning, any more than a historian who consults a digital archive relinquishes historical judgment. What matters is whether the researcher retains the understanding necessary to formulate the problem, select an appropriate method, recognize

[27] Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. Trends in Cognitive Sciences, 20(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

its limitations, and evaluate the results. The relationship becomes more complicated when the instrument begins to perform functions that previously belonged to the researcher's own reasoning, particularly the formulation of explanations, the organization of arguments, and the evaluation of competing interpretations.

Generative artificial intelligence introduces precisely this complication. Unlike conventional search engines or bibliographic databases, large language models do not merely retrieve information that the researcher must subsequently organize. They can produce complete argumentative structures, formulate plausible hypotheses, summarize disagreements, propose methodological approaches, and generate conclusions in polished academic language. Such capabilities can considerably improve research efficiency when used by an intellectually competent investigator. They can also create an environment in which the visible results of research appear before the intellectual work ordinarily associated with producing them has taken place.

Consider a graduate student beginning a dissertation in an unfamiliar field. Traditionally, the initial months of research would involve reading foundational works, discovering competing schools of thought, identifying unresolved questions, and gradually developing a defensible position. The process would almost certainly include misunderstandings, abandoned hypotheses, incomplete bibliographies, and repeated revisions. None of these difficulties would be particularly attractive to a university administrator concerned primarily with measurable productivity. Yet they would constitute an important part of the student's intellectual formation. By encountering and resolving difficulties, the student would gradually acquire the ability to distinguish an established proposition from a disputed interpretation, a promising hypothesis from an unsupported assertion, and an original contribution from the repetition of familiar ideas.

A generative system can now provide the same student with an apparently coherent research proposal, a theoretical framework, a provisional bibliography, and a sequence of conclusions within minutes. The result may be useful, particularly when treated as a collection of suggestions requiring independent examination. But if the student adopts the proposed framework before developing sufficient knowledge to evaluate it, the normal relationship between intellectual development and academic production becomes inverted. Instead of producing a scholarly argument through research, the student begins with an argument whose intellectual foundations remain to be investigated. The completed-looking document may create an illusion of progress even when the underlying understanding has scarcely advanced.

This distinction between receiving an answer and developing an understanding lies at the centre of Alejandro Sinopoli's *EPISTEMÍA: Governing the Human Process of Understanding* (2026).[28] Sinopoli approaches the problem primarily from legal philosophy and epistemology, arguing that the process through which human beings formulate

[28] Sinopoli, A. (2026). EPISTEMÍA: Governing the human process of understanding: A proposal for a new object of legal protection in the age of artificial intelligence. The Journal Scientific Platform XXI Century, No. 5 (14), 40–72. https://episs.ru/wp-content/uploads/2026/09/5-14-ISSUE-2026-English-FINAL.pdf

questions, establish relevance, compare alternatives, revise interpretations, and form judgments possesses significance independent of the correctness of the final result.

Understanding, in his account, cannot be reduced to the possession of accurate information. It develops through an active and revisable process in which the individual participates in the construction and evaluation of meaning.[29]

Sinopoli's contribution is particularly valuable because it draws attention to a feature of generative systems that conventional discussions of automation frequently overlook: the temporal character of human understanding. Human reflection requires opportunities for hesitation, disagreement, reconsideration, and the gradual integration of unfamiliar information. Generative systems, by contrast, can produce answers and elaborate arguments at a speed that bears little relationship to the time required for a person to understand and independently assess them. The resulting difficulty is not simply that machines produce too much information. It is that the process of intellectual evaluation may become detached from the process of textual production.[30]

A researcher may receive a twenty-page analysis in a few minutes, but cannot responsibly verify twenty pages of unfamiliar legal, historical, or scientific argument merely because the material has arrived in a convenient and persuasive form. Checking references may require access to original publications; evaluating legal claims may require examination of legislation and judicial decisions; assessing scientific conclusions may demand familiarity with experimental methods, statistical assumptions, and competing interpretations. Some statements can be checked quickly, while others require substantial independent investigation. The apparent efficiency of automated drafting therefore creates a potential asymmetry between the speed at which scholarly claims can be generated and the human resources required to establish their reliability.

This asymmetry should not be mistaken for a universal mathematical relationship or an experimentally demonstrated law of cognitive decline. Its significance depends upon the complexity of the task, the availability of reliable sources, the competence of the reviewer, and the consequences of error. Certain routine checks can themselves be automated, and appropriately designed tools may improve rather than weaken human verification. Nevertheless, increasing the volume of plausible claims does not automatically increase the human capacity to understand them. Where institutions encourage rapid production without providing corresponding resources for independent evaluation, the imbalance can become a significant organizational problem.

Sinopoli develops another concept that deserves particular attention in this context: *cognitive friction*. He uses the expression to describe forms of productive resistance that prevent the premature closure of understanding. Doubt, disagreement, exposure to alternative explanations, the need to reformulate an argument, and the obligation to justify a conclusion can all contribute to the development of intellectual judgment. Such resistance is not valuable merely because it consumes time.

[29] Sinopoli. EPISTEMÍA, pp.40-45.

[30] Sinopoli. EPISTEMÍA, pp.43-48.

Needless administrative procedures, inaccessible information, and defective research tools can obstruct inquiry without improving understanding. The relevant distinction concerns difficulties that stimulate meaningful intellectual engagement rather than those that simply make academic work slower.[31]

The importance of this distinction becomes apparent when we consider the educational function of research. A student who discovers that an apparently convincing argument contradicts a primary source has encountered something more valuable than a corrected paragraph. The experience may transform their understanding of the subject, reveal the limitations of a particular method, or encourage greater caution in evaluating future claims. A professor who identifies an unexpected weakness in a long-established hypothesis may likewise be compelled to reconsider assumptions that have shaped years of work. Intellectual progress frequently depends upon these moments of resistance, when an explanation ceases to be satisfactory and a more demanding investigation becomes necessary.

Generative systems can assist such processes by introducing counterarguments, identifying overlooked literature, or helping researchers compare competing explanations. They may also reduce intellectual friction by delivering immediately acceptable answers that discourage further investigation. Whether assistance becomes substitution depends substantially upon how the instrument is used and upon the intellectual and institutional conditions surrounding its use. The availability of a rapid answer does not compel its acceptance, but an academic environment increasingly organized around speed, measurable output, and shortened deadlines may make independent examination less attractive and more difficult to sustain.

Empirical research provides reasons to take these concerns seriously, although the evidence requires careful interpretation. In *The Impact of Generative AI on Critical Thinking* (2025), Hao-Ping Lee and colleagues surveyed 319 knowledge workers and examined 936 reported examples of GenAI use in professional activities. They found that greater confidence in the generative system was associated with lower reported engagement in critical thinking, while greater confidence in one's own ability was associated with higher engagement. Importantly, the researchers also observed that generative tools can change the nature of critical-thinking activity, shifting attention from gathering information and solving problems toward checking generated material, integrating responses, and supervising the overall task.[32]

These findings do not establish that generative AI necessarily diminishes intelligence or that its users cease to think critically. The research concerns self-reported professional experience, and the observed relationships do not by themselves demonstrate long-term causal deterioration. Indeed, many respondents reported engaging in critical thinking when

[31] Sinopoli. EPISTEMÍA, pp.49-51.

[32] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (Article 1121, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

using such systems. The more significant implication is that intellectual work may be reorganized rather than simply eliminated.

A person who previously developed an answer through investigation may increasingly spend time evaluating an answer prepared by a machine. Whether this reorganization strengthens or weakens professional competence depends upon the quality of that evaluation and the extent to which the individual continues to exercise independent judgment.

A different perspective emerges from the experimental study *Your Brain on ChatGPT* by Nataliya Kosmyna and colleagues (2025). The researchers examined essay-writing tasks under different conditions, including unaided writing, search-engine assistance, and the use of an LLM. Their investigation combined behavioural analysis with electroencephalographic measurements, involving 54 participants in the initial three sessions and 18 in a fourth session in which selected participants changed conditions. The reported differences in patterns of neural connectivity, recall, and engagement suggest that the cognitive experience of producing a text may vary substantially according to the instruments employed.[33]

The study is important, but its limitations are equally important. It is a research preprint based on a comparatively small experimental sample and a particular kind of writing task. Differences in neural activity during essay composition cannot legitimately be translated into claims that LLM users have suffered a general reduction in intelligence, nor can the findings establish the permanent cognitive consequences of sustained AI use across an entire generation. Their relevance to the present inquiry is narrower and more defensible: the method by which a written product is created may matter to the cognitive activity of its author, even when the final product is formally acceptable.

The problem also has a history considerably older than generative artificial intelligence. In *Ironies of Automation* (1983), Lisanne Bainbridge drew attention to the paradoxical position of human operators in increasingly automated systems. Automation may reduce the need for continuous human intervention while leaving people responsible for exceptional situations in which specialized knowledge and effective judgment are most urgently required. Yet the organization of automated work may itself make it harder to maintain the skills needed to respond to those exceptions.[34] Subsequent research by Parasuraman and Manzey (2010) demonstrated that automation-related complacency and bias can affect both inexperienced and expert users, particularly where competing demands for attention, imperfect decision aids, and habitual reliance on automation influence human performance.[35]

---

[33] Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task (arXiv:2506.08872, version 2). arXiv. Research preprint. https://doi.org/10.48550/arXiv.2506.08872

[34] Bainbridge, L. (1983). Ironies of automation. Automatica, 19(6), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8

[35] Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381–410. https://doi.org/10.1177/0018720810376055

Although these studies were not conducted in the contemporary academic environment of generative LLMs, the underlying human-factors problems are directly relevant. A professor may possess extensive knowledge of a discipline yet lack sufficient time to examine every assertion in a lengthy AI-assisted manuscript. An editor may understand the principles of scholarly verification but face a workload that encourages reliance upon superficial indicators of quality. A junior researcher may accept a generated argument because it resembles the academic literature they have been trained to imitate, without recognizing the assumptions that an experienced specialist would immediately question. In each case, the presence of a human participant provides no guarantee that meaningful intellectual evaluation has occurred.

The implications extend beyond the immediate accuracy of published research. Universities are responsible not only for evaluating existing knowledge but also for cultivating the capacities through which future generations will produce and examine it. If students increasingly complete demanding intellectual tasks through systems whose outputs they cannot independently reconstruct, the institution may continue to award degrees and recognize academic achievement without obtaining adequate evidence that the corresponding competencies have developed. This possibility does not imply that such substitution is universal or unavoidable. It identifies a risk that becomes particularly significant when educational assessment rewards the finished product while paying insufficient attention to the intellectual process through which it was obtained.

The same concern applies to established researchers. Expertise cannot be inferred merely from the ability to produce a well-structured document, just as the possession of a substantial bibliography does not demonstrate familiarity with the works it contains. A scholarly argument acquires credibility through the defensibility of its propositions, the quality of its evidence, the appropriateness of its methods, and the author's capacity to respond to criticism. Generative assistance can contribute to each of these activities when directed and evaluated by a competent researcher. It cannot, by the mere production of coherent language, establish that the person submitting the work has exercised the judgment that academic authorship ordinarily presupposes.

This distinction leads naturally to Sinopoli's subsequent proposal of *Cognitive Dignity* (2026),[36] which moves the discussion from the conditions of human understanding toward the possible normative protection of the person as an active participant in the formation of knowledge. His later analysis distinguishes legitimate assistance and voluntary delegation from situations in which technological or institutional arrangements materially displace the cognitive process while continuing to attribute its results to the human participant. The proposal does not establish that a new autonomous human right has already been recognized, nor does it require individuals to perform every intellectual operation without assistance. Its importance for the present study lies in identifying a problem that conventional measures of

[36] Sinopoli, A. (2026). Cognitive dignity: A proposal for an emerging human right: Protecting the human process of sense-making and knowledge formation against technological substitution. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(2), 89–117. https://www.crail.ca/CRAIL_Vol.1-No.2-2026-October.pdf

productivity and procedural compliance may fail to capture: the preservation of meaningful human participation in activities whose legitimacy depends upon human understanding.[37]

For academic institutions, this question is especially demanding. A university cannot reasonably prohibit intellectual tools without undermining the very tradition of scientific innovation it exists to advance. Yet neither can it assume that the increasing sophistication of those tools automatically strengthens the intellectual capacities of their users. The institutional challenge is to preserve opportunities for independent investigation, criticism, and the acquisition of expertise while taking advantage of the genuine benefits of computational assistance. That challenge requires more than instructions to disclose the use of artificial intelligence, important though such disclosure may be. It requires serious consideration of what universities teach, what they evaluate, and what kind of intellectual achievement their qualifications and publications are intended to certify.

The resulting issue reaches beyond cognitive psychology into the foundations of academic authority. If a machine can formulate an argument, a human can approve it without fully understanding it, and an institution can subsequently recognize the document as scholarship, the ordinary signs of academic accomplishment may remain intact while the intellectual process underlying them becomes increasingly difficult to identify. Before examining the institutional consequences of that possibility, however, it is necessary to consider a more fundamental question: how do language, interpretation, and authority acquire the power to transform an unsupported proposition into something that appears to deserve belief?

4.2. The Humpty Dumpty Academy: Meaning, Authority, and Proof

It is difficult to imagine that Lewis Carroll, when introducing Humpty Dumpty to his readers in 1871, anticipated that his extraordinary creation would eventually become a useful guide to some of the intellectual habits of modern academic institutions. Yet the famous conversation between Alice and the self-confident gentleman perched upon a wall has lost remarkably little of its relevance. Indeed, in an age when almost any proposition can be dressed in sophisticated terminology, furnished with references, and presented as an original theoretical contribution, Humpty Dumpty might feel entirely at home in certain university departments.

The episode is familiar, but its philosophical implications deserve attention. Humpty Dumpty informs Alice that he can make a word mean whatever he chooses. When Alice questions whether words can be assigned such arbitrary meanings, he dismisses the objection with a declaration that reveals the real issue: “The question is,” said Humpty Dumpty, “which is to be master—that’s all”.[38] What begins as an amusing disagreement about language becomes a dispute over intellectual authority. Alice assumes that words belong to a shared system of meaning, within which individual speakers may propose definitions but cannot simply abolish the conditions of mutual understanding. Humpty

[37] Sinopoli. Cognitive dignity, pp. 102-111.
[38] Carroll, L. (1871). Through the Looking-Glass, and What Alice Found There. London: Macmillan (first published December 1871; title-page dated 1872). Chapter VI, “Humpty Dumpty.” Available at: https://www.gutenberg.org/files/12/12-h/12-h.htm.

Dumpty proceeds from the opposite assumption. Meaning, for him, is ultimately a privilege of the person who claims the power to determine it.

There is, of course, nothing objectionable about introducing new terminology. Scientific progress frequently requires the refinement of existing concepts or the creation of entirely new ones. The history of mathematics, physics, medicine, and philosophy would be impossible to reconstruct without acknowledging the importance of conceptual innovation. A researcher who proposes a new definition is not necessarily engaging in intellectual arbitrariness. The decisive question is whether that definition is sufficiently precise, logically consistent, distinguishable from existing concepts, and capable of performing a meaningful explanatory function. A genuine theoretical contribution must do more than rename something already familiar. It must help us understand a problem that previous conceptual instruments could not adequately address.

The distinction becomes less comfortable when the language of innovation begins to substitute for innovation itself. A familiar proposition may acquire an imposing new designation, be placed within an elaborate classification, and emerge as the foundation of an allegedly unprecedented theory. Another author may introduce a concept so broad that it accommodates almost every conceivable example while remaining conveniently immune to criticism. The result can be rhetorically impressive without being intellectually productive. What appears to be a conceptual breakthrough may amount to little more than a redistribution of familiar ideas among newly invented categories.

Such practices are not necessarily fraudulent, and their existence long predates generative artificial intelligence. Academic history offers numerous examples of terminological fashions, speculative systems, and ambitious theories whose explanatory value proved less substantial than their authors initially supposed. Indeed, intellectual experimentation requires a measure of freedom, and even mistaken hypotheses may contribute to the development of knowledge. The problem arises when the ordinary distinction between proposing an interpretation and establishing its validity begins to disappear. At that point, the authority of the author may be invoked to protect the concept from precisely the scrutiny that would establish whether it deserves a place in scholarly discourse.

A remarkable historical illustration of this problem emerges from English jurisprudence rather than philosophy. In *Liversidge v Anderson*, decided by the House of Lords on 3 November 1941 and reported in [1942] AC 206, the judges considered the interpretation of Regulation 18B of the Defence (General) Regulations 1939. The regulation authorized the Home Secretary to order detention where he had reasonable cause to believe that a person fell within specified categories associated with threats to public safety or national defence. The dispute concerned the meaning of those apparently straightforward words and, consequently, the extent to which an executive decision could be subjected to judicial scrutiny.[39]

[39] Liversidge v Anderson [1941] UKHL 1; [1942] AC 206. House of Lords, judgment of 3 November 1941. Dissenting opinion of Lord Atkin, especially pp. 225–247, 244–245. Available at: https://uniset.ca/other/cs5/1942AC206.html.

The historical circumstances were exceptional. Britain was engaged in a war that threatened its existence, and the government possessed emergency powers whose justification could not be divorced from the dangers confronting the country. The majority of the House of Lords accepted an interpretation under which the relevant judgment about reasonable cause belonged substantially to the Home Secretary, provided he acted in good faith. Their reasoning reflected concern that judicial examination of sensitive security information and executive assessments would interfere with the exercise of emergency powers. Whatever criticism may subsequently be directed at that approach, it would be historically inaccurate to portray the majority simply as judges indifferent to national survival or incapable of understanding ordinary English.

Lord Atkin nevertheless dissented, and his opinion has acquired a significance extending far beyond the particular wartime dispute. He insisted that the statutory expression concerning reasonable cause could not simply be reduced to the minister's personal conviction that his reasons were sufficient. The existence of reasonable grounds, he argued, was a condition capable of objective examination. His distinction was elementary in grammatical form but profound in legal consequence. To possess reasonable grounds for a belief is not necessarily the same thing as believing that one possesses them. If the two propositions are treated as interchangeable, a condition intended to limit governmental authority risks becoming little more than the government's declaration that it has complied with the condition.

Atkin recognized that the issue was not merely linguistic. The interpretation adopted by the majority had consequences for the liberty of individuals detained without ordinary criminal proceedings. By making the executive's own assessment effectively decisive, the law risked losing an important external safeguard against arbitrary detention. His dissent therefore addressed the relationship between the meaning of legal language, the allocation of institutional authority, and the possibility of independent examination. It was in this context that he introduced his celebrated literary comparison, observing that the proposed interpretation reminded him of the peculiar authority claimed by Humpty Dumpty.[40]

The reference was devastating because it exposed the logical structure of the dispute without requiring a lengthy philosophical excursion. The executive, in Atkin's analysis, could not be permitted to transform a legally prescribed condition into an exclusively subjective judgment merely by asserting its preferred interpretation. His criticism did not establish that every legal expression possesses only one possible meaning, irrespective of context. Nor did it eliminate the ordinary difficulties of judicial interpretation. Its force lay in the insistence that interpretation must remain accountable to language, legal structure, and the purposes of the conditions imposed by law. Otherwise, an apparent limitation upon authority may be converted into an extension of that authority through the very process intended to define it.

There is an important analogy here with the production and assessment of academic knowledge. A researcher who announces that a proposition has been demonstrated must be able to identify the grounds upon which that claim rests. The fact that an argument appears

[40] Liversidge v Anderson, pp.244-245.

convincing to its author does not establish that its premises are true, its sources authentic, or its conclusions justified.

Similarly, an editor who declares that a manuscript has undergone adequate scholarly verification cannot establish that fact merely by repeating the declaration. The relevant question concerns what was actually examined, by whom, according to which standards, and with what demonstrable result. The distinction resembles Atkin's opposition between possessing reasonable grounds and merely believing that such grounds exist, although the legal standards and consequences are, of course, different.

This analogy does not turn academic publishing into a system of constitutional adjudication. Courts exercise coercive authority under law, whereas journals operate through professional standards, editorial agreements, and scholarly conventions. The difference is fundamental. Nevertheless, both activities depend upon distinguishing a justified conclusion from the unsupported assertion of a person or institution entitled to pronounce it. Once that distinction weakens, the formal appearance of authority can begin to conceal the absence of the intellectual work from which legitimate authority should arise.

The philosophical foundations of this problem are considerably older than either Carroll's literary invention or Atkin's dissent. In his 1784 essay *Beantwortung der Frage: Was ist Aufklärung?* (*An Answer to the Question: What Is Enlightenment?*), Immanuel Kant described enlightenment through the individual's emergence from dependence upon the guidance of others in the exercise of understanding. His famous invocation of *Sapere aude* expressed the courage required to employ one's own reason rather than surrender its exercise to an external authority.[41]

Kant was addressing the intellectual and political conditions of the Enlightenment, not the operation of digital technologies, and his argument should not be retrospectively transformed into a theory of artificial intelligence. Nevertheless, the distinction he established between possessing the capacity for understanding and being willing to exercise it remains relevant. Intellectual dependence does not always arise because a person lacks the ability to think. It may also persist because established habits, convenient authorities, or institutional arrangements discourage the effort required for independent judgment. In the contemporary university, this problem acquires an additional dimension when external assistance is not provided by a tutor, a religious authority, or an established scholarly tradition, but by an automated system capable of presenting an apparently complete argument almost instantaneously.

One can readily imagine Kant's surprise at discovering that an instrument capable of supplying answers to an extraordinary range of questions might become an occasion for further intellectual dependence rather than intellectual emancipation. Yet such an outcome is neither inevitable nor inherent in the technology. A researcher may use generative AI precisely to challenge accepted interpretations, expose weaknesses in conventional arguments, or expand access to unfamiliar fields of knowledge. The danger lies in confusing

[41] Kant, I. (1784). Beantwortung der Frage: Was ist Aufklärung? Berlinische Monatsschrift, 4 (December), 481–494. Original German text available at: https://de.wikisource.org/wiki/Beantwortung_der_Frage:_Was_ist_Aufklärung%3F.

the availability of explanations with the possession of understanding. An instrument may greatly enlarge the materials available for independent thought while simultaneously making it tempting to avoid the intellectual effort through which those materials become knowledge.

This is also where the distinction between intellectual authority and intellectual appearance becomes particularly important. Academic writing has developed a recognizable professional language. It values precision, conceptual consistency, engagement with relevant literature, and the explicit presentation of arguments. These conventions are indispensable when they help communicate complex research. They become problematic when their visible characteristics are mistaken for evidence that the underlying intellectual work has actually been performed.

A manuscript may contain an impressive introduction, elaborate theoretical distinctions, numerous references, and conclusions expressed with perfect confidence. None of these features establishes that the author understands the relevant literature, has examined the primary sources, or can defend the proposed argument against informed criticism. Conversely, a genuinely original investigation may initially be written in awkward prose, contain imperfect transitions, or require substantial editorial assistance before its significance becomes apparent. A serious scholarly evaluation must therefore distinguish literary presentation from intellectual substance without dismissing the importance of either.

Generative language models make this distinction more demanding because they can reproduce the recognizable linguistic conventions of academic discourse with remarkable fluency. They can assemble plausible theoretical frameworks, imitate the tone of scholarly disagreement, produce explanations that appear carefully qualified, and construct elaborate verbal distinctions. These capabilities are valuable when used as instruments under informed human direction. Yet linguistic sophistication does not provide an independent guarantee that the statements produced are well grounded. A coherent argument can be mistaken; a convincing citation can be inaccurate; an elegant classification can be conceptually empty. The difficulty arises when the reader, author, or reviewer begins to treat formal sophistication as a substitute for the verification of substance.

A particularly illuminating philosophical distinction was developed by Harry G. Frankfurt in his essay *On Bullshit*, originally published in 1986. Frankfurt examined a form of discourse whose defining characteristic is not necessarily deliberate falsehood, but a fundamental indifference to the truth of what is being communicated. A liar ordinarily operates with some conception of the truth, which the lie is designed to conceal or distort. The person engaged in what Frankfurt analyzes as bullshit need not maintain such a relationship with the facts. The immediate objective may instead be to create a particular impression, sustain a performance, or communicate an appearance of knowledge without the discipline required to establish its accuracy.[42]

Frankfurt's analysis requires careful application to artificial intelligence. His account concerns human communicative conduct and the intentions associated with it. A language

[42] Frankfurt, H. G. (1986). On Bullshit. Raritan, 6(2), 81–100. Reissued as: Frankfurt, H. G. (2005). On Bullshit. Princeton: Princeton University Press. Available at: https://philpapers.org/rec/FRAOBS.

model should not be described as a liar or a producer of bullshit in precisely the same philosophical sense, since the relevant intentional and epistemic conditions cannot simply be attributed to the machine. The analogy concerns the relationship between the resulting discourse and the standards by which claims are assessed. Machine-generated statements may be accurate, inaccurate, or insufficiently grounded without the system independently assuming responsibility for distinguishing among those possibilities. Whether the text subsequently becomes part of a responsible scholarly argument depends upon the human and institutional processes through which it is examined and adopted.

The more disturbing possibility, therefore, is not that artificial intelligence has invented a new form of intellectual dishonesty. It is that academic institutions may increasingly encounter polished discourse whose relationship to truth has not been adequately established, while the formal procedures of publication continue to confer upon it an appearance of scholarly legitimacy. This situation does not require deliberate deception by every participant. A manuscript may be submitted in good faith, reviewed superficially under severe time pressure, and accepted by an editor who assumes that the reviewer has performed a more comprehensive examination. Each person may believe that the necessary intellectual work has been completed somewhere else in the process. The final publication nevertheless emerges with the institutional marks of scholarly approval.

Here the distinction identified by Lord Atkin returns in a different setting. To assert that a proposition has been verified is not equivalent to having verified it. To describe an article as peer reviewed is not sufficient to demonstrate the competence, independence, or substantive adequacy of the review. Nor does the inclusion of an impressive bibliography establish that the cited works support the propositions for which they have been invoked. These differences may appear self-evident, yet modern academic administration often relies upon simplified indicators that can obscure precisely the distinctions upon which scientific credibility depends.

The sociology of science offers an important counterweight to the tendency to treat such indicators as ends in themselves. In *The Normative Structure of Science*, Robert K. Merton identified organized skepticism as a defining element of the scientific ethos. Scientific claims, on this account, must remain open to critical examination according to empirical and logical criteria, independently of the personal standing of those who advance them. Merton's analysis was normative rather than a claim that every scientific institution consistently behaved according to these principles. Its significance lies in describing the standards against which institutional scientific practice may properly be assessed.[43]

Organized skepticism is not a ceremonial expression of doubt. It requires the possibility of discovering that an influential scholar is mistaken, that a fashionable theory lacks evidence, or that an accepted explanation must be revised. Such criticism may be uncomfortable, particularly in environments where professional advancement depends upon reputation, institutional affiliation, and publication records. Yet the willingness to submit

[43] Merton, R. K. (1973). The Normative Structure of Science. In: The Sociology of Science: Theoretical and Empirical Investigations. Chicago: University of Chicago Press, pp. 267–278. Originally published as "Science and Technology in a Democratic Order," Journal of Legal and Political Sociology, 1 (1942), 115–126. Available at: https://sciencepolicy.colorado.edu/students/envs_5110/merton_sociology_science.pdf.

intellectual authority to independent examination remains one of the conditions that distinguish scientific inquiry from the production and circulation of authoritative opinion.

This is precisely why the expansion of AI-assisted academic writing deserves serious attention. The problem cannot be reduced to the frequency of factual errors or the presence of fabricated references. Even a manuscript containing entirely genuine sources may be intellectually deficient if those sources have been misunderstood, selectively interpreted, or assembled into an argument that its author cannot independently defend. More sophisticated generative systems may reduce some forms of technical error while leaving this deeper problem unresolved. Indeed, improvements in fluency and factual reliability could make superficial assessments less capable of identifying the absence of genuine scholarly contribution.

Nor should we assume that extensive publication necessarily indicates either exceptional intellectual achievement or academic misconduct. Productive researchers may legitimately publish frequently, especially when working in collaborative disciplines or developing closely related research projects. Equally, a relatively small publication record does not guarantee originality or depth. The number of papers is an administrative fact; their scientific value requires independent evaluation. Confusing the two would reproduce the very error under examination, substituting a convenient numerical indicator for the more demanding task of scholarly judgment.

The Humpty Dumpty Academy is therefore not an institution in which every professor invents definitions or every manuscript is written by a machine. It is a metaphor for a more subtle institutional disorder: the gradual weakening of the distinction between intellectual authority and the evidence required to justify it. In such an environment, familiar concepts may acquire new names without acquiring greater explanatory power; elaborate arguments may be accepted without sufficient examination of their premises; and the formal procedures designed to certify scholarly quality may become detached from the intellectual activity they were intended to protect.

The image is amusing because its central character is so magnificently certain of his own authority. He does not need to investigate the meaning of a word or persuade Alice that his interpretation is defensible. He simply announces that the word means whatever he wishes. Yet the real danger appears when that attitude ceases to be a literary absurdity and becomes an institutional habit. A university that allows academic standing to substitute for argument, a journal that treats formal compliance as proof of scientific merit, or an author who mistakes the confident output of a generative model for an independently established conclusion has accepted, in different ways, the proposition that authority may determine what ought to be demonstrated.

There remains an essential difference, however, between the world of Carroll's satire and the institutions of contemporary scholarship. Alice was free to question Humpty Dumpty, and his pretensions did not survive the ordinary logic of their conversation. Academic institutions possess more consequential mechanisms of recognition. Publication, professional accreditation, editorial endorsement, and university affiliation can influence which arguments receive attention and which researchers acquire further opportunities. These mechanisms are not inherently defective; scientific communities require procedures

for evaluating work and allocating limited resources. Their legitimacy nevertheless depends upon the quality of the judgments through which recognition is conferred.

The central question is therefore not whether academic authority should exist, but whether its exercise remains connected to independent intellectual examination. This question becomes particularly pressing when automated systems can produce far more convincing scholarly material than existing procedures may be able to evaluate. The challenge is not simply to identify machine-generated language, which may be neither necessary nor sufficient for establishing scientific quality. It is to determine whether the person responsible for a scholarly claim has examined its grounds and whether the institution certifying that claim has exercised a form of judgment worthy of its authority.

At this point, the philosophical and legal discussion returns to the practical organization of academic research. The university cannot preserve its intellectual mission merely by insisting that a human being remains formally responsible for every published manuscript. Responsibility requires the capacity to understand what is being approved, the opportunity to question it, and the willingness to reject it when the evidence proves insufficient. Without these conditions, the forms of scholarly judgment may survive while their substance becomes increasingly uncertain. The consequences of that separation are most visible in the contemporary systems of manuscript production, peer review, and academic recognition, where the distinction between genuine verification and its procedural imitation is no longer a purely theoretical concern.

4.3. Scholarship Without Scholars? The Crisis of Academic Verification

The difficulties examined in the preceding sections become particularly consequential when they enter the institutions responsible for evaluating and transmitting scientific knowledge. An individual researcher may misunderstand a source, misinterpret a statistical result, or accept an argument that proves less convincing upon closer examination. Such mistakes have always accompanied intellectual inquiry. Science does not require its practitioners to be infallible; it requires their claims to remain open to criticism, correction, and independent verification. The more serious problem begins when the institutions established to perform these functions become increasingly unable to distinguish genuine research from its convincing imitation.

Academic publishing occupies a central position in this process. A scientific journal does more than distribute manuscripts. Through editorial selection, peer review, publication, and subsequent correction, it contributes to the construction of a scholarly record upon which other researchers may rely. Publication does not transform a hypothesis into established truth, nor does peer review guarantee that every argument or result has been independently reproduced. Nevertheless, these procedures create a reasonable expectation that the work has undergone some form of competent intellectual examination. When that expectation becomes unreliable, the consequences extend beyond the reputation of an individual author or journal. They affect the conditions under which scientific knowledge accumulates.

The vulnerability of this system did not begin with generative artificial intelligence. Long before ChatGPT became widely available, academic publishing was already confronting organized forms of manuscript production, manipulated peer review, fabricated

research data, and commercial arrangements designed to secure publication rather than advance knowledge. The 2022 joint report of the Committee on Publication Ethics (COPE) and the International Association of Scientific, Technical and Medical Publishers (STM) documented the activities of so-called *paper mills*: organizations offering manufactured manuscripts, authorship positions, and related publication services. The report described the pressures that encourage participation in such arrangements and the difficulties publishers face when attempting to identify fraudulent submissions.[44]

The chronology matters. The existence of a commercial market for questionable academic publications before the widespread adoption of generative LLMs demonstrates that the underlying incentives and institutional vulnerabilities cannot reasonably be attributed to those technologies alone. An academic system that rewards publication counts without adequately evaluating intellectual contribution may create opportunities for the production of low-value or fraudulent research regardless of the available writing tools. Generative AI can make certain parts of that production faster and cheaper, but it did not create the demand for scholarly credentials, the competition for academic recognition, or the temptation to substitute formal achievement for substantive investigation.

The persistence of the problem is illustrated by the international United2Act initiative, launched in 2023 to coordinate responses to paper mills. Its first-phase report, published in September 2025, described the development of guidance, investigative practices, training materials, and collaborative measures intended to strengthen research integrity across the scholarly publishing system. The initiative reflects an important institutional recognition: fraudulent or systematically unreliable publications cannot be addressed adequately by isolated editors acting independently of publishers, universities, research organizations, and other participants in scholarly communication.[45]

There is nevertheless a danger in treating these developments solely as a struggle between honest publishers and dishonest manuscript factories. The market for manufactured academic credentials does not exist in an institutional vacuum. It responds to expectations and rewards established within the academic system itself. When careers depend heavily upon publication records, when research performance is evaluated through numerical indicators, and when institutions seek visible evidence of productivity, the distinction between producing knowledge and producing evidence of having produced knowledge may become increasingly difficult to maintain. Such incentives do not make misconduct inevitable, nor do they render quantitative measures useless. They do, however, create conditions in which the appearance of scholarly accomplishment can acquire an exchange value of its own.

Generative artificial intelligence enters this environment with a capability that differs significantly from earlier instruments of academic production. A manuscript factory once

[44] COPE & STM. (2022). Paper Mills: Research Report from COPE & STM. Committee on Publication Ethics and International Association of Scientific, Technical and Medical Publishers. Version 1, June 2022, 15 pp. https://doi.org/10.24318/jtbG8IHL

[45] United2Act. (2025). United2Act Phase 1: Report and Recommendations. International Association of Scientific, Technical and Medical Publishers (STM). Published September 3, 2025. https://stm-assoc.org/united2act-phase-1-report-and-recommendations/

required a combination of human writers, templates, fabricated or recycled materials, and administrative coordination.

Many of these activities can now be assisted by systems capable of producing technically sophisticated prose, assembling references, summarizing literature, and reorganizing arguments within remarkably short periods. The resulting text may be entirely legitimate when supervised by a competent scholar, but the same technological capabilities can also be used to increase the volume and apparent sophistication of work whose substantive foundations remain weak.

Evidence of the transformation of scholarly writing is already available. In a study published in *Science Advances* in 2025, Dmitry Kobak and colleagues examined changes in vocabulary across more than fifteen million biomedical abstracts indexed in PubMed between 2010 and 2024. Their analysis identified an abrupt increase in words disproportionately associated with LLM-assisted writing, leading the researchers to estimate that at least 13.5 percent of the abstracts published in 2024 had been processed with LLM assistance. The estimated proportion varied considerably across fields, countries, and publication venues.[46]

These findings document a significant change in the production of scientific language, but they must not be confused with evidence that an equivalent proportion of scientific studies were fabricated, intellectually deficient, or written without meaningful human participation. The method identifies statistical features associated with machine-assisted textual production; it does not reconstruct the research process behind each publication. A scientist may employ an LLM to improve grammar or clarity while retaining complete responsibility for the underlying investigation. Another may use the same instrument to generate explanations or conclusions that have never been independently examined. The presence of machine assistance does not, by itself, establish which of these situations occurred.

This distinction is indispensable to any responsible discussion of academic integrity. The use of artificial intelligence in scientific writing is not inherently incompatible with intellectual independence. Indeed, prohibiting such assistance indiscriminately would risk confusing the form of scholarly expression with its substantive origins. Researchers who work in a language other than their native one may benefit substantially from competent linguistic assistance. Others may use computational systems to organize extensive material, identify possible counterarguments, or improve the accessibility of complex findings. These applications can support scholarship without transferring the responsibility for scientific judgment to the instrument.

The real question concerns which intellectual functions remain attributable to the human researcher. Did the author formulate and understand the central problem? Were the cited sources examined independently? Were the methods selected and evaluated by someone competent to assess their suitability? Can the author explain the reasoning

[46] Kobak, D., González-Márquez, R., Horvát, E.-Á., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. Science Advances, 11(27), eadt3813. https://doi.org/10.1126/sciadv.adt3813

connecting evidence to conclusions, respond to informed objections, and identify the limitations of the study?

These questions are not answered by determining whether a particular paragraph was written with an LLM. Nor are they resolved by a general declaration that artificial intelligence was used during manuscript preparation.

The difficulty is recognized, at least in principle, by contemporary scholarly publishing policies. The Committee on Publication Ethics states that AI tools cannot qualify as authors because they cannot assume responsibility for scholarly work, manage conflicts of interest, or undertake the legal obligations associated with authorship. Human authors remain responsible for the submitted manuscript, including its machine-assisted components, and must disclose relevant uses of AI.[47] In August 2026, *Nature Computational Science* similarly emphasized transparency, human accountability, and the preservation of independent scholarly judgment throughout manuscript preparation and peer review.[48]

These policies represent an important normative foundation, but they do not by themselves establish that meaningful verification has occurred. A disclosure statement tells the reader that an instrument was used and, ideally, identifies its role in the research process. It cannot demonstrate that every source was examined, that the analytical framework was independently developed, or that the author possesses sufficient expertise to defend the conclusions. An author may disclose the use of AI with exemplary transparency while nevertheless submitting work of limited intellectual value. Conversely, a manuscript prepared without any generative assistance may contain equally serious methodological weaknesses. Disclosure contributes to accountability, but it is not a substitute for verification.

A closely related problem of non-delegable responsibility has been examined in the author's earlier legal research. *Procedural Liability in the Age of LLMs* addressed the duty of legal professionals to independently verify AI-generated authorities, factual assertions, and procedural submissions.[49] *The Tool Has a Manufacturer: Producer Liability for Legal AI* extended this analysis to the responsibilities of developers and commercial providers whose systems are incorporated into professional decision-making.[50] Although academic publishing and judicial proceedings operate under different legal and institutional regimes, both confront the same underlying problem: the delegation of technical or linguistic operations does not, by itself, transfer responsibility for the substantive validity of the

[47] COPE Council. (2024). Authorship and AI Tools. Committee on Publication Ethics. Position statement, version 1, August 20, 2024. https://doi.org/10.24318/cCVRZBms

[48] Responsible and Transparent Use of AI in Scientific Publishing. (2026). Nature Computational Science, 6(8), 803. Published August 20, 2026. https://doi.org/10.1038/s43588-026-01043-4

[49] Kildeev, A. (2026a). Procedural Liability in the Age of LLMs. Scientific Platform XXI Century, No. 7 (16) (July 2026), pp. 27–51. Published journal version. https://episs.ru/

[50] Kildeev, A. The Tool Has a Manufacturer: Producer Liability for Legal AI. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(2), 118–153. Published journal version. https://www.crail.ca/CRAIL_Vol.1-No.2-2026-October.pdf

resulting claims. Applied to scholarly publishing, this distinction requires attention to the respective duties of authors, reviewers, editors, and technological providers.

Legitimate scholarly collaboration must be distinguished from the abdication of judgment. No single author or reviewer is required to reproduce every specialist contribution personally; what matters is whether material claims remain subject to identifiable, competent and accountable evaluation at the appropriate level. Delegating performance does not, by itself, transfer the duty to assess and justify what an institution ultimately certifies. Related questions of verification and allocated responsibility have been examined in the context of AI-assisted legal practice.

The distinction becomes even more significant when the scholarly manuscript passes from its author to the reviewer. Peer review is frequently described as the principal safeguard of scientific publishing, but its effectiveness depends upon conditions that cannot be guaranteed merely by establishing a formal procedure. The reviewer must possess relevant expertise, have sufficient time to examine the work, retain independence from improper influences, and understand the limits of the evaluation being performed. Even under favourable circumstances, a reviewer cannot necessarily reproduce an entire experimental investigation, inspect every underlying dataset, or independently verify every historical and legal proposition. A responsible review therefore requires both substantive examination and a clear understanding of what remains unverified.[51]

The growing capacity for automated manuscript production may place additional pressure upon this already imperfect arrangement. Where the number and complexity of submitted papers increase faster than the availability of qualified reviewers, editorial offices face a difficult choice. They may extend review periods, reject more manuscripts before external evaluation, distribute work among additional specialists, or gradually reduce the depth of examination expected from each reviewer. None of these consequences follows automatically from the use of LLMs; journals differ substantially in resources, submission practices, and editorial standards. Nevertheless, the underlying asymmetry identified earlier in this chapter remains relevant. Producing an apparently complete argument can be considerably faster than independently establishing whether the argument is correct.

A survey of legal academics provides a useful illustration of the professional distinctions involved. In *AI Use in Legal Scholarship*, Wayne Unger reported preliminary findings from 268 responses concerning the acceptable uses of generative AI across different stages of scholarly writing. According to the author's publicly available research memorandum, 78 percent of respondents considered AI assistance with grammar acceptable, while only 11 percent considered AI drafting analytical sections containing novel arguments acceptable. The survey also reported strong support for developing norms governing scholarly AI use, alongside relatively limited reported adoption of formal institutional policies.[52]

[51] Kildeev A., Cited works.

[52] Unger, W. (2026). AI Use in Legal Scholarship. SSRN Working Paper. Preliminary research memorandum, abstract and appendices, 15 pp. Posted August 17, 2026; last revised August 25, 2026. https://doi.org/10.2139/ssrn.7290238

These findings are instructive, but their evidentiary status requires qualification. The publicly available SSRN document is a fifteen-page work-in-progress memorandum containing an abstract, questionnaire, and preliminary results rather than the completed research article. The available materials do not yet establish all the methodological details necessary to assess representativeness, including a sufficiently documented sampling frame, response rate, and subgroup analysis. The reported attitudes therefore describe the respondents rather than proving a consensus across the entire legal academy. Even with these limitations, the contrast between acceptance of linguistic assistance and concern about machine-generated original arguments is directly relevant to the distinction developed in the present study.

The problem is not that every scholarly argument must originate in complete intellectual isolation. Research has always involved consultation, collaboration, disagreement, and the development of ideas through engagement with other minds. Scientific authorship is rarely an exercise in solitary invention. What matters is the relationship between the contribution attributed to an author and the intellectual work that the author actually performed. A researcher may legitimately adopt an argument proposed by a colleague, provided it is examined, understood, appropriately acknowledged, and incorporated into a defensible scholarly position. The same standards of understanding and responsibility should apply when an argument originates from a computational instrument.

Yet the formal attribution of responsibility becomes increasingly fragile when the person expected to exercise judgment cannot reconstruct the reasoning upon which a conclusion depends. This is the problem identified in Alejandro Sinopoli's discussion of *Cognitive Dignity*: a person may remain formally present within a process while the cognitively significant functions attributed to that person have been materially displaced. The distinction between assistance, voluntary delegation, and substantive substitution becomes particularly important when institutions treat formal approval as evidence of independent understanding.[53]

An academic reviewer who receives an AI-generated manuscript may understand the relevant discipline but lack the time necessary to verify its numerous claims. Another reviewer may possess impressive professional credentials but insufficient competence in the particular subject under examination. A third may accept the text because its terminology, references, and conclusions resemble familiar forms of academic argument. These situations are not equivalent. The first concerns the practical limits of qualified verification under increasing workloads; the second concerns the inappropriate allocation of expertise; the third involves the danger of mistaking formal familiarity for substantive reliability. In each case, however, the existence of a nominal human approval stage may conceal the absence of meaningful evaluation.

At this point, the familiar expression *human-in-the-loop* deserves a moment of reconsideration. The phrase suggests that the presence of a human participant provides an essential safeguard against the limitations of automated systems. But the effectiveness of that safeguard depends upon the person's competence, independence, attention, and genuine opportunity to reject the machine's conclusions. Without these conditions, one might

[53] Sinopoli. Cognitive Dignity, pp.102-108.

imagine Homer Simpson sitting contentedly before a control panel, pressing the appropriate button while remaining cheerfully unaware of what the machinery is actually doing. The procedure has retained its human operator, but not necessarily the judgment for which that operator was supposed to be present. What was intended as ***human-in-the-loop*** has quietly become ***Homer-in-the-loop***.

The humour is familiar; the institutional problem is not. A reviewer who lacks the competence to assess a manuscript does not become qualified merely by occupying an editorial position. Equally, a highly competent reviewer cannot perform unlimited verification simply because an administrative system assigns additional manuscripts. The distinction matters because the corresponding remedies differ. Inadequate expertise requires better reviewer selection and clearer disciplinary boundaries. Excessive workloads require realistic allocation of time, resources, and expectations. Neither problem can be solved by attaching another automated approval stage to a process whose underlying deficiencies remain unchanged.

Nor would it be sensible to dismiss machine assistance in editorial work altogether. Automated reference checks, plagiarism screening, manuscript-format validation, and identification of potentially problematic images can provide useful preliminary safeguards. Computational tools may help reviewers locate relevant sources, compare versions of a document, or identify inconsistencies requiring further examination. Their contribution is valuable precisely because it allows human attention to be directed toward questions that require interpretation and judgment. What these instruments cannot establish merely through their operation is that the research has been independently understood and found intellectually defensible.

The distinction between technical screening and scholarly evaluation becomes clearer when one considers the difficulty of assessing original theoretical work. A citation can often be checked against an identifiable publication, but verifying its existence does not determine whether it supports the proposition for which it is cited. A mathematical formula can be syntactically correct without being relevant to the phenomenon under investigation. A philosophical concept may be expressed with great precision while contributing little beyond the restatement of established arguments. In such cases, evaluating the work requires knowledge of the field, intellectual independence, and a willingness to challenge the author's assumptions rather than merely inspect the formal properties of the document.

The evaluation of scholarly originality presents a further difficulty that no automated screening procedure can fully resolve. A theoretical argument may be elegantly constructed, extensively referenced, and expressed in the language of an established discipline while contributing little to its development. Conversely, an unfamiliar or unconventional proposition may challenge accepted assumptions and deserve serious consideration precisely because it does not conform to prevailing academic expectations. The responsibility of a competent reviewer is to distinguish these possibilities through careful examination of the argument, its sources, and its actual contribution to knowledge. Such evaluation requires intellectual independence as well as disciplinary expertise. Without them, peer review risks becoming an exercise in recognizing familiar academic forms rather than assessing the substance of scholarly work.

A related danger concerns the institutional treatment of preprints and working papers. Repositories perform a valuable function by allowing researchers to circulate preliminary results, establish public records of their work, and invite criticism before formal publication. Yet repository availability should not be confused with independent scholarly endorsement. A paper may be publicly accessible, assigned a persistent identifier, and indexed by search engines without having undergone substantive external peer review. Such dissemination is entirely legitimate, provided its status is clear. The problem arises when the visibility and formal presentation of a manuscript encourage readers or institutions to attribute to it a level of scientific validation that has not actually occurred.

The same distinction applies to the growing use of AI-generated research summaries, automated literature reviews, and academic discovery systems. These tools may help readers navigate the expanding scientific literature, but they can also reproduce the appearance of scholarly consensus without adequately distinguishing original research, preliminary hypotheses, disputed interpretations, and unverified claims. The resulting risk is cumulative. A speculative proposition may be summarized as a finding, cited by another manuscript as established background, and eventually appear in later literature stripped of the qualifications that originally accompanied it. Such processes existed long before generative AI, but automated production and redistribution can alter their speed and scale.

The possibility of recursively incorporating machine-generated material into subsequent computational systems introduces a further concern, although it must be distinguished from the direct deterioration of human scholarly competence. Shumailov and colleagues demonstrated that repeated training on recursively generated data can produce forms of model collapse under the conditions examined in their study. This result concerns the statistical properties of machine-learning systems and should not be presented as proof that scientific knowledge itself must inevitably collapse when scholars use generative AI.[54] Its relevance is more limited: the quality and provenance of information entering subsequent processes matter, whether those processes are computational or institutional. The analogy does not eliminate important differences between human scientific inquiry and machine learning.

At the level of scholarly publishing, the corresponding institutional question concerns the preservation of reliable distinctions within the scientific record. A journal must be able to distinguish original findings from commentary, verified evidence from inference, preliminary research from established results, and editorial acceptance from independent reproduction. These distinctions do not guarantee that scientific knowledge will progress without error, but they provide the conditions under which errors can be identified and corrected. When publication systems lose the ability or willingness to maintain them, the accumulation of documents may increasingly diverge from the accumulation of dependable knowledge.

It would nevertheless be a mistake to describe contemporary academia as a uniformly collapsing institution. Strong universities, demanding editorial boards, independent research

[54] Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631, 755–759. https://doi.org/10.1038/s41586-024-07566-y

groups, and exceptional scholars continue to make substantial contributions across numerous disciplines.

Developments in medicine, physics, mathematics, engineering, and other fields provide ample evidence that scientific inquiry remains vigorous and productive. The available evidence does not establish an irreversible or universal decline in academic competence. It does, however, justify concern about institutional practices that may weaken the conditions necessary for maintaining that competence over time.

The question is therefore not whether scientific knowledge has ceased to advance, but whether the institutions responsible for its reproduction can sustain the intellectual standards upon which future advances depend. A university may preserve excellent research laboratories while simultaneously encouraging superficial forms of academic assessment. A journal may publish outstanding articles while struggling to identify questionable submissions elsewhere in its portfolio. A researcher may possess genuine expertise in one field while accepting unreliable automated summaries in another. These possibilities are not mutually exclusive. Institutional strength and institutional vulnerability can coexist within the same academic system.

This coexistence makes the task of reform more demanding than either technological enthusiasm or generalized cultural pessimism would suggest. Simply prohibiting generative AI would not eliminate publication incentives, commercial manuscript factories, inadequate reviewer selection, or the misuse of academic credentials. Equally, unrestricted adoption accompanied only by statements about human responsibility would not ensure that the relevant judgment is actually exercised. Responsible integration requires an institutional environment in which authors remain answerable for their arguments, reviewers receive tasks appropriate to their competence, editors exercise independent authority, and scholarly claims remain open to examination after publication.

Such an environment also requires intellectual honesty about the limits of verification. No editor can guarantee the truth of every published proposition, and no reviewer should be expected to possess universal disciplinary competence. The purpose of scholarly evaluation is not to create an illusion of infallibility, but to establish credible, transparent, and corrigible procedures through which intellectual claims may be examined. A responsible journal should be able to explain the kind of review a manuscript received, identify the standards applied, manage conflicts of interest, and correct the scholarly record when significant errors emerge. These duties are not innovations demanded by artificial intelligence; they are established responsibilities whose importance becomes more visible as automated production expands.

There is a final irony in the contemporary debate. Generative AI has frequently been presented as a technology capable of democratizing access to intellectual production. In significant respects, this promise is real. Researchers without substantial institutional resources can obtain linguistic assistance, explore unfamiliar literature, and prepare work that might previously have been difficult to communicate effectively. Yet wider access to the means of producing scholarly text does not automatically provide the education, expertise, or critical judgment necessary to produce scholarship. Democratizing the tools of intellectual expression and cultivating the capacity for independent thought are related but

distinct objectives. The former can be achieved through technological distribution; the latter requires sustained educational and institutional effort.

The historical achievement of the university was never merely its ability to preserve books, organize examinations, or confer academic titles. Its deeper contribution lay in cultivating communities capable of understanding, criticizing, extending, and transmitting knowledge. Scientific publications were among the visible results of that activity, but they were not its ultimate purpose. When the production of those results becomes increasingly independent of the intellectual capacities traditionally required to produce them, the university faces a question that cannot be resolved through publication metrics or improvements in automated writing alone.

The challenge extends beyond the accuracy of today's scholarly record. Universities reproduce themselves through the training of future researchers, the transmission of disciplinary standards, and the continuity of intellectual traditions. If successive generations become increasingly dependent upon systems whose conclusions they cannot independently examine, the consequences may eventually affect not only the quality of individual publications but also the capacity of academic institutions to preserve and renew their own intellectual foundations. This possibility requires empirical examination before its wider institutional consequences, including the weakening of intellectual continuity, can be assessed.

Chapter 5. Empirical Evidence and the Limits of Cognitive Delegation

5.1. Cognitive Offloading and the Question of Independent Judgment

The preceding analysis of the abdication of academic judgment raises a question that legal philosophy cannot resolve through conceptual argument alone. If scholarly institutions increasingly accept intellectual products generated with the assistance of large language models, is there any empirical reason to question whether the people responsible for those products will retain the knowledge and evaluative competence necessary to assess them independently? This chapter addresses that limited question. It does not diagnose cognitive decline among students or university teachers, assume that generative artificial intelligence causes a general deterioration of intelligence, or attempt an independent investigation in psychology or educational science. Its purpose is to determine whether relevant findings in those disciplines identify a sufficiently credible risk to warrant institutional attention.

An appropriate starting point is the established concept of cognitive offloading. Risko and Gilbert describe how people use external actions and resources to reduce the demands placed on internal cognitive processes.[55] Such practices long predate generative AI: writing, calculation, indexing, and electronic reminders all redistribute cognitive work. Offloading is neither inherently deficient nor inherently harmful. It may release attention for more demanding intellectual activity. Its consequences depend upon the task, the individual's existing competence, and the activity being delegated. This distinction matters because the mere use of an external instrument cannot itself establish the erosion of judgment.

[55] Risko, E. F. Cognitive Offloading / E. F. Risko, S. J. Gilbert // Trends in Cognitive Sciences. 2016. Vol. 20, no. 9. pp. 676–688. DOI: 10.1016/j.tics.2016.07.002.

Generative systems nevertheless introduce a particular institutional problem. They can supply not only facts or computational assistance but apparently complete explanations, arguments, evaluations, and written conclusions. Where the purpose of an educational or scholarly exercise is to develop the very operation being outsourced, the immediate quality of the resulting product may cease to be a reliable measure of the person's acquired competence. A student may submit an accurate answer without being able to reproduce its reasoning; a researcher may receive a plausible synthesis without independently testing its sources. These are possibilities, not findings about all users. They identify the question against which the empirical studies must be read: what happens when assistance improves the performance of a task but may reduce the intellectual work through which independent mastery is formed?

5.2. Experimental Evidence: Performance, Retention, and the Conditions of Assistance

The study designs, principal measurements, and inferential limits are summarized in Exhibit D. One of the clearest experimental illustrations comes from Bastani and colleagues, who studied generative-AI tutoring in mathematics with nearly one thousand secondary-school students.[56] They distinguished between a relatively unrestricted GPT-4-based interface and a tutor designed to support the student's reasoning rather than simply provide solutions. Access to both systems improved performance during assisted practice. The critical distinction emerged when access was withdrawn: students assigned to the less restricted interface performed approximately 17 per cent worse than those who had never received AI assistance on the subsequent unassisted assessment. The experiment therefore separates two outcomes often conflated in public discussion—successful completion of assisted work and learning that remains available without the tool. Its scope must, however, be stated precisely. This was a school mathematics experiment, not an investigation of university faculty or a demonstration of long-term cognitive decline. The result is relevant because it establishes that the design of assistance can change the relationship between immediate performance and subsequent independent performance.

Barcaui examined a related question in higher education through a randomized study of undergraduate learning and later knowledge retention.[57] Initially, 120 students were assigned to ChatGPT-assisted or traditional study conditions. At a surprise assessment 45 days later, the mean score among the 85 participants who completed follow-up was 57.5 per cent in the assisted condition and 68.5 per cent in the comparison condition. The observed difference was statistically significant. Unlike an immediate test of productivity, this outcome concerns retention after the initial learning activity. Yet the study does not justify a claim about durable cognitive impairment. Its particular instructional setting limits generalization, and the loss of 35 participants—29.2 per cent of the original sample—creates

[56] Bastani, H. Generative AI without guardrails can harm learning: Evidence from high school mathematics / H. Bastani, O. Bastani, A. Sungu [et al.] // Proceedings of the National Academy of Sciences. 2025. Vol. 122, no. 26. Article e2422633122. DOI: 10.1073/pnas.2422633122.

[57] Barcaui, A. ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention / A. Barcaui // Social Sciences & Humanities Open. 2025. Vol. 12. Article 102287. DOI: 10.1016/j.ssaho.2025.102287.

a substantial attrition concern, notwithstanding the reported balance of follow-up completion across conditions. The result is best treated as evidence of a possible learning cost under specified circumstances, not as a general law of AI-supported study.

Any assessment confined to adverse findings would itself fail the standard of independent judgment that this article defends. Makransky and colleagues reported two controlled experiments involving 175 university students and 234 secondary-school students.[58] They evaluated a theory-informed conversational tutor that encouraged learners to explain material and actively construct understanding, comparing it with a generic AI interface and relevant control conditions. The structured tutor yielded stronger conceptual-knowledge outcomes than generic ChatGPT in both immediate assessments and in a delayed assessment in one of the studies; comparisons with the non-AI controls were more variable. These findings neither prove that generative AI always improves learning nor erase the adverse results reported elsewhere. They demonstrate something more analytically useful: the educational consequences of the same broad technological category cannot be assessed without examining whether its design supports or substitutes the learner's own cognitive activity.

The issue is not limited to retention. Yan's randomized, scenario-based experiment enrolled 240 preservice teachers at a Chinese university; 218 were retained in the principal cleaned analysis.[59] Participants faced pedagogical cases containing both accurate and misleading generative-AI recommendations. Under misleading advice, those who recorded an independent judgment before seeing the recommendation achieved higher expert-keyed accuracy than those who saw the recommendation before recording their first judgment (53.5 per cent versus 37.3 per cent on the relevant items). The controlled result concerns susceptibility to erroneous recommendations rather than any demonstrated decline of intelligence. Moreover, the participants were teachers in training, the tasks were simulated, and transfer to genuine classroom decisions remains unproven. The experiment nevertheless identifies a practically significant variable: the sequence in which a person first reasons and then consults an algorithmic recommendation may affect the integrity of the final judgment.

Taken together, these experiments support a cautious but important proposition. Generative-AI assistance is not a unitary educational intervention. It may facilitate understanding when designed and used to sustain active explanation, while forms of assistance that replace reasoning may produce gains in visible performance without comparable gains in independent competence. The empirical claim is conditional. It does not follow that every use of an LLM weakens cognition, still less that an entire generation of students or scholars is undergoing cognitive deterioration.

[58] Makransky, G. Beyond the "Wow" Factor: Using Generative AI for Increasing Generative Sense-Making / G. Makransky, B. M. Shiwalia, T. Herlau, S. Blurton // Educational Psychology Review. 2025. Vol. 37. Article 60. DOI: 10.1007/s10648-025-10039-x.

[59] Yan, D. Independent judgment before viewing generative AI recommendations reduces susceptibility to misleading advice: A randomized scenario-based experiment with preservice teachers / D. Yan // Frontiers in Psychology. 2026. Vol. 17. Article 1936522. DOI: 10.3389/fpsyg.2026.1936522.

5.3. From Limited Findings to Academic Responsibility

Evidence concerning university teachers must be handled with particular restraint. Qiang and Wang analyzed multi-source accounts from 25 university teachers recruited through an AI-focused professional learning community and associated networks in China.[60] Their qualitative analysis addresses teachers' reported experiences of cognitive delegation, professional judgment, and strategies for checking generated material. Such accounts are valuable for identifying how practitioners describe the shifting boundaries of their work, but they are not measurements of diminished intellectual ability. The sample is small and selectively recruited; the materials are uneven across cases; and self-reports cannot establish the effects of AI use on actual classroom performance. It would therefore be unjustified to convert this study into evidence that university teaching as a profession has cognitively deteriorated. Its narrower significance lies in documenting that the preservation of independent assessment has become an explicit concern within teachers' accounts of their own practice.

A broader methodological perspective is supplied by the systematic review and meta-analysis of Boolzen and colleagues, published in 2026.[61] The authors identified 85 eligible studies of generative AI and cognitive learning outcomes in STEM education and included 49 studies, representing 59 effect sizes, in the quantitative synthesis. Although a conventional random-effects model indicated an overall positive effect, the findings displayed extreme heterogeneity. An analysis designed to account for possible publication bias substantially reduced the estimated average benefit. The authors also reported that none of the included randomized or non-randomized intervention studies was assessed as having a low overall risk of bias. These observations preclude simple declarations either that AI reliably damages learning or that it reliably improves it. Of particular relevance is the review's distinction between AI used to augment learners' cognitive activity and AI used to substitute for it: this distinction helps account for variation in outcomes, while leaving considerable uncertainty unresolved.

Several limits of inference follow. First, a short-term difference in test performance or knowledge retention is not equivalent to a persistent change in general cognitive capacity. Secondly, experimental findings involving school pupils, undergraduates, and preservice teachers cannot be transferred without qualification to experienced scholars and university professors. Thirdly, heterogeneous interventions—ordinary chat interfaces, specially designed tutors, different tasks, and different degrees of human verification—cannot be treated as one uniform exposure. Finally, the existing literature does not establish that changes attributed to AI are irreversible, widespread across academic systems, or caused by the technology independently of prior weaknesses in teaching, assessment, research

[60] Qiang, Z. Between Human and Machine: University Teachers' Accounts of Judgment, Metacognitive Defense, and Cognitive Offloading in Generative AI Use / Z. Qiang, Z. Wang // Journal of Intelligence. 2026. Vol. 14, no. 10. Article 254. DOI: 10.3390/jintelligence14100254.

[61] Boolzen, C. Evidence of impact and interpretational limits of generative AI in STEM education: A systematic review and meta-analysis on cognitive learning outcomes / C. Boolzen, J. Kuhn, S. Flegr [et al.] // Artificial Intelligence Review. 2026. Vol. 59. Article 223. DOI: 10.1007/s10462-026-11665-9.

incentives, or institutional governance. The cited findings support a reason for caution; they do not provide a scientific warrant for a diagnosis of academia as cognitively degraded.

That conclusion does not deprive the evidence of institutional significance. Academic institutions routinely assign responsibility to identifiable persons: the student must demonstrate what has been learned, the author must substantiate claims, the reviewer must examine the submission, and the editor must make an accountable decision. These duties presuppose more than a human name attached to a finished document. They presuppose an effective capacity to inspect arguments, recognize uncertainty, evaluate sources, and refuse an unsupported conclusion. The studies considered here do not show that this capacity has generally disappeared. They do, however, reveal circumstances in which successful use of a cognitive instrument may be difficult to distinguish from the independent competence that academic procedures are intended to cultivate and verify.

The resulting issue is juridical and institutional rather than clinical. A journal may accept AI-assisted drafting while still requiring authors to verify every material assertion and citation. A university may permit AI-assisted learning while preserving forms of assessment that test knowledge without the tool. Reviewers and editors may use automated assistance without treating the resulting text as an independent act of scholarly judgment. None of these safeguards demands hostility to technology. Their common purpose is to ensure that the allocation of responsibility remains connected to a demonstrable human capacity to exercise it.

In the terminology developed in this article, the empirical literature does not establish *abdicatio iudicii* as an accomplished, universal condition. It identifies conditions under which the risk of such abdication becomes intelligible. The doctrinal concern begins when institutional procedures continue to attribute judgment to a human actor while allowing the actual formation and verification of that judgment to be displaced by generated outputs. Whether, how often, and with what consequences this occurs must remain open to further empirical investigation. The legal and philosophical obligation is already clear: apparent fluency or assisted productivity must not be mistaken for independent verification, and procedural responsibility must not be reduced to the ceremonial presence of a human signatory. This distinction provides the bridge to the following examination of institutional memory and the economics of intellectual discontinuity.

Chapter 6. Institutional Amnesia and the Economics of Intellectual Discontinuity

Universities are commonly understood as institutions devoted to the preservation and transmission of knowledge. Their libraries, laboratories, archives, and increasingly sophisticated digital repositories give this mission an impressive material form. A modern university can preserve millions of publications, provide access to scientific databases, and maintain an institutional record extending across generations. Yet the continuity of scientific knowledge depends upon something more demanding than the survival of documents. A scientific tradition persists only when successive generations acquire the intellectual capacities necessary to understand, question, reproduce, and extend the work of their predecessors.

This distinction becomes particularly important when considering the possibility of institutional amnesia. The expression does not necessarily imply that universities have forgotten their histories or abandoned their libraries. Scientific knowledge can remain extensively documented while the practical competence required to interpret and develop it becomes increasingly difficult to reproduce. An archive may preserve the results of a research programme without preserving every methodological judgment, experimental adjustment, unsuccessful hypothesis, or intellectual disagreement through which those results were obtained. The published record remains available, but its effective continuation depends upon the education and experience of those who inherit it.

The transmission of scientific knowledge has therefore always involved relationships between experienced researchers and those entering their disciplines. Laboratory practices, professional judgment, methodological caution, and the ability to recognize significant problems are developed through sustained participation in research communities. Formal instruction provides indispensable foundations, but it cannot fully replace the experience of working alongside scholars who possess the accumulated knowledge of a field. The importance of these relationships is supported by the National Academies of Sciences, Engineering, and Medicine in *The Science of Effective Mentorship in STEMM* (2019), which identifies mentorship as a significant element of scientific training and emphasizes the responsibility of institutions to support its quality and effectiveness.[62]

The same institutional concern appeared in the National Academies' earlier report *Graduate STEM Education for the 21st Century* (2018). That study acknowledged the substantial achievements of American graduate education while calling for improvements in professional preparation, mentoring, transparency, and the development of competencies suitable for different research careers.[63] These recommendations are particularly instructive because they were formulated before the widespread adoption of generative language models. The challenge of reproducing scientific expertise was already recognized as an institutional problem, independent of the technologies now transforming academic writing.

The conditions under which this transmission occurs have nevertheless been changing. Data published by the American Association of University Professors show that approximately 47 percent of faculty members in American higher education held contingent appointments outside the tenure system in 1987, compared with approximately 68 percent in 2023. Over the same period, the proportion holding full-time tenured or tenure-track appointments declined from approximately 53 percent to 32 percent. These figures describe substantial changes in the organization of academic employment, although the proportions differ considerably among institutional categories and should not be interpreted as a uniform transformation of every research university.[64]

---

[62] National Academies of Sciences, Engineering, and Medicine. (2019). The Science of Effective Mentorship in STEMM. Washington, DC: The National Academies Press. https://doi.org/10.17226/25568

[63] National Academies of Sciences, Engineering, and Medicine. (2018). Graduate STEM Education for the 21st Century. Washington, DC: The National Academies Press. https://doi.org/10.17226/25038

[64] American Association of University Professors. (2025). Data Snapshot: Tenure and Contingency in US Higher Education, Fall 2023. Academe. https://www.aaup.org/academe/issues/spring-2025/data-snapshot-tenure-and-contingency-us-higher-education-fall-2023, see also Exhibit B.

The significance of these developments requires careful interpretation. A temporary or part-time appointment does not establish intellectual inferiority, insufficient professional commitment, or poor teaching. Many accomplished scholars work outside traditional tenure arrangements, while permanent academic appointments provide no automatic guarantee of scientific originality. Nor do aggregate employment figures directly measure the quality of research mentorship. Nevertheless, employment arrangements influence the organizational conditions within which scholars teach, supervise students, conduct long-term investigations, and participate in departmental life. Where continuity of employment is uncertain or professional responsibilities are fragmented among several institutions, maintaining sustained research relationships may become more difficult. Whether these risks materialize depends upon institutional practices and cannot be inferred solely from appointment categories.

Changes in the destinations of doctoral graduates provide another perspective on the reproduction of the academic workforce. According to the U.S. National Center for Science and Engineering Statistics, the proportion of doctorate recipients with definite non-postdoctoral employment commitments in the United States who expected to work in academia declined from 56 percent in 2004 to 40 percent in 2024. In the latter year, industry and business accounted for a comparable proportion of those employment commitments.[65] These percentages concern a specifically defined group of graduates reporting definite employment commitments; they are not percentages of all doctorate recipients and should not be treated as measures of permanent career outcomes.

The movement of highly educated researchers into industry is not necessarily evidence of institutional weakness. Research-intensive companies, technological enterprises, medical organizations, and other employers offer important opportunities for scientific work. The ability of universities to prepare researchers for such careers may itself be regarded as a substantial educational achievement. The institutional question is different: how does the university maintain the professional capacity required to train future researchers when the distribution of scientific employment changes? A successful research system may benefit from extensive movement between universities and industry, but the continuity of academic research communities still requires sufficient opportunities for experienced scholars to supervise new generations.

England introduces a third institutional configuration, one which should not be treated as interchangeable with either the United States or Canada. Its universities operate within a distinctive arrangement of tuition fees, student finance, public funding bodies and research grants. England-specific statistics also have to be distinguished from figures reported for the United Kingdom as a whole. This matters particularly for employment and participation data, where superficially similar categories may refer to different populations or appointment regimes. The relevant comparison is consequently institutional rather than a forced equivalence between American tenure, Canadian academic contracts and British employment classifications.

---

[65] National Center for Science and Engineering Statistics. (2026). Doctorate Recipients from U.S. Universities: 2024. NSF 26-315. https://ncses.nsf.gov/pubs/nsf26315/report/postgraduation-trends, see also Exhibit B.

These developments acquire particular significance when considered alongside the financial expansion of research activity. The American National Center for Science and Engineering Statistics reported that higher-education research and development expenditures reached $117.7 billion in fiscal year 2024, representing an 8.1 percent nominal increase over the preceding year. Between 2014 and 2024, university research expenditure grew at an average annual rate of approximately 3 percent after adjustment for inflation.[66] This is substantial evidence of the continuing scale and financial strength of American academic research. It would be misleading to portray the entire university research sector as an institution undergoing straightforward contraction.

Canadian statistics likewise demonstrate considerable research activity. Statistics Canada reported higher-education research and development expenditure of C$19.0 billion in 2023/24, an annual increase of 7.2 percent.[67] Although the American and Canadian measures differ in currency, institutional coverage, and reporting methodology, they independently establish that significant financial resources continue to support university-related research in both countries. These expenditures cannot, however, be converted into direct measurements of scholarly originality, methodological competence, or intellectual independence.

The English financial picture requires a different indicator. According to the House of Commons Library's September 2026 briefing, higher-education providers in England received £46.1 billion in total income in 2024/25. This represents institutional income from multiple sources, not research-and-development expenditure, and therefore must not be placed in a three-country ranking beside the American and Canadian R&D figures. The same parliamentary analysis traces a pronounced alteration in England's funding structure: tuition fees accounted for 24 percent of provider income in 2005/06 but exceeded one half from 2019/20, while funding-body grants declined from 39 percent in 2005/06 to 11 percent in 2024/25. These observations describe a transformation in the financing of English higher education, without by themselves demonstrating changes in the quality of scholarship.

The distinction between expenditure and institutional capacity deserves closer examination. Research budgets finance equipment, laboratories, technical personnel, scientific projects, and numerous activities essential to discovery. An increase in expenditure may support valuable new investigations or compensate for rising costs without producing a proportionate change in the number of researchers, the quality of training, or the continuity of particular research programmes. Conversely, an institution may maintain highly productive research groups despite financial constraints. Aggregate expenditure therefore provides necessary economic context, but it cannot independently establish whether the conditions for reproducing scientific expertise are improving.

[66] National Center for Science and Engineering Statistics. (2026). Universities Report 8.1% Growth in R&D Expenditures in FY 2024, Reaching Over $117 Billion. NSF 26-305. https://ncses.nsf.gov/pubs/nsf26305, see also Exhibit B.

[67] Statistics Canada. (2025). Spending on Research and Development in the Higher Education Sector, 2023/2024. The Daily, 17 October. https://www150.statcan.gc.ca/n1/daily-quotidien/251017/dq251017e-eng.htm, see also Exhibit B.

The Canadian university financing system illustrates the complexity of this relationship. According to Statistics Canada, total university revenues reached C$52.4 billion in 2023/24, expressed in constant 2024 dollars. Provincial funding represented 32.4 percent of revenues, while tuition accounted for 30.1 percent. Over the preceding decade, the tuition share had increased by approximately 5.4 percentage points.[68] These figures reveal changes in financial composition, not a direct deterioration of educational quality. Universities differ in their dependence upon particular revenue sources, institutional mandates, and allocation of funds. Moreover, apparent financial surpluses may include resources restricted to purposes other than ordinary academic operations.

English institutions exhibit a further tension between substantial aggregate revenues and uneven financial resilience. The House of Commons Library records that 48 publicly funded universities in England, representing 40 percent of that defined group, reported expenditure exceeding income in 2024/25.[69] In a separate regulatory population, the Office for Students reported deficits at 35.8 percent of higher-education providers in 2024/25, while estimating that 42.7 percent could report deficits in 2025/26 without mitigating action. The OfS also recorded £218.2 million in restructuring costs in 2024/25.[70] The percentages must not be blended: the parliamentary university subgroup and the OfS provider population have different denominators, and the 2025/26 figure is a conditional forecast rather than an observed result. Moreover, the regulator noted an improvement in aggregate financial performance for 2024/25. The simultaneous presence of improvement and significant institutional pressure is central to a balanced interpretation.

England is also instructive because the financial structure distinguishes direct institutional support for research from revenue derived from student fees. The parliamentary briefing reports that direct funding through Research England, which supports research capacity and infrastructure rather than only specific projects, broadly maintained its real value from 2010 to 2024 before falling in real terms during the following two years.[71] This does not establish an inevitable decline in scientific capacity. It does, however, bring the economics of continuity into focus: research infrastructure and established scholarly communities may be maintained by funding streams whose purposes and time horizons differ from those of student-facing operations.

Financial composition matters because universities must make decisions about competing institutional priorities. Resources devoted to facilities, technology, administrative systems, immediate educational demand, and research infrastructure may all be necessary. The difficulty is that their economic benefits become visible over different

[68] Statistics Canada. (2025). University Revenues Increase More Than Expenditures in 2023/2024. The Daily, 29 July. https://www150.statcan.gc.ca/n1/daily-quotidien/250729/dq250729b-eng.htm, see also Exhibit B.

[69] Bolton, P., and Lewis, J. (2026). Higher Education Finances and Funding in England. House of Commons Library, Research Briefing CBP-10037, 4 September. https://commonslibrary.parliament.uk/research-briefings/cbp-10037/

[70] Office for Students. (2026). Financial Sustainability of Higher Education Providers in England: 2026. 14 May 2026, updated report. https://www.officeforstudents.org.uk/publications/financial-sustainability-of-higher-education-providers-in-england-2026/, see also Exhibit B.

[71] Bolton, P., and Lewis, J., Cited article.

periods. A completed building, a new laboratory instrument, or a digital platform can be identified in an annual budget. The development of a researcher capable of independently formulating significant scientific problems may require many years, and its eventual benefits may emerge outside the institution that financed the original training.

This difference in time horizons creates a potentially important economic problem. The immediate costs of scientific education are concentrated and measurable, while some of its most valuable returns are uncertain, delayed, and distributed across institutions, industries, and generations. A university that supports demanding research training may benefit society through discoveries made decades later by graduates working elsewhere. Such benefits are difficult to capture through annual institutional performance indicators. The resulting incentive problem does not establish that universities systematically neglect fundamental research or mentorship, but it helps explain why these activities require deliberate institutional protection.

Canadian federal advisory research has addressed related concerns. In its 2023 *Report of the Advisory Panel on the Federal Research Support System*, an independent expert panel emphasized the importance of investigator-initiated research, sustainable support for graduate students and postdoctoral researchers, and the long-term development of scientific talent. The panel warned that inadequate support for research trainees and insufficient investment in fundamental research could weaken Canada's capacity to attract and retain researchers.[72] These observations represent a documented institutional assessment at the time of the report, rather than proof that all identified problems remained unchanged in 2026.

The importance of investigator-initiated research extends beyond the allocation of grants. Scientific progress cannot always be organized around problems whose practical significance has already been established. Many discoveries emerge from investigations that initially appear remote from immediate economic or administrative priorities. Yet a balanced research system must also accommodate public priorities, technological development, and practical applications. The relevant economic question concerns the institutional conditions under which short-term objectives and long-term scientific inquiry can coexist without one progressively displacing the other.

It is equally necessary to distinguish concerns about academic continuity from broader claims of educational decline. The available evidence is heterogeneous. In the OECD's 2023 Survey of Adult Skills, Canadian adults achieved average scores of 271 in literacy and numeracy, both above the OECD averages. Compared with the preceding survey cycle, literacy performance remained broadly stable while numeracy improved.[73] These findings do not measure original scientific judgment or university research competence, but they caution against narratives portraying an uninterrupted and universal deterioration in intellectual capacities.

---

[72] Advisory Panel on the Federal Research Support System. (2023). Report of the Advisory Panel on the Federal Research Support System. Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/panel-federal-research-support/en/report-advisory-panel-federal-research-support-system

[73] OECD. (2024). Survey of Adult Skills 2023: Canada, Country Note. https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/canada_5ecab9d9-en.html, see also Exhibit A.

The English results reinforce the need for such caution. In the same OECD adult-skills programme, adults aged 16–65 in England obtained mean scores of 272 in literacy, 268 in numeracy and 259 in adaptive problem solving. The OECD reported broadly stable literacy and improved numeracy relative to the 2011/12 survey cycle.[74] School-based PISA observations address a different age group and cannot be combined numerically with adult-skills results; England-specific PISA 2022 reporting also contains explicit sampling qualifications.[75] None of these assessments measures independent academic scholarship directly. Their value is methodological: improvements and stability must remain in view alongside financial and institutional vulnerabilities.

What emerges from the three cases is a set of institutional tensions rather than a single statistical trajectory. Universities in the United States, Canada and England continue to support substantial scientific investigation, educate researchers, and preserve extensive intellectual resources. At the same time, differing employment structures, doctoral career opportunities, public funding arrangements and financial pressures raise questions about the sustainability of the conditions under which scholarly expertise is reproduced. The available indicators do not establish a uniform degree of institutional decline, nor do they justify assuming a single causal mechanism across three distinct university systems.

Generative artificial intelligence adds a new dimension to these existing challenges. Digital systems can make scientific literature more accessible, assist the preservation of technical documentation, facilitate collaboration, and reduce the time required for routine research operations. Properly integrated into education, they may strengthen the ability of researchers to engage with difficult problems. Yet the preservation of information should not be confused with the preservation of competence. A university may acquire increasingly powerful systems for retrieving and generating scientific explanations without ensuring that successive generations retain the ability to examine their assumptions, reproduce their methods, or identify their limitations.

The economic attraction of such systems is understandable. Where institutions face pressure to increase measurable output while controlling expenditure, automation offers opportunities to accelerate routine processes and reduce particular categories of work. These benefits deserve serious evaluation. Their long-term value, however, depends upon whether the resulting efficiencies strengthen intellectual training or merely increase production. If automation reduces the time required for scholarly administration and allows researchers to devote greater attention to difficult questions, its contribution may be substantial. If it instead becomes an excuse to weaken the institutional conditions through which scientific expertise is developed, apparent efficiency may conceal costs that become visible only later.

There is no reliable basis for calculating a universal economic cost of such intellectual discontinuity, still less for predicting the inevitable disappearance of scientific schools. The empirical material considered here does not establish that generative AI has already caused

[74] OECD. (2024). Survey of Adult Skills 2023: England (United Kingdom), Country Note. https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/united-kingdom_02bc78e4-en.html, see also Exhibit A.

[75] Department for Education. (2023). PISA 2022: National Report for England. Research Report RR1392. https://www.gov.uk/government/publications/pisa-2022-national-report-for-england

a measurable decline in research competence in any of the three cases. Much of the institutional transformation discussed in this chapter began decades before contemporary LLMs entered academic practice. The technology must therefore be understood within existing economic and organizational environments rather than treated as their original cause.

The deeper question concerns the relationship between the preservation of intellectual products and the preservation of the human capacities required to create them. A scientific institution may accumulate publications without adequately supporting the formation of new investigators; it may acquire advanced instruments while neglecting the skills necessary to evaluate their results; and it may celebrate research productivity without sufficiently examining the conditions under which future discoveries will become possible. These outcomes are neither universal nor inevitable. They are institutional risks whose seriousness depends upon decisions made by universities, public authorities, funding bodies, and scholarly communities.

Institutional amnesia, in this sense, begins when an organization loses sight of the processes that allow knowledge to remain a living intellectual practice. Its consequences need not be immediate. Scientific institutions may continue to function successfully for considerable periods while depending upon competencies developed through earlier arrangements. The task of responsible academic governance is to ensure that the conditions required to renew those competencies remain available to future generations.

The evidence from the United States, Canada and England does not justify announcing the collapse of the university. It does justify examining how different research systems allocate resources, sustain scientific careers, support mentorship, and distinguish intellectual achievement from its measurable appearances. The relevant uncertainty concerns the future, but institutional choices must be made in the present. The concluding chapter therefore turns from the diagnosis of possible intellectual discontinuity to the principles through which universities may preserve the independence of human judgment while continuing to benefit from technological progress.

Chapter 7. Preserving the Human Academy

The central finding of this study is neither the imminent collapse of the university nor the inevitable degradation of human intelligence through artificial intelligence. The evidence examined across the United States, Canada, and England supports neither conclusion. Universities continue to generate significant scientific discoveries, sustain extensive research programmes, and educate new generations of scholars. Their financial structures, employment practices, and educational outcomes reveal substantial differences and, in some instances, contradictory tendencies. What emerges from this investigation is a more specific institutional problem: the growing possibility that the visible production of academic knowledge may become detached from the human capacities required to understand, verify, and reproduce it.

This distinction provides the conceptual basis for the present inquiry. Academic knowledge has always involved an interaction between intellectual activity and its material expression. A scientific publication records an investigation; a degree recognizes an

educational achievement; a research institution preserves and develops disciplinary expertise.

These relationships have never been perfect. An incompetent scholar could publish before the invention of the computer, just as a distinguished university could confer recognition upon work whose significance later proved negligible. Generative artificial intelligence did not create these weaknesses. It has, however, introduced an instrument capable of producing many of the external characteristics of scholarly accomplishment without requiring a corresponding intellectual achievement from the person who employs it.

The resulting danger extends beyond the familiar problem of inaccurate machine-generated information. A fabricated citation can be identified, a statistical error corrected, and a defective argument rejected. More difficult to recognize is the situation in which the document appears satisfactory while the intellectual competence normally presumed behind it is absent or insufficient. An article may contain authentic references, plausible conclusions, and grammatically impeccable explanations without establishing that its nominal author understands the argument. An editor may accept the article after completing the required procedural stages without establishing that its central claims have received competent scrutiny. A university may subsequently recognize the publication as evidence of scholarly achievement. At every stage, the formal requirements may appear satisfied, while the substantive basis of that recognition remains uncertain.

This is the mechanism of academic abdication examined in the preceding chapters. It occurs when responsibility for intellectual judgment remains formally assigned to a human being, but the relevant person either cannot exercise that judgment or is not provided with the conditions necessary to exercise it meaningfully. Its causes may include inadequate expertise, excessive workloads, inappropriate institutional incentives, or habitual reliance upon automated systems. The presence of artificial intelligence is neither a necessary nor a sufficient condition. What makes contemporary generative technologies significant is their capacity to widen the distance between the speed of intellectual production and the resources required for independent evaluation.

The first consequence concerns the meaning of scholarly authorship. Authorship cannot reasonably require a researcher to perform every calculation personally, compose every sentence without assistance, or possess exhaustive knowledge of every discipline relevant to a collaborative investigation. Modern science depends upon instruments, specialization, and intellectual cooperation. Nevertheless, the attribution of a scholarly contribution presupposes a defensible relationship between the author and the claims presented under that author’s name. A researcher must be capable of explaining the central reasoning, identifying the evidentiary foundation, understanding the relevant methodological limitations, and responding responsibly to substantive criticism. Where specialist contributions are involved, their provenance and responsibility must remain identifiable. A polished manuscript cannot substitute for these conditions.

Disclosure of artificial intelligence assistance is an important element of transparency, but it cannot establish intellectual competence. The statement that a manuscript was prepared with the assistance of an LLM tells the reader something about its production.

It does not demonstrate that its sources were examined or its arguments understood. Equally, the absence of such a disclosure cannot establish scholarly merit. Scientific publishing must therefore avoid confusing the identification of a technological instrument with the evaluation of the intellectual work performed through it. The central question is not simply who or what produced the words, but whether the person responsible for the scientific claims can justify them.

The second consequence concerns peer review. A review procedure acquires scholarly legitimacy through the quality of the judgment exercised within it. Editorial appointments, institutional affiliations, submission platforms, and standardized evaluation forms cannot create expertise where it is absent. Nor can they guarantee that a qualified reviewer has sufficient time to examine a difficult manuscript. Effective verification requires the appropriate allocation of disciplinary competence, reasonable opportunities for independent examination, access to relevant evidence, and procedures through which uncertainty, disagreement, and error can be recorded and addressed. These conditions impose costs, but eliminating them does not necessarily eliminate the costs of scientific evaluation. It may merely transfer those costs to future researchers, institutions, and readers who must confront inadequately examined claims after publication.

This does not mean that every reviewer must reproduce an entire research project or personally verify every proposition within a multidisciplinary paper. Such an expectation would be unrealistic and could make scientific publishing impossible. The appropriate standard is proportionate, competent, and transparent evaluation. Reviewers should examine the claims they are qualified to assess, identify material questions requiring additional expertise, and avoid representing unperformed verification as completed work. Editors must exercise judgment in selecting reviewers, interpreting their reports, and deciding whether unresolved deficiencies justify rejection, revision, or further investigation. Accountability requires an identifiable human decision and a reconstructible basis for that decision. A signature, an approval button, or a completed checklist is not itself evidence that the corresponding intellectual examination occurred.

The third consequence concerns the educational mission of universities. Academic qualifications have meaning because they are understood to recognize capabilities developed through study and research. If students can satisfy increasingly demanding formal requirements without acquiring the corresponding competence, the relationship between academic certification and intellectual achievement becomes weaker. This possibility cannot be resolved solely through prohibitions on generative technologies or attempts to identify machine-produced prose. Universities must retain credible methods of establishing whether students understand their subjects, can defend their reasoning, recognize competing interpretations, and respond when evidence contradicts their initial conclusions. Substantive supervision, methodological examinations, discussion of research decisions, and the critical assessment of intellectual development may be more informative than reliance upon the finished document alone.

Such measures must be designed with respect for academic freedom and the legitimate use of technological assistance. The objective is not to compel researchers to demonstrate every moment of independent thought or expose their private intellectual processes to

institutional surveillance. It is to ensure that the achievement recognized by an academic qualification corresponds to an actual human capability. The use of computational instruments is compatible with that objective when the researcher retains the competence necessary to direct, interpret, and evaluate their contribution.

The fourth consequence is institutional and economic. Scientific knowledge is reproduced through people, relationships, and organizations whose development frequently exceeds the horizon of annual budgets. Research schools require continuity; mentorship requires time; fundamental investigation may yield its most important results only after prolonged uncertainty. The statistical material examined in this study documents changes in academic employment, doctoral career pathways, research expenditure, university revenues, and financial conditions across three national systems. These observations do not establish a common trajectory of intellectual decline. They do, however, demonstrate why the sustainability of academic expertise cannot be evaluated through expenditure totals or publication volumes alone.

A university may legitimately seek to reduce administrative costs, improve operational efficiency, and employ automation to support research. The relevant economic question is what happens to the resources thereby released. If technological efficiencies allow experienced scholars to devote more time to teaching, mentorship, experimental work, and difficult intellectual problems, automation may strengthen the very capacities whose preservation concerns us. If the same efficiencies are treated exclusively as opportunities to increase measurable output while reducing support for the development of expertise, the institution risks weakening its long-term intellectual foundations. The possibility of such a trade-off is a reason for systematic evaluation, not a finding that every university has already made the wrong choice.

Here the limitations of the present research must remain explicit. Neither university budgets nor standardized educational assessments directly measure the independent judgment of scholars. Employment data do not establish the quality of mentorship, and changes in research expenditure do not reveal whether individual scientific traditions are becoming more or less sustainable. The evidence assembled for the United States, Canada, and England identifies institutional conditions, historical transformations, and questions requiring investigation. It does not establish that generative artificial intelligence has caused a measurable decline in scientific competence or that such a decline is inevitable. Those propositions would require longitudinal research capable of distinguishing technological effects from pre-existing institutional developments.

Such research is both possible and necessary. Universities could examine whether students retain the ability to explain and reconstruct arguments developed with computational assistance, whether different approaches to AI-supported education affect independent problem-solving over time, and whether the growth of automated manuscript production changes the accuracy, duration, and substantive depth of peer review. Longitudinal studies of research training could investigate the relationship between employment continuity, mentorship, and the development of independent scientific expertise. These investigations would require carefully defined populations, appropriate comparison groups, and a willingness to accept results that contradict expectations. The

concern defended here is a researchable institutional hypothesis; it should not become an ideology protected from empirical criticism.

The practical implications are nevertheless sufficiently clear to justify action without awaiting a definitive historical verdict. Universities and scholarly publishers already possess responsibilities that do not depend upon resolving the future capabilities of artificial intelligence. They must distinguish formal compliance from substantive evaluation, require accountable authorship, support competent independent review, and maintain conditions under which scientific expertise can be transmitted. They must also preserve the capacity to acknowledge institutional errors and correct the scholarly record. These obligations arise from the purpose of scientific inquiry itself, rather than from a particular technological controversy.

The deeper issue is the legitimacy of academic authority. A university possesses the power to confer qualifications, distribute recognition, and influence the development of knowledge. A scientific journal may determine which arguments enter the published record under its editorial endorsement. These forms of authority are necessary, but their exercise must remain answerable to standards external to the preferences of those who hold them. Neither an academic title nor an institutional seal can transform an unsupported proposition into a demonstrated conclusion. A persistent identifier establishes a durable reference to a publication; it does not certify the truth of its contents. Peer review may justify serious attention to a manuscript, but it cannot exempt that manuscript from subsequent criticism.

This is why the philosophical and legal examples considered earlier are more than literary ornament. Kant's insistence upon the independent exercise of understanding and Lord Atkin's refusal to allow institutional authority to determine the meaning of its own legal limitations illuminate a common intellectual principle. A claim to authority cannot replace the grounds upon which that authority is justified. Humpty Dumpty's extraordinary confidence becomes genuinely dangerous only when the power to declare is mistaken for the power to establish what is true. Scholarship survives through the opposite discipline: the willingness to submit even confident and authoritative propositions to independent examination.

The preservation of the human academy is therefore neither a programme of technological resistance nor an appeal to an idealized past. It is an institutional commitment to maintaining the relationship between knowledge, understanding, judgment, and responsibility. That commitment does not exclude independent researchers working beyond traditional universities, collaborative research organizations, or increasingly sophisticated computational instruments. It requires that scholarly authority remain connected to demonstrable intellectual competence wherever research is undertaken and wherever its results are evaluated.

The title *Academia Without Scholars* should consequently be understood as a warning about a possible institutional condition, not as a declaration that such a condition has already become universal. A university without scholars need not resemble an abandoned institution. It may possess impressive buildings, extensive budgets, sophisticated digital systems, flourishing publication statistics, and an international reputation. What it would lack is something less immediately visible and more difficult to replace: a sufficiently

independent community of people capable of understanding the knowledge it preserves, judging the claims it circulates, and creating knowledge that does not yet exist.

We cannot predict precisely what tomorrow's universities or artificial intelligence systems will become. That uncertainty provides no reason to suspend the standards by which scholarship is distinguished from its imitation. The responsibility to preserve independent inquiry belongs to the institutions and people who possess the authority to shape academic life today.

The decisive question is no longer whether machines can write like scholars. It is whether universities will continue to require that someone understands, verifies, and takes responsibility for what is written.

ACADEMIA WITHOUT SCHOLARS: ARTIFICIAL INTELLIGENCE AND THE EROSION OF INDEPENDENT JUDGMENT

EXHIBITS A–D | UNITED STATES • CANADA • ENGLAND

Empirical documentation supporting Chapter V, Exhibits A-C relate to Chapter 6, Exhibit D relates to Chapter 5. Reference date: 15 October 2026. Historical anchor: approximately 1990; contemporary observations reflect the latest compatible published data identified for each measure.

Reading rule. Enrolments, test results, academic employment, finances and R&D expenditure describe different phenomena. None directly measures the originality of scholarship, independent judgment or the causal effects of generative AI. The data include improvements and stability as well as adverse changes.

EXHIBIT A — Educational Scale and Measured Skills

A.1. Higher-education participation: observations, not forced growth rates

System / population	Historical observation	Recent observation	Comparison status
United States: enrolments at degree-granting institutions	13,818,600 (1990)	19,659,600 (2024)	NCES published series; historic institutional coverage has changed [A01]
Canada: university registrations / enrolments	841,328 (1990)	1,439,850 (2024/25)	Historical full/part-time categories differ from current domicile-based counts [A02–A03]
England: UK-domiciled higher-education students in England	800,600 (1990/91)	—	No definition-matched contemporary comparison established [A04]
England: English-domiciled higher-level enrolments in England	—	2,040,275 (2024/25)	DfE category includes apprenticeships and other higher-level learning; distinct from historical HE count [A05]

Canada: 1990 = 532,131 full-time + 309,197 part-time. In 2024/25, 1,172,976 domestic and 266,874 international university enrolments were recorded. These are different collection frameworks; no percentage growth is calculated. England: the two listed populations differ in domicile, programme coverage and institutional universe. The UK-wide HESA total of 2,863,180 in 2024/25 is not an England-only comparator [A06]. The 2024/25 figures represent programme enrolments at the autumn reference date, not necessarily unique persons [A03].

A.2. School assessments: longitudinal observations within each programme

Assessment	Baseline	Recent	Difference / qualification
United States: NAEP grade-4 mathematics	213 (1990)	237 (2024)	+24 scale points [A07]
United States: NAEP grade-4 reading	217 (1992)	215 (2024)	−2; not statistically significant [A08]
Canada: PISA reading, age 15	534 (2000)	507 (2022)	−27 points [A09]
Canada: PISA mathematics, age 15	532 (2003)	497 (2022)	−35 points [A09]
England: PISA reading, age 15	496 (2006)	496 (2022)	0 observed change; 2022 sampling caveats [A10–A11]
England: PISA mathematics, age 15	495 (2006)	492 (2022)	−3 observed points; 2022 sampling caveats [A10–A11]

NAEP, PISA and PIAAC involve distinct ages, scales and populations; their raw scores must not be pooled or read as measures of university research quality. The NAEP reading series above is the national all-schools result, not the public-schools-only estimate. England's 2022 PISA sampling qualifications are disclosed by the Department for Education.

A.3. Adult skills: evidence against an undifferentiated decline thesis

OECD Survey of Adult Skills, 2022/23	Literacy	Numeracy	Adaptive problem solving
Canada, adults 16–65	271	271	259
England, adults 16–65	272	268	259

OECD country notes [A12–A13]. Canadian and English literacy was broadly stable against the preceding survey cycle, while numeracy improved. Adaptive problem solving is presented as a contemporary observation. These are adult population skills, not measures of independently produced scholarship.

EXHIBIT B — Academic Workforce and Research Economics

B.1. United States: structure of academic appointments

Appointment category	1987	2023	Interpretation
Contingent appointments, outside tenure line	~47%	~68%	Employment structure [B01]
Full-time tenured or tenure-track	~53%	~32%	Complementary category [B01]
Part-time faculty	~33%	~49%	Subset overlapping contingent appointments [B01]

AAUP analysis of federal academic-workforce datasets. Categories may overlap: part-time is not an additional population to add to contingent. Employment type alone is not evidence of competence, teaching quality, disciplinary expertise, or mentorship performance. The AAUP population comprises appointments of nonmedical instructional staff, with or without formal faculty status, at degree-granting nonprofit institutions; graduate student employees are reported separately [B01].

B.2. Doctorate recipients: definite non-postdoctoral employment in the United States

Destination measure	2004	2024	Denominator
Academic employment sector	56%	40%	Recipients with definite U.S. non-postdoc employment commitments [B02]
Industry / business sector	—	~40%	Same 2024 group; not all doctorate recipients [B02]

In the 2024 definite-commitment group, 8,098 recipients identified academia and 8,266 industry/business among 20,469 recipients. Movements toward industry may signify opportunities outside universities, not a decline in scholarly ability. These sector shares do not, by themselves, establish a decline in the absolute number of academic positions [B02].

B.3. University research expenditure: distinct national series

Indicator	Latest observed	Reported change	Measurement boundary
United States: higher-education R&D	US$117.7bn (FY 2024)	+8.1% nominal year-on-year	US institutional HERD; +~3% annual real growth 2014–24 [B03]
Canada: higher-education R&D	C$19.0bn (2023/24)	+7.2% year-on-year	Canadian HERD sector; not an exchange-rate-adjusted comparator [B04]

U.S. and Canadian R&D totals cannot be ranked against English higher-education-provider income. Research expenditure can rise without an identical change in scholarly quality, mentorship or continuity.

B.4. University finances: Canada and England have different reporting frameworks

Population / measure	Reported observation	Scope and status
Canada: university revenue, 2023/24	C$52.4bn	Inflation-adjusted 2024 Canadian dollars [B05]
Canada: provincial grants / tuition share	32.4% / 30.1%	Share of Canadian university revenues; tuition share up 5.4 pp over ten years [B05]
England: HE-provider income, 2024/25	£46.1bn	Total England provider revenue, not R&D spending [B06]
England: tuition / funding-body grants	Tuition >50%; grants 11%	2024/25 provider income composition [B06]
England: deficits — publicly funded universities	48 institutions; 40%	Parliamentary/HESA university subgroup, 2024/25 [B06]
England: deficits — OfS-regulated providers	35.8%	Different OfS denominator, 2024/25 [B07]
England: predicted deficits, 2025/26	42.7%	OfS forward-looking provider forecasts; not observed outcomes [B07]
England: restructuring costs, 2024/25	£218.2m	OfS sector costs; 20.7% above 2023/24 [B07]

England: 48 universities/40% [B06] and OfS 35.8% [B07] concern different institutional populations and must not be combined. The OfS describes 42.7% as a prospective deficit exposure derived from submitted forecasts and its assessment of potential financial outcomes, including the role of mitigating actions. Research England's direct support for research capacity is a separate series from total provider income [B06].

EXHIBIT C — Primary Sources, Provenance and Limits of Inference

The identifiers below apply only to these Exhibits; the article-wide References remain separately numbered. Links identify official statistical or institutional sources. Dates refer to the relevant observation or publication, not to an inferred common census year.

C.1. Source register

[A01] NCES, Digest of Education Statistics, Table 306.10 (1990 and 2024 enrolments). https://nces.ed.gov/programs/digest/d25/tables/dt25_306.10.asp

[A02] Senate of Canada, Subcommittee on Post-Secondary Education, final report (1997), historical enrolment counts. https://sencanada.ca/en/content/sen/committee/361/post/rep/repfinaldec97part2-e

[A03] Statistics Canada, Canadian postsecondary enrolments and graduates 2024/25, The Daily (1 Oct 2026). https://www150.statcan.gc.ca/n1/daily-quotidien/261001/dq261001a-eng.htm

[A04] House of Commons, Hansard (20 June 1994), UK-domiciled higher education students studying in England, 1990/91. https://api.parliament.uk/historic-hansard/written-answers/1994/jun/20/student-numbers

[A05] Department for Education, Higher-Level Learners in England, 2024/25 (8 Oct 2026); English-domiciled learners. https://explore-education-statistics.service.gov.uk/find-statistics/higher-level-learners-in-england/2024-25

[A06] HESA, Higher Education Student Statistics: UK, 2024/25; UK-wide only as context. https://www.hesa.ac.uk/data-and-analysis/sb273/figure-7

[A07] NCSES / NSF, STEM Talent, NAEP mathematics, Table STAL-15, all grade-4 students, 1990–2024. https://ncses.nsf.gov/pubs/nsb20261/table/STAL-15

[A08] NCES, The Nation's Report Card: Reading 2024, grade 4, national trends. https://www.nationsreportcard.gov/reports/reading/2024/g4_8/national-trends/

[A09] OECD, PISA 2022 Results, Vol. I, Canadian country trends and historical assessments. https://doi.org/10.1787/53f23881-en

[A10] Department for Education, PISA 2018: National Report for England (historical trend tables). https://www.gov.uk/government/publications/pisa-2018-national-report-for-england

[A11] Department for Education, PISA 2022: National Report for England, RR1392, incl. sampling caveats. https://www.gov.uk/government/publications/pisa-2022-national-report-for-england

[A12] OECD, Survey of Adult Skills 2023: Canada, country note (2024). https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/canada_5ecab9d9-en.html

[A13] OECD, Survey of Adult Skills 2023: England (United Kingdom), country note (2024). https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/united-kingdom_02bc78e4-en.html

[B01] AAUP, Data Snapshot: Tenure and Contingency in U.S. Higher Education, Fall 2023 (2025). https://www.aaup.org/academe/issues/spring-2025/data-snapshot-tenure-and-contingency-us-higher-education-fall-2023

[B02] NCSES, Doctorate Recipients from U.S. Universities: 2024, postgraduation trends. https://ncses.nsf.gov/pubs/nsf26315/report/postgraduation-trends

[B03] NCSES, University R&D expenditure in FY 2024; NSF 26-305 (2026). https://ncses.nsf.gov/pubs/nsf26305

[B04] Statistics Canada, Higher Education R&D 2023/24, The Daily (17 Oct 2025). https://www150.statcan.gc.ca/n1/daily-quotidien/251017/dq251017e-eng.htm

[B05] Statistics Canada, University Revenues 2023/24, The Daily (29 July 2025). https://www150.statcan.gc.ca/n1/daily-quotidien/250729/dq250729b-eng.htm

[B06] Bolton & Lewis, Higher Education Finances and Funding in England, Commons Library CBP-10037 (4 Sep 2026). https://commonslibrary.parliament.uk/research-briefings/cbp-10037/

[B07] Office for Students, Financial Sustainability of Higher Education Providers in England: 2026 (updated June). https://www.officeforstudents.org.uk/publications/financial-sustainability-of-higher-education-providers-in-england-2026/

C.2. Inference boundaries

- Geography: only the United States, Canada and England are substantive cases. UK-wide data do not replace England-specific observations.
- Categories: Canada 1990/2024–25 student counts and England 1990/91/2024–25 higher-learning counts are non-matched observations; no percentage growth is reported for those pairs.
- Instruments: NAEP and PISA are school measures; OECD PIAAC is an adult-skills measure. None is a direct test of research judgment or of scientific originality.
- Employment: AAUP part-time and contingent categories overlap. The doctorate employment figures refer to graduates with definite U.S. non-postdoc job commitments, not all doctoral graduates.
- Finances: U.S. and Canadian HERD, Canadian university revenue and England HE-provider income have different institutional scopes, currencies, accounting bases and purposes; no ranking or conversion is justified.
- Forecasts: the OfS 2025/26 deficit statistic is forward-looking, not an observed loss rate. Its provider population is distinct from the publicly funded university group reported in Parliament.
- Causality: pre-LLM institutional trends cannot be ascribed to generative AI. Descriptive endpoints do not establish a common causal path or a general collapse of academic competence.
- Balance: increases in research resources, stable literacy and improved adult numeracy are retained alongside documented financial and institutional pressures.

Editorial scope statement. The selected indicators support inquiry into the conditions for reproducing scientific expertise; they are not measurements of scholarly merit. The series have been sourced to official records, while the excluded like-for-like comparisons are intentionally withheld rather than supplied with unsupported estimates.

EXHIBIT D — Selected Evidence on Cognitive Delegation and Independent Judgment

Methodological matrix for Chapter 5 · selected published studies · reference numbers 1-7.

Study / source	Research design and population	Principal observation	Scope and inferential limits
Risko & Gilbert (2016) [1] · 10.1016/j.tics.2016.07.002	Conceptual review of cognitive offloading; no experiment or participant sample.	External aids redistribute demands on memory and cognition. Offloading may support rather than impair thinking.	General cognitive framework, not LLM-specific evidence. Neither cognitive decline nor causation is tested.
Bastani et al. (2025) [2] · 10.1073/pnas.2422633122	Field experiment in secondary-school mathematics; approximately 1,000 pupils; unrestricted GPT-4 assistance compared with a guided tutor and controls.	Assisted exercise performance improved, yet unrestricted assistance was associated with about 17% worse unassisted assessment performance than the no-AI control.	School mathematics and particular tutoring interfaces; a learning-performance gap, not evidence of permanent impairment or university-wide effects.
Barcaui (2025) [3] · 10.1016/j.ssaho.2025.102287	Randomized undergraduate study; 120 initially allocated; 85 completed a surprise retention assessment after 45 days.	Mean retained-knowledge scores: 57.5% with ChatGPT assistance versus 68.5% in the comparison condition (reported p = 0.002).	35/120 participants absent at follow-up (29.2% attrition); one teaching setting. Does not establish generalized cognitive deterioration.
Makransky et al. (2025) [4] · 10.1007/s10648-025-10039-x	Two controlled experiments: 175 university students and 234 school students; theory-guided conversational tutoring versus generic assistance and comparators.	Structured prompting to explain and construct understanding produced stronger conceptual learning outcomes than generic ChatGPT in the tested settings.	Effects depend on instructional design and outcome measure. Material counterevidence against treating all AI use as harmful.
Yan (2026) [5] · 10.3389/fpsyg.2026.1936522	Randomized scenario experiment with 240 preservice teachers (218 in principal analysis); accurate and deliberately misleading AI advice.	For misleading items, independent judgment before AI advice was associated with higher expert-keyed accuracy (53.5% versus 37.3%).	Single training context and simulated pedagogical decisions; examines advice susceptibility, not longitudinal decline.
Qiang & Wang (2026) [6] · 10.3390/jintelligence14100254	Qualitative accounts and accompanying materials from 25 university teachers using generative AI.	Teachers described cognitive delegation, professional judgment and deliberate verification or "metacognitive defense" strategies.	Small purposive/self-selecting sample; narratives do not quantify competence loss or establish causality.
Boolzen et al. (2026) [7] · 10.1007/s10462-026-11665-9	Systematic review: 85 eligible studies; meta-analysis: 49 studies and 59 effect sizes in STEM learning.	Average positive effects in conventional synthesis were strongly qualified by very high heterogeneity and sensitivity to publication-bias adjustment.	No included study had a low overall risk-of-bias rating. Neither universal learning gains nor generalized harm is established.

Interpretive boundary. These studies address cognitive offloading, assisted performance, retention, advice reliance and scholarly practice. They do not establish generalized or irreversible cognitive deterioration among students or university faculty.

Note. Bibliographic details and DOI are provided in Chapter 5, reference numbers 1-7. Evidence types and denominators are deliberately kept separate.

References

1. Kildeev, A. (2025a). Artificial Intelligence Is Slave of the Lamp. Scientific Platform XXI Century, No. 6 (September 2025), 3–31. Published journal version. https://episs.ru/
2. Kildeev, A. (2025b). LLM Is Not Artificial Intelligence. Scientific Platform XXI Century, No. 7 (October 2025), 3–44. Published journal version. https://episs.ru/
3. Liebig, T., & Kildeev, A. (2026). Instrumentum Vocale and the Architecture of Responsibility: From Liability to Embedded Accountability. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(1), 11–84. Published journal version. https://crail.ca/
4. University of Oxford. Our History. https://www.ox.ac.uk/about/the-university/history; University of Cambridge. The Medieval University. https://www.cam.ac.uk/about-the-university/history/the-medieval-university

5. The Royal Society. History of the Royal Society. https://royalsociety.org/about-us/who-we-are/history/
6. Harvard University. The History of Harvard University. https://www.harvard.edu/about/history/
7. University of Toronto Archives & Records Management Services. The University's Original Charter. https://utarms.library.utoronto.ca/history/university-original-charter ; Université Laval. Origin and History. https://www.ulaval.ca/en/about-us/origin-and-history
8. Johns Hopkins University. History & Mission. https://www.jhu.edu/about/history/
9. Polanyi M. The Tacit Dimension. Garden City, NY: Doubleday, 1966. Reissued by the University of Chicago Press, 2009.
10. Lave J., Wenger E. Situated Learning: Legitimate Peripheral Participation. Cambridge: Cambridge University Press, 1991. DOI: 10.1017/CBO9780511815355.
11. National Academies of Sciences, Engineering, and Medicine. The Science of Effective Mentorship in STEMM. Washington, DC: National Academies Press, 2019. DOI: 10.17226/25568.
12. Boyer Commission on Educating Undergraduates in the Research University. Reinventing Undergraduate Education: A Blueprint for America's Research Universities. Stony Brook, NY: State University of New York at Stony Brook, 1998. https://repo.library.stonybrook.edu/xmlui/handle/11401/57972
13. Natural Sciences and Engineering Research Council of Canada. Discovery Grants Program (Individual). Official programme documentation. https://nserc.canada.ca/en/funding-opportunity/discovery-grants-program-individual
14. UK Research and Innovation / Research England. How We Fund Higher Education Providers. https://www.ukri.org/publications/research-england-how-we-fund-higher-education-providers/how-we-fund-higher-education-providers/
15. State Higher Education Executive Officers Association. State Higher Education Finance: FY 2025. Published 2026. https://shef.sheeo.org/report/
16. Statistics Canada. University Finances, 2002/03. The Daily, 19 August 2004; University Revenues Increase More Than Expenditures in 2023/2024. The Daily, 29 July 2025.
17. Further and Higher Education Act 1992, c. 13. https://www.legislation.gov.uk/ukpga/1992/13/contents
18. House of Commons Library. Tuition Fees in England: History, Debates, and International Comparisons. Research Briefing CBP-10155, 23 February 2026; Higher Education and Research Act 2017, c. 29.
19. Bolton, P., and Lewis, J. Higher Education Finances and Funding in England. House of Commons Library, Research Briefing CBP-10037, 4 September 2026.

20. Slaughter, S., and Rhoades, G. Academic Capitalism and the New Economy: Markets, State, and Higher Education. Johns Hopkins University Press, 2004.
21. Bok, D. Universities in the Marketplace: The Commercialization of Higher Education. Princeton University Press, 2003.
22. San Francisco Declaration on Research Assessment (DORA), 2012/2013. https://sfdora.org/read/
23. Hicks, D., Wouters, P., Waltman, L., de Rijcke, S., and Rafols, I. "Bibliometrics: The Leiden Manifesto for Research Metrics." Nature, 520 (2015): 429–431. https://doi.org/10.1038/520429a
24. UK Research and Innovation. Research Excellence Framework. Official documentation, updated 2026. https://www.ukri.org/who-we-are/research-england/research-excellence/research-excellence-framework/
25. Colby, G. "Data Snapshot: Tenure and Contingency in US Higher Education." Academe, American Association of University Professors, Spring 2023.
26. Statistics Canada. "Part-time University Faculty, 1990/91 to 1997/98." The Daily, 8 May 2002.
27. Sinopoli, A. (2026). EPISTEMÍA: Governing the human process of understanding: A proposal for a new object of legal protection in the age of artificial intelligence. The Journal Scientific Platform XXI Century, No. 5 (14), 40–72. https://episs.ru/wp-content/uploads/2026/09/5-14-ISSUE-2026-English-FINAL.pdf
28. Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (Article 1121, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778
29. Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task (arXiv:2506.08872, version 2). arXiv. Research preprint. https://doi.org/10.48550/arXiv.2506.08872
30. Bainbridge, L. (1983). Ironies of automation. Automatica, 19(6), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8
31. Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381–410. https://doi.org/10.1177/0018720810376055
32. Sinopoli, A. (2026). Cognitive dignity: A proposal for an emerging human right: Protecting the human process of sense-making and knowledge formation against technological substitution. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(2), 89–117. https://www.crail.ca/CRAIL_Vol.1-No.2-2026-October.pdf
33. Carroll, L. (1871). Through the Looking-Glass, and What Alice Found There. London: Macmillan (first published December 1871; title-page dated 1872).

Chapter VI, “Humpty Dumpty.” Available at: https://www.gutenberg.org/files/12/12-h/12-h.htm.
34. Liversidge v Anderson [1941] UKHL 1; [1942] AC 206. House of Lords, judgment of 3 November 1941. Dissenting opinion of Lord Atkin, especially pp. 225–247, 244–245. Available at: https://uniset.ca/other/cs5/1942AC206.html.
35. Kant, I. (1784). Beantwortung der Frage: Was ist Aufklärung? Berlinische Monatsschrift, 4 (December), 481–494. Original German text available at: https://de.wikisource.org/wiki/Beantwortung_der_Frage:_Was_ist_Aufklärung%3F.
36. Frankfurt, H. G. (1986). On Bullshit. Raritan, 6(2), 81–100. Reissued as: Frankfurt, H. G. (2005). On Bullshit. Princeton: Princeton University Press. Available at: https://philpapers.org/rec/FRAOBS.
37. Merton, R. K. (1973). The Normative Structure of Science. In: The Sociology of Science: Theoretical and Empirical Investigations. Chicago: University of Chicago Press, pp. 267–278. Originally published as “Science and Technology in a Democratic Order,” Journal of Legal and Political Sociology, 1 (1942), 115–126. Available at: https://sciencepolicy.colorado.edu/students/envs_5110/merton_sociology_science.pdf.
38. COPE & STM. (2022). Paper Mills: Research Report from COPE & STM. Committee on Publication Ethics and International Association of Scientific, Technical and Medical Publishers. Version 1, June 2022, 15 pp. https://doi.org/10.24318/jtbG8IHL
39. United2Act. (2025). United2Act Phase 1: Report and Recommendations. International Association of Scientific, Technical and Medical Publishers (STM). Published September 3, 2025. https://stm-assoc.org/united2act-phase-1-report-and-recommendations/
40. Kobak, D., González-Márquez, R., Horvát, E.-Á., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. Science Advances, 11(27), eadt3813. https://doi.org/10.1126/sciadv.adt3813
41. COPE Council. (2024). Authorship and AI Tools. Committee on Publication Ethics. Position statement, version 1, August 20, 2024. https://doi.org/10.24318/cCVRZBms
42. Responsible and Transparent Use of AI in Scientific Publishing. (2026). Nature Computational Science, 6(8), 803. Published August 20, 2026. https://doi.org/10.1038/s43588-026-01043-4
43. Kildeev, A. Procedural Liability in the Age of LLMs. Scientific Platform XXI Century, No. 7 (16) (July 2026), 27–51. Published journal version. https://episs.ru/
44. Kildeev, A. The Tool Has a Manufacturer: Producer Liability for Legal AI. Canadian Review of Artificial Intelligence, Law & Society (CRAIL), 1(2), 118–153. Published journal version. https://www.crail.ca/CRAIL_Vol.1-No.2-2026-October.pdf

45. Unger, W. (2026). AI Use in Legal Scholarship. SSRN Working Paper. Preliminary research memorandum, abstract and appendices, 15 pp. Posted August 17, 2026; last revised August 25, 2026. https://doi.org/10.2139/ssrn.7290238
46. Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631, 755–759. https://doi.org/10.1038/s41586-024-07566-y
47. Risko, E. F. Cognitive Offloading / E. F. Risko, S. J. Gilbert // Trends in Cognitive Sciences. 2016. Vol. 20, no. 9. pp. 676–688. DOI: 10.1016/j.tics.2016.07.002.
48. Bastani, H. Generative AI without guardrails can harm learning: Evidence from high school mathematics / H. Bastani, O. Bastani, A. Sungu [et al.] // Proceedings of the National Academy of Sciences. 2025. Vol. 122, no. 26. Article e2422633122. DOI: 10.1073/pnas.2422633122.
49. Barcaui, A. ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention / A. Barcaui // Social Sciences & Humanities Open. 2025. Vol. 12. Article 102287. DOI: 10.1016/j.ssaho.2025.102287.
50. Makransky, G. Beyond the “Wow” Factor: Using Generative AI for Increasing Generative Sense-Making / G. Makransky, B. M. Shiwalia, T. Herlau, S. Blurton // Educational Psychology Review. 2025. Vol. 37. Article 60. DOI: 10.1007/s10648-025-10039-x.
51. Yan, D. Independent judgment before viewing generative AI recommendations reduces susceptibility to misleading advice: A randomized scenario-based experiment with preservice teachers / D. Yan // Frontiers in Psychology. 2026. Vol. 17. Article 1936522. DOI: 10.3389/fpsyg.2026.1936522.
52. Qiang, Z. Between Human and Machine: University Teachers’ Accounts of Judgment, Metacognitive Defense, and Cognitive Offloading in Generative AI Use / Z. Qiang, Z. Wang // Journal of Intelligence. 2026. Vol. 14, no. 10. Article 254. DOI: 10.3390/jintelligence14100254.
53. Boolzen, C. Evidence of impact and interpretational limits of generative AI in STEM education: A systematic review and meta-analysis on cognitive learning outcomes / C. Boolzen, J. Kuhn, S. Flegr [et al.] // Artificial Intelligence Review. 2026. Vol. 59. Article 223. DOI: 10.1007/s10462-026-11665-9.
54. National Center for Science and Engineering Statistics. (2026). Doctorate Recipients from U.S. Universities: 2024. NSF 26-315. https://ncses.nsf.gov/pubs/nsf26315/report/postgraduation-trends
55. National Center for Science and Engineering Statistics. (2026). Universities Report 8.1% Growth in R&D Expenditures in FY 2024, Reaching Over $117 Billion. NSF 26-305. https://ncses.nsf.gov/pubs/nsf26305
56. Statistics Canada. (2025). Spending on Research and Development in the Higher Education Sector, 2023/2024. The Daily, 17 October. https://www150.statcan.gc.ca/n1/daily-quotidien/251017/dq251017e-eng.htm

57. Statistics Canada. (2025). University Revenues Increase More Than Expenditures in 2023/2024. The Daily, 29 July. https://www150.statcan.gc.ca/n1/daily-quotidien/250729/dq250729b-eng.htm
58. Advisory Panel on the Federal Research Support System. (2023). Report of the Advisory Panel on the Federal Research Support System. Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/panel-federal-research-support/en/report-advisory-panel-federal-research-support-system
59. OECD. (2024). Survey of Adult Skills 2023: Canada, Country Note. https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/canada_5ecab9d9-en.html
60. Bolton, P., and Lewis, J. (2026). Higher Education Finances and Funding in England. House of Commons Library, Research Briefing CBP-10037, 4 September. https://commonslibrary.parliament.uk/research-briefings/cbp-10037/
61. Office for Students. (2026). Financial Sustainability of Higher Education Providers in England: 2026. 14 May 2026, updated report. https://www.officeforstudents.org.uk/publications/financial-sustainability-of-higher-education-providers-in-england-2026/
62. OECD. (2024). Survey of Adult Skills 2023: England (United Kingdom), Country Note. https://www.oecd.org/en/publications/survey-of-adults-skills-2023-country-notes_ab4f6b8c-en/united-kingdom_02bc78e4-en.html
63. Department for Education. (2023). PISA 2022: National Report for England. Research Report RR1392. https://www.gov.uk/government/publications/pisa-2022-national-report-for-england
64. Colby, G. (2025). Data Snapshot: Tenure and Contingency in US Higher Education, Fall 2023. Academe, 111 (Spring). American Association of University Professors. https://www.aaup.org/article/data-snapshot-tenure-and-contingency-us-higher-education-fall-2023
65. National Academies of Sciences, Engineering, and Medicine. (2018). Graduate STEM Education for the 21st Century. Washington, DC: The National Academies Press. https://doi.org/10.17226/25038