

AI AND LEGAL INEQUALITIES IN HIGHER EDUCATION COURSES

R. GEORGE WRIGHT*

The rise of artificial intelligence (AI) elements in higher education coursework has resulted in new and complex forms of legally actionable inequality. Some such forms of inequality in the higher education process are obvious. Others are less obvious, but extremely important. Herein, we focus mainly upon the latter, more subtle, ways in which AI tends to increase unjustified inequalities in the pedagogy of higher education.

Certainly, there are ways in which AI may diminish existing inequalities in higher education. Consider, for example, historic forms of inequalities in meaningful access to higher education faced, in different ways, by various groups. There are obvious ways in which persons with limited movement, various physical impairments, and a wide range of neurological and psychological conditions might benefit substantially from an AI academic tutor, as opposed to, or in addition to, a human tutor.¹ AI is, of course, incapable of either patience or impatience. But AI tutors can certainly be experienced as not only instantly available at any hour, but as infinitely patient, and as supportive and encouraging as it is instructed to be.²

Equally clearly, AI may benefit higher education students who are not yet fluent in the university's primary language of instruction. AI's abilities to translate various languages into academically appropriate English, for example, continues to improve.³ AI can, in this sense, reduce some barriers to higher education success,⁴ and more generally, foster genuine inclusiveness in higher

* Professor of Law, Indiana University Robert H. McKinney School of Law.

1. See, e.g., Lauren Bialystock, *AI and the Future of (Philosophy of) Education*, 76 EDUC. THEORY 133, 134 (2026) (citing Beth McMurtrie, *AI to the Rescue*, CHRON. HIGHER EDUC. (June 20, 2025), www.chronicle.com/special-projects/the-different-voices-of-student-success/ai-to-the-rescue [<https://perma.cc/7WP3-89NT>]); Rachel Curry, *People with ADHD, autism, dyslexia say AI agents are helping them succeed at work*, CNBC (Nov. 8, 2025, at 9:00 EST), www.cnbc.com/2025/11/08/adhd-autism-dyslexia-jobs-careers-ai-agents-success.html? [<https://perma.cc/P7Z9-KZ3Z>]; Tayiba Rasheed et al., *Leveraging AI to Mitigate Educational Inequality: Personalized Learning Resources, Accessibility, and Student Outcomes*, 3 CRITICAL REV. SOC. SCI. STUD. 2399, 2400 (2025) (“[A]rtificial intelligence enhances accessibility by facilitating support for learners with disabilities, language learners, and those who have limited exposure to qualified instructors.”).

2. Thus, opening the door to, among other problems, what is known as AI sycophancy. See *infra* note 20 and accompanying text. For the moment, consider the positive and negative dimensions of one student's report that AI “was so patient . . . I don't think anyone has ever paid such pure attention to me and my thinking and my questions.” D. Graham Burnett, *Will the Humanities Survive Artificial Intelligence?*, NEW YORKER (Apr. 26, 2025), www.newyorker.com/culture/the-weekend-essay/will-the-humanities-survive-artificial-intelligence [<https://perma.cc/TN3U-LJ3S>].

3. See, e.g., Dragan Filimonovic, Christian Rutzer & Conny Wunsch, *Can GenAI Improve Academic Performance? Evidence From the Social and Behavioral Sciences* 19 (Oct. 6, 2025), <https://arxiv.org/abs/2510.02408> [<https://doi.org/10.48550/arXiv.2510.02408>] (“GenAI tools help reduce . . . structural frictions, such as linguistic barriers . . . particularly among disadvantaged groups.”); Rasheed et al., *supra* note 1, at 2400.

4. See NAESP Staff, *Generative AI in Education: Early Adoption, Equity, and the Road Ahead*, NAT'L ASS'N OF ELEMENTARY SCH. PRINCIPLES (Oct. 4, 2025), <https://www.naesp.org/resource/generative-ai-in-education-early-adoption-equity-and-the-road-ahead/> [<https://perma.cc/7WP3-89NT>].

education.⁵

Consider as well matters of pricing and inequality. We may assume that the overall price of a largely AI-based higher education program will continue to decrease, inevitably putting increased competitive price pressure on higher education programs with lesser reliance on AI.⁶ This development would, through reducing tuition costs to students, further expand realistic access, across economic classes, to higher education.⁷

As well, some features of a largely AI-based higher education may help all students, across the board, while helping some or all disadvantaged students in distinctively greater measure. AI programs can thus provide specifically, and even individually, targeted subject-matter, motivational, and emotional support, especially to students with various sorts of disadvantaged backgrounds.⁸ AI's ability to facilitate learning, in private and at the student's own pace and schedule, may benefit all students to some degree, but, again, to the benefit of many disadvantaged students especially.⁹ In particular, students with demanding commutes from rural areas, or with demanding work or family commitments, may distinctively benefit from immediate and continuous access to, and substantive feedback from, AI programs.¹⁰

These elements help make up the positive side of the ledger on AI's impact on inequality in higher education. There is, however, a negative side of the ledger, in which AI perpetuates, exacerbates, and creates significant inequalities

cc/8UHH-STHZ]; Cornelia C. Walther, *The Future of Education: Will AI Be the Great Equalizer?*, FORBES (Apr. 22, 2025, at 14:07 EDT), <https://www.forbes.com/sites/cornelia-walther/2025/03/16/the-future-of-education-will-ai-be-the-great-equalizer/> [https://perma.cc/A38Q-ELCF] ("AI-driven tools can . . . bridge gaps in educational access.").

5. See Walther, *supra* note 4.

6. See *id.*; Aviva Legatt, *AI, Cost, and Confidence are Rewriting Higher Ed's Value Proposition*, FORBES (Oct. 15, 2025, at 15:34 EDT), <https://www.forbes.com/sites/avivalegatt/2025/10/15/ai-cost-and-confidence-are-rewriting-higher-eds-value-proposition/> [https://perma.cc/N7AW-C8ZG].

7. We set aside here the possibility of AI reducing the sheer time required to complete any given higher education program. On financial cost issues, see Walther, *supra* note 4.

8. See, e.g., Blanka Klimova & Marcel Pikhart, *Exploring the Effects of Artificial Intelligence on Student and Academic Well-Being in Higher Education: A Mini-Review*, 16 FRONTIERS IN PSYCH., Feb. 3, 2025, at 3; Hoang Pham et al., *How will AI Impact Racial Disparities in Education?*, STAN. L. SCH. (June 29, 2024), <https://law.stanford.edu/2024/06/29/how-will-ai-impact-racial-disparities-in-education/> [https://perma.cc/93E9-KNCH] ("AI can be used to reduce the harmful impacts of . . . disparities by personalizing learning based on students' needs.").

9. See the authorities cited in *supra* note 8; see also *How artificial intelligence in education is transforming classrooms*, SMU SIMMONS SCH. EDUC. & HUM. DEV. (Jan. 9, 2025), <https://learningsciences.smu.edu/blog/artificial-intelligence-in-education> [https://perma.cc/4LKC-246F] ("Students can learn at their own speed, receive immediate feedback, and focus on areas they need to improve.").

10. See *id.*; Lucille Favero et. al., *Do AI Tutors Empower or Enslave Learners? Toward a Critical Use of AI in Education* 1 (July 9, 2025), <https://arxiv.org/pdf/2507.06878> [https://doi.org/10.48550/arXiv.2507.06878]. At the moment, many relevant studies are confined to small sample sizes, short term effects, and the inability to rule out confounders such as sheer novelty and extra institutional attention. See Anjali Singh et al., *Protecting Human Cognition in the Age of AI* 11 (Sep. 29, 2025), <https://arxiv.org/html/2502.12447> [https://doi.org/10.48550/arXiv.2502.12447].

in higher education. Some of these dimensions are obvious, but others less so. It is to the negative side of the ledger to which we now turn, with our initial focus being on the most obvious ways in which AI may contribute to inequalities in higher education coursework.

A moment's thought suggests that AI in general, and typical AI higher education programs, inevitably foreground certain beliefs, values, and vocabularies that are not universally shared within and across cultures. We should expect dominant cultures to exert corresponding media and academic publication dominance. As a default, dominant cultures should tend to have a disproportionate presence, if not dominance, in many AI higher education programs. This disproportionate representation of dominant cultures in AI training materials and curricular programs should be expected to obtain unless distinctive, conscientious efforts are undertaken to the contrary. All else equal, AI's content derives from word predictivity, and word predictivity depends upon overall frequency and volume of word appearance, which reflects patterns of media dominance.¹¹

Thus in general, "AI systems are trained on massive datasets that reflect existing social and cultural patterns—some of which are biased."¹² All else equal, AI responses to students tend to reflect the cumulative, perhaps unintended, cultural biases absorbed into the training sets and broader cultural output from which those training sets are constructed.¹³ Worse, AI training models "may even amplify existing societal biases embedded in their training data."¹⁴ AI, as it is officially adopted, may legitimize such biases. More subtly, any deliberate, but inevitably imperfect, remedial human intervention at any point in the overall AI generation process may incorrectly imply a lack of any remaining bias in the student-facing AI program.¹⁵

The biases in question can take effect in many forms and contexts and can have prejudicial effect along many dimensions, including, but not limited to, disability, race, ethnicity, sex and sexual orientation, economic class, language,

11. See, e.g., Jennifer Monahan, *Artificial Intelligence, Explained*, CARNEGIE MELLON UNIV. HEINZ COLL. (July 2023), <https://www.heinz.cmu.edu/media/2023/July/artificial-intelligence-explained> [<https://perma.cc/77GK-RQCF>]; Cole Stryker & Eda Kavlakoglu, *What is artificial intelligence (AI)?*, IBM (Aug. 9, 2024), <https://www.ibm.com/think/topics/artificial-intelligence> [<https://perma.cc/94AU-Y5B9>].

12. Ind. Univ. Kelley Sch. Bus., *AI Playbook: Principles, Practices, and Policies for Faculty* 8 (Oct. 2025), https://kelley.iu.edu/Kelley_AI_Playbook.pdf [<https://perma.cc/Q3CC-HWBF>].

13. See Xin Zhang et al., *A Systematic Literature Review of Empirical Research on Applying Generative Artificial Intelligence in Education*, 1 FRONTIERS DIGIT. EDUC. 223, 235 (2024) (referring to findings that "an AI-powered educational platform inadvertently perpetuated social biases present in its training data"); Giovanni Vindigni, *Data-Driven Disparities: How AI Applications in Education May Perpetuate or Mitigate Inequality*, 1 EUR. J. SCI. & MOD. TECH., Aug. 2025, at 1 (on "algorithmic bias" as opposed to "an impartial algorithmic design").

14. Chenyu Hou, Gaoxia Zhu & Vidya Sudarshan, *The Role of Critical Thinking on Undergraduates' Reliance Behaviours on Generative AI in Problem-Solving*, 56 BRIT. J. EDUC. TECH. 1919, 1935 (2025).

15. See Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 674 (2016) ("[E]ven honest attempts to certify the absence of prejudice on the part of those involved in the data mining process may wrongly confer the imprimatur of impartiality.").

culture, and national origin.¹⁶ Dominant beliefs, values, vocabularies, and policy concerns tend to arrive as the AI pedagogical defaults.¹⁷ And even more obviously, higher education students may have formally equal access to AI programs, but remain substantively disadvantaged by any lesser familiarity, let alone any lesser fluency, with the sheer mechanics of AI learning programs and processes.¹⁸

There are, however, also a range of less obvious, but important, aspects of inequality in higher education AI. These important dimensions of inequality may be more difficult to detect in practice, isolate, and optimally redress. The difficulties in meaningfully addressing these dimensions of inequality reflect their rootedness in much broader forms of cultural and economic discrimination and inequality.

Consider the concrete, lived experiences of historically disadvantaged students entering into higher education, perhaps as their family's first such generation, alongside a cohort of advantaged students, for whom higher education has perhaps been a familiar family experience. The norms, expectations, and experiences of higher education in general are more familiar to the latter group than to the former. And now, there is the additional element of a largely AI-dominant curricular experience. Should we expect both disadvantaged and advantaged groups of students to be off to an equal start, as the semester unfolds, in taking full advantage, in practice, of the wide range of AI-educational possibilities? Should we expect both groups to equally avoid the

16. See Mohammad Atari et al., *Which Humans?*, HARV. UNIV. (2024), https://coevolution.fas.harvard.edu/sites/g/files/omnuum5841/files/culture_cognition_coevol_lab/files/which_humans_09222023.pdf [<https://perma.cc/XX4N-4Q3S>]. See also *id.* at 15 (“When researchers claim that LLMs give ‘human’-like responses, they need to specify which humans they are talking about.”); Sahan Bulathwela et al., *Artificial Intelligence Alone Will Not Democratise Education: On Educational Inequality, Techno-Solutionism, and Inclusive Tools*, 16 SUSTAINABILITY, Jan. 2024, at 5–6.

17. For a sampling of the expressed concern, see, e.g., Earle Abrahamson, *AI in Higher Education: Bridging the Divide Between Access, Equality, and Opportunity*, INT’L CTR. FOR ACAD. INTEGRITY (Mar. 10, 2025), https://academicintegrity.org/aws/ICAI/pt/sd/news_article/587198/_PARENT/layout_interior/false [<https://perma.cc/9SMY-JGNZ>]; Donna Campbell, *AI enters the classroom as law schools prep students for a tech-driven practice*, NAT’L JURIST (July 9, 2025), <https://nationaljurist.com/ai-enters-the-classroom> [<https://perma.cc/74VF-42WD>] (“You can’t ignore the equity issues.”); Josie Cox, *The most radical act of feminism? Using AI*, THE TIMES (Sep. 10, 2025, at 12:20 BST), https://www.thetimes.com/us/news-today/article/the-most-radical-act-of-feminism-using-ai-gpx2nxtc3?gaa_at=eafs&gaa_n=AWetsqeF6jleyN9Y7cNgmEvHwG0cWqx7wfJbBz643HaDtF6TiPmyvq4nw3Wxmak5Rfw%3D&gaa_ts=6990e9b8&gaa_sig=1K-Ezos3z8IMP-Scs2UCQVrqT9vjhTbykhsnAKRnk4d07yz-kHXXcwVa4KzhfZdFlbz-DJnueU2CTmJWZKulWg%3D%3D [<https://perma.cc/HL8K-X5T5>] (referring to “all the sexism, racism and bias that’s embedded in the technology”); Josh Freeman, *Student Generative AI Survey 2025*, HIGHER EDUC. POL’Y INST. (Feb. 26, 2025), <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/> [<https://perma.cc/2H5J-L6V6>] (finding male students and “more socioeconomically advantaged students more likely to use AI than others”); Pham et al., *supra* note 8; Mohan Yang et al., *Analysing Nontraditional Students’ ChatGPT Interaction, Engagement, Self-Efficacy and Performance: A Mixed-Methods Approach*, 56 BRIT. J. EDUC. TECH. 1973, 1975 (2025) (new AI-based educational technologies “risk widening the existing digital divide” if not thoughtfully deployed and utilized).

18. See the authorities cited in *supra* note 17.

non-obvious traps, pitfalls, and suboptimal uses of academic AI? In particular, should we expect each group, with or without specific training, to equally quickly master the subtle arts of devising the best possible succession of precisely formulated AI user prompts, for all purposes and contexts?

Certainly, many of the less obvious harms of AI use in higher education will afflict all groups of students, to one degree or another. But there are likely to be what the law refers to as disparate impacts¹⁹ among different groups of students. Higher educational institutions, and the law, should continuously explore and attempt to remedy any such disparate impact.

Consider first some of the subtle, but important, effects of what is known as AI sycophancy.²⁰ At a first approximation, AI sycophancy refers to AI responses

19. For useful background on disparate impact discrimination, as distinct from disparate treatment or intentional discrimination, see, e.g., *Dep't of Homeland Sec. v. Regents of the Univ. of Cal.*, 591 U.S. 1, 34 (2020); *EEOC v. Abercrombie & Fitch Stores, Inc.*, 575 U.S. 768, 781–82 (2015) (Thomas, J., concurring in part and dissenting in part) (“[D]isparate impact claims . . . involve employment practices that are facially neutral in their treatment of different groups but that in fact fall more heavily on one group than another and cannot be justified by business necessity.” (alteration in original) (quoting *Raytheon Co. v. Hernandez*, 540 U.S. 44, 52 (2003))). Importantly, disparate impact and disparate treatment theories differ as to their elements, liability standards, defenses, and remedies. See, e.g., *Young v. UPS*, 575 U.S. 206, 247–48 (2015) (Scalia, J., dissenting) (citing 42 U.S.C. §§ 1981a, 2000e-2k, 2000e-5g) (noting the availability, by statute, of compensatory, punitive, and equitable relief in disparate treatment Title VII cases, but only injunctive or other equitable relief in disparate impact Title VII cases). Disparate impact claims are also possible in some federal housing and age discrimination civil rights contexts. See, e.g., *Tex. Dep't of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519 (2015). For a sense of some general problems and possibilities associated with AI and disparate impact theories, see Christine Polek & Shastri Sandy, *The Disparate Impact of Artificial Intelligence and Machine Learning*, 21 COLO. TECH. L.J. 85 (2023) (focusing on business employees and business customers in particular); Vincent Calderon, *Unintentional Algorithmic Discrimination: How Artificial Intelligence Undermines Disparate Impact Jurisprudence*, 24 DUKE L. & TECH. REV. 28 (2024) (noting the often difficult problems of sufficiently proving causation of a harm, and of offering a less discriminatory and otherwise satisfactory alternative business practice). See also Andrew D. Selbst, *Disparate Impact in Big Data Policing*, 52 GA. L. REV. 109 (2017) (creatively advocating for what amounts to algorithmic impact statements, by analogy to required environmental impact statements, and for a notice-and-comment process allowing for persons to weigh in prior to government adoption of algorithmic-driven policing technologies). There are, as well, questions as to the availability of a private right of action in Title VII and other disparate impact contexts. See *Alexander v. Sandoval*, 532 U.S. 275, 293 (2001).

20. For background, particularly in educational contexts, see Malihe Alikhani, *Breaking the AI Mirror*, BROOKINGS (Apr. 15, 2025), <https://www.brookings.edu/articles/breaking-the-ai-mirror/> [<https://perma.cc/D6KH-3CZ7>] (describing “AI sycophancy” as affirming, rather than challenging, a user’s misleading statistics, thus reinforcing the user’s misunderstanding); Phillip Alcock, *The Sycophancy Problem*, A.I. FOR BEGINNERS (July 31, 2025), <https://aiforbeginners.substack.com/p/the-sycophancy-problem> [<https://perma.cc/P3HF-YAPT>] (highlighting how sycophantic AI emphasizes the prompter’s strengths while “glossing over obvious problems”); Simar Bajaj, *Next Time You Consult an AI Chatbot, Remember One Thing*, N.Y. TIMES (Sep. 29, 2025), <https://www.nytimes.com/2025/09/26/well/is-ai-validation-healthy.html> [<https://perma.cc/Z92F-WD6Y>] (referring to AI non-judgmentalism both where appropriate and where inappropriate, sheer person-oriented relativism and subjectivism, and the process of psychologically rewarding users with an amplified version of social media “likes”); Rebecca Bellan, *AI sycophancy isn’t just a quirk, experts consider it a ‘dark pattern’ to turn users into*

that would, if they were actually intended, count as flattering, ingratiating, toadying, or similarly manipulative. We may begin with, but then expand upon, that understanding.

AI sycophancy is understandable from the standpoint of the competitive market needs of AI program suppliers, and from the standpoint of AI users, particularly under distinctive circumstances. A vicious circle can be seen to develop. Private AI program sellers face a competitive market imperative to develop programs that schools or students are willing to pay for. User adoption and loyalty may depend upon subjective user experiences. And positive user experiences may stem disproportionately from AI's sycophantic relationships with users.²¹ User feedback can then promote and further incentivize AI sycophancy by highly ranking, whether consciously or not, instances of AI sycophancy, as distinct from appropriately challenging AI responses.²² So, unless the AI programs in question are appropriately prompted, usually in some non-obvious way, they may not appropriately prioritize a meaningful, arms-length broadening of a user's pre-existing judgments and beliefs.²³ There is thus

profit, TECH CRUNCH (Aug. 25, 2025, at 13:50 PST), <https://techcrunch.com/2025/08/25/ai-sycophancy-isnt-just-a-quirk-experts-consider-it-a-dark-pattern-to-turn-users-into-profit/> [<https://perma.cc/3SSY-6CE9>] (a chatbot's 'rewarding' pseudo-interactionism "can make people feel understood or cared for"); Mike Caulfield, *AI Is Not Your Friend*, ATLANTIC (May 9, 2025), <https://www.theatlantic.com/technology/archive/2025/05/sycophantic-ai/682743/> [<https://perma.cc/V8AQ-7PJJ>] (one bot version in particular "wouldn't stop telling users how brilliant their bad ideas were"); Yaqub Chaudhary & Jonnie Penn, *Beware the Attention Economy: Collection and Commodification of Intent Via Large Language Modules*, HARV. DATA SCI. REV., Dec. 30, 2024, at 2 (discussing LLM "sycophancy, ingratiation, and emotional infiltration"); Myra Cheng et al., *Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence* 2 (Oct. 1, 2025), <https://arxiv.org/abs/2510.01395> [<https://doi.org/10.48550/arXiv.2510.01395>] ("[P]articipants rated sycophantic responses as higher quality, trusted the sycophantic model more, and were more willing to use it again"); *id.* at 3 (citing an AI response of "You did what [was] right for you"); *id.* at 11 ("[W]e find a clear user preference for AI that provides unconditional validation."); Lihua Du et al., *Alignment Without Understanding: A Message-and-Conversation-Centered Approach to Understanding AI Sycophancy* 2 (Sep. 25, 2025), <https://arxiv.org/html/2509.21665v1> [<https://doi.org/10.48550/arXiv.2509.21665>] (on AI program encouragement of "over-reliance and emotional dependency, weakening resilience, and diminishing autonomous judgment"); Shubhangi Dua, *AI Agents Will Know Your 'Intentions' Before You Do. It's Called the 'Intention Economy'*, ENTER. MGMT. 360 (Feb. 28, 2025), <https://em360tech.com/tech-articles/ai-agents-will-know-your-intentions-you-do-its-called-intention-economy> [<https://perma.cc/QRK3-QD65>] (on building AI user comfort and trust); Erie Meyer et al., *Tech Brief: AI Sycophancy & Open AI*, GEO. L. (July 30, 2025), <https://www.law.georgetown.edu/tech-institute/research-insights/insights/tech-brief-ai-sycophancy-openai-2/> [<https://perma.cc/7WXL-2S7B>]; Amy Winecoff, *Artificial Sweeteners: The Dangers of Sycophantic AI*, TECH. POL'Y PRESS 2 (May 14, 2025), <https://www.techpolicy.press/artificial-sweeteners-the-dangers-of-sycophantic-ai/> [<https://perma.cc/FJ3M-4NVX>] (AI can be "too quick to validate and too hesitant to challenge users").

21. See Meyer et al., *supra* note 20.

22. See Miryam Naddaf, *AI chatbots are sycophants—researchers say it's harming science*, NATURE (Oct. 24, 2025), <https://www.nature.com/articles/d41586-025-03390-0> [<https://perma.cc/4GBS-U2Y8>].

23. See Ian Sample, *Sycophantic AI chatbots tell users what they want to hear, study shows*, THE GUARDIAN (Oct. 24, 2025, at 11:00 EDT), <https://www.theguardian.com/technology/2025/>

a perverse drive toward agreeableness rather than optimal long-term educational value in higher education AI programs.²⁴

These tendencies tend, to some degree, to cut across economic class and other social group lines. But they are likely, one might suspect, to disparately impact groups with historically less advantage and less meaningful familiarity with institutions of higher education. All higher education students desire some sort of balance between AI's apparent, if unreal, warmth, encouragement, validation, and acceptance on the one hand, and sheer AI accuracy, completeness, expertise, boundary-pushing, challenge, and critique on the other.²⁵

But we might expect, in general, students from outsider groups to seek more of what is subjectively experienced as validation than students who feel themselves to be relatively more at home in the higher educational institution environment. And this and related disparities open the door to disparate adverse educational consequences of AI sycophancy, to the greater prejudice, generally, of those groups less familiar with higher educational institution environments.²⁶

Consider now a further, non-obvious inequality dimension. This problem can disparately affect all disadvantaged groups, but for clarity, let us focus in particular on students who are in the process of learning English, their main language of instruction, as a second or subsequent language. Any reasonable prompt from such a student may well generate, instantaneously, an impressive looking, apparently polished document, perhaps suitable with minimal further student input, for submission as a creditable completion of a course assignment.²⁷ From the perspective of such a student, the document in question, perhaps with minor tinkering, seems to be, in itself, better than whatever the limited English speaker would be likely to generate after hours of personal effort.²⁸

As a general matter, however, any such student reaction short-circuits a meaningful overall educational process, disproportionately so for disadvantaged student groups. As one study has indicated, “when students use ChatGPT to generate complete assignments, rather than for assistance such as suggesting revisions for their own critical review, they miss essential learning

oct/24/sycophantic-ai-chatbots-tell-users-what-they-want-to-hear-study-shows [https://perma.cc/Y7TQ-SY4Z].

24. See Amy Lotz, *AI sycophancy: The downside of a digital yes-man*, AXIOS (July 7, 2024), www.axios.com/2025/07/07/ai-sycophancy-chatbots-mental-health [https://perma.cc/3FTA-ZW4K].

25. See Shaper Richter et al., *Chatbots in Tertiary Education: Exploring the Impact of Warm and Competent Avatars on Self-Directed Learning*, 56 BRIT. J. EDUC. TECH. 2102 (2025).

26. See *supra* note 20 and accompanying text.

27. See Nathalie A. Smuha, *The Broader Picture: Generative AI, Education, and the Future of Democracy*, ICONNECT (Oct. 17, 2025), <https://www.iconnectblog.com/the-broader-picture-generative-ai-education-and-the-future-of-democracy/> [https://perma.cc/UZ5Z-2LE8].

28. See *id.*

opportunities.”²⁹ Given their objectively impressive initial AI-generated product, disadvantaged students may tend to see relatively lesser value in any further succession of multiple, precisely formulated, carefully calibrated, and sensitive responses to each iteration of AI-generated product, including invitations to meaningfully refine, challenge, or critique the student’s own assumptions and conclusions.³⁰

Unfortunately, any group differences in this short-circuiting the optimal reflection, research, writing, revising, and editing process are likely to lead to disparities in actual educational outcomes, whether the educational institution measures them or not. At this point, we certainly cannot be sure about the nuances of any such effects, their precise lines of causation, their magnitudes, and their permanence.³¹

However, there is already some evidence that if AI is not optimally structured and used in education, it may promote serious student over-reliance effects,³² weaker argumentative reasoning,³³ lower student engagement with content,³⁴ reduced critical thinking,³⁵ excessive off-loading of cognitive responsibilities,³⁶ reduced ability to remember even one’s own work-product,³⁷ reduced creativity and innovation,³⁸ and even, perversely, an increase in student

29. Elizabeth Bauer et al., *Looking Beyond the Hype: Understanding the Effects of AI on Learning*, 37 EDUC. PSYCH. REV., Apr. 2025, at 10. For discussion in the law school context, see Steve Lorteau & Douglas Sarro, *Artificial Intelligence in Legal Education: A Scoping Review* (U. Ottawa Fac. L. Sch. Working Paper, Paper No. 2025-21, 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5762982.

30. For background, see Teresa Luther, Joachim Kimmerle & Ulrike Cress, *Teaming Up with an AI: Exploring Human-AI Collaboration in a Writing Scenario with ChatGPT*, 5 AI 1357, 1369 (2024); Lixiang Yan et al., *Beyond Efficiency: Empirical Insights on Generative AI’s Impact on Cognition, Metacognition and Epistemic Agency in Learning*, 56 BRIT. J. EDUC. TECH. 1675, 1677 (2025).

31. For some possible complications, see Ruiqi Deng et al., *Does ChatGPT Enhance Student Learning? A Systematic Review & Meta-Analysis of Experimental Studies*, 227 COMPUTS. & EDUC., Apr. 2025.

32. See Favero, *supra* note 10; Anastasia Berg, Opinion, *Why Even Basic A.I. Use Is So Bad for Students*, N.Y. TIMES (Oct. 29, 2025), <https://www.nytimes.com/2025/10/29/opinion/ai-students-thinking-school-reading.html> [<https://perma.cc/MJV9-2LV4>].

33. See, e.g., Lybi Ma, *AI Eases Our Mental Load at the Expense of Critical Thinking*, PSYCH. TODAY (Oct. 30, 2024), <https://www.psychologytoday.com/us/blog/the-art-of-critical-thinking/202410/ai-eases-our-mental-load-at-the-expense-of-critical> [<https://perma.cc/8DF6-SLP6>].

34. See *id.*

35. See *id.*; see also Michael Gerlich, *AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking*, 15 SOCIETIES, Sep. 2025, at 12; Roman Yavich, *Will the Use of AI Undermine Students Independent Thinking?*, 15 EDUC. SCIS., May 2025, at 13–14.

36. See Ma, *supra* note 33.

37. See especially the oft-cited Nataliya Kosmyna et al., *Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant For an Essay Writing Task* 19 (June 10, 2025), <https://arxiv.org/pdf/2506.08872> [<https://doi.org/10.48550/arXiv.2506.08872>].

38. See Chunpeng Zhai, Santoso Wilbowo & Lily D. Li, *The Effects of Over-Reliance on AI Dialogue Systems on Students’ Cognitive Abilities: A Systematic Review*, 11 SMART LEARNING ENV’T, June 18, 2024, at 4.

anxiety and stress levels.³⁹

Some students, certainly, are to some degree aware of one or more harms, as well as the benefits, of their experiences with educational AI.⁴⁰ But it is reasonable to expect that students less familiar with the possibilities of higher education programs to have a lesser sense of what their own optimal higher educational experience might concretely involve. And in practice, greater degrees of initial stress, anxiety, and time constraints may reinforce inequalities in the disparately adverse effects of AI in higher education.⁴¹

In the background, the likely increasing cost advantages of AI in higher education virtually ensure its widespread embrace.⁴² Guiding disadvantaged students in particular through the processes of, e.g., genuinely optimal AI prompt formulation and adaptive revision⁴³ would require individual and small group human-led tutorials and counseling and would be costly. Inevitably, developing an AI-infused higher education program that respects meaningful equal opportunities as among student groups will require substantial investment by educational institutions as well as by the students.

In some measure, costs incurred to reduce disparities in opportunities in higher education reflect inequalities of educational opportunity already experienced by the affected students at earlier grade levels. Many students arrive at higher educational institutions without ever having had a meaningful opportunity to develop the storehouses of knowledge⁴⁴ and the corresponding educationally oriented memory⁴⁵ needed to take full advantage of advanced

39. See Juan Garzon, Eddy Patino & Camilio Marulanda, *Systematic Review of Artificial Intelligence in Education: Trends, Benefits, and Challenges*, 9 MULTIMODAL TECH. & INTERACTION, Aug. 2025, at 14.

40. See Richard Adams, *Pupils fear AI is eroding their ability to study*, THE GUARDIAN (Oct. 15, 2025), <https://www.theguardian.com/technology/2025/oct/15/pupils-fear-ai-eroding-study-ability-research> [<https://perma.cc/Y9VQ-43GY>]; Helen Pearson, *Will AI Make Students Smarter, Or Stop Them From Thinking at All?*, 646 NATURE 788, 790 (2025).

41. For the possibilities of either greater educational equality or else the creation of an educational underclass consigned to educational screens without appropriate guidance, see Daniel Mollenkamp, *Will AI Shrink Disparities in Schools, or Widen Them?*, EDSURGE (Aug. 19, 2024), <https://www.edsurge.com/news/2024-08-19-will-ai-shrink-disparities-in-schools-or-widen-them> [<https://perma.cc/K22K-FWFY>].

42. See Walther, *supra* note 4.

43. See the authorities cited *supra* note 30.

44. See, e.g., E.D. HIRSCH, JR., *WHY KNOWLEDGE MATTERS* 83 (2016) (“When looking *anything* up, there is simply no substitute for possessing a lot of knowledge relevant to the subject.”).

45. See ANTHONY T. KRONMAN, *EDUCATION’S END: WHY OUR COLLEGES AND UNIVERSITIES HAVE GIVEN UP ON THE MEANING OF LIFE* 50 (2007). Professor Kronman argues that “[t]o have a well-stocked memory is to be well-equipped for the challenges of life . . . as we face our choices and evaluate the alternatives before us.” In the AI educational context, see Brian Klaas, *The Death of the Student Essay—and the Future of Cognition*, at n.4 (June 19, 2025), www.forkingpaths.co/p/the-death-of-the-student-essay [<https://perma.cc/Y8FM-3TNP>] (“[T]hose people who already have advanced critical thinking skills can unleash even greater potential in collaboration with AI tools.”); Barbara Oakley et al., *The Memory Paradox: Why Our Brains Need Knowledge in an Age of AI*, in *THE ARTIFICIAL INTELLIGENCE REVOLUTION: CHALLENGES AND OPPORTUNITIES* (Max

educational programs. Mitigating these inequalities is clearly not a matter solely of tutoring students in the optimal use of AI prompts. Broader and greater educational efforts are needed. Institutional failure in this regard ensures that the educational experience will be less valuable, overall, for less privileged than for more privileged student groups.

One element of the problem, we may infer, is that disadvantaged students may tend disproportionately to fail to appreciate the differences between their perhaps institutionally acceptable work-product and what the product might have been. But as Barbara Oakley and her colleagues have observed, the real inequalities remain. Thus:

[L]earners still building foundational knowledge face a significant risk: mistaking AI fluency for their own. Without a robust internal framework for comparison, they may readily accept plausible-sounding output without realizing what's missing or incorrect. This bypasses the mental effort—retrieval, error detection, integration—that neuroscience shows is essential for forming lasting memory engrams and flexible schemas.⁴⁶

Professor Oakley concludes that “the result is a false sense of understanding: the learner feels accomplished, but the underlying work hasn’t been done.”⁴⁷

To be clear, this often unrecognized deprivation occurs whether the student’s submitted coursework contains any missing or incorrect representations or is in itself flawless. The same overestimation by the student of what has been accomplished obtains in either case. To the extent that any education at all relies in any respect on any such systematic misimpressions on the part of its students, they are thereby complicitous, at least, in an important contemporary form of what has been called “false consciousness.”⁴⁸

Realistically, it is understandable that an educational institution, like most other institutions, might not wish to dispel any unjustified favorable impressions as to its own real accomplishments. But in failing to clarify what its own disadvantaged students have, and have not, actually accomplished, the higher education institution betrays its own purported purposes and the educational interests of its disadvantaged students in particular.

Let us then assume that one or more groups of students have brought a plausible disparate impact claim against their institution of higher education.

Rangeley & Nicholas Fairfax eds., forthcoming 2026); Matthias Stadler, Maria Bannert & Michael Sailer, *Cognitive Ease at a Cost: LLMs Reduce Mental Effort But Compromise Depth in Student Scientific Inquiry*, 160 COMPUTS. HUM. BEHAV., July 30, 2024, at 2 (“Prior knowledge in the domain plays a crucial role; students with a deeper understanding of a topic can navigate information more strategically; . . . complex searches require a high degree of cognitive flexibility, metacognitive skills, and a robust working memory capacity.”).

46. Oakley et al., *supra* note 45, at 39.

47. *Id.*

48. For useful discussions, see, e.g., Steven Lukes, *In Defense of “False Consciousness”*, 2011 U. CHI. LEGAL F. 19; W.G. Runciman, *False Consciousness*, 44 PHIL. 303 (1969); David Miller, *Ideology and the Problem of False Consciousness*, 20 POL. STUD. 432 (1972).

We cannot here trace out all of the possible theories, remedies, and defenses to such claims.⁴⁹ The matrix of possibilities is expansive and partly indeterminate.⁵⁰ But several points can be made about litigating the inequalities associated with AI in higher education.

First, the law should recognize that in other higher education contexts, the distinction between a disparate impact claim and other, less controversial sorts of discrimination claims can be hazy.⁵¹ Second, all parties should recognize that flexible, sensitive, well-crafted injunctive relief is, at the level of the public interest, more valuable than specific awards of money damages to particular student plaintiffs in our kinds of cases. This recognition accords nicely with the more general legal sentiment that injunctive relief should often be the preferred remedy in disparate impact cases.⁵² And third, universities should recognize that particularly under the foreseeable economic market circumstances, a school's adversarialism, or even arms-length distancing, with respect to their less advantaged students is neither ethically appropriate nor institutionally prudent. Student success after graduation, and the flourishing, if not the survival, of the educational institution itself are increasingly inseparable. As a general matter, the university and any student groups disproportionately adversely affected by AI should jointly and cooperatively investigate, responsibly experiment with, and effectively reform the AI programs and practices at issue.

49. For background, see the authorities cited *supra* note 19.

50. For a much broader sense of some of the indeterminacies, see, e.g., Alexandra Brodsky, *Discrimination on the Basis of Consensual Sex*, 99 N.Y.U. L. REV. 1487, 1524–25 (2024); Jessica A. Clarke, *Sex Discrimination Formalism*, 109 VA. L. REV. 1699 (2023); Christopher S. Havasy, *Social Justice Conflicts in Public Law*, 113 CALIF. L. REV. 1315, 1380–81 (2025); Elaine McCabe et al., *Rights Under Title IX*, 26 GEO. J. GENDER & L. 967, 994 (2025); Tammi Walker, *No Department, No Enforcement: Title IX After the Collapse of the Department of Education*, 59 U.C. DAVIS L. REV. ONLINE 69, 77 (2025).

51. One might, for example, reconceive of a claimed intentional denial of constitutional equal protection, as in *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 600 U.S. 181 (2023), as involving an unintended disparate impact, on some groups, of university policies that are intended to promote overall fairness, equality, and other worthy university goals.

52. See, e.g., *Young v. UPS*, 575 U.S. 206, 241, 247–49 (2015) (Scalia, J., dissenting).