

The University's Unseen Hand: How Campus AI Contracts Benefit from Global Labor Exploitation

Weekly Analysis — <https://ainews.social>

When a provost signs a multi-year licensing agreement with an AI vendor, what exactly is being purchased? On paper, the line items are familiar: per-seat access, single sign-on integration, an enterprise tenancy with a privacy addendum, perhaps a faculty pilot stipend. The procurement office logs the spend under "instructional technology" or "academic productivity," the press office drafts a release about innovation and student success, and the contract joins the long shelf of vendor relationships that quietly run the modern campus. What does not appear on the line items is the actual substrate of the service — the click-workers in Nairobi labeling toxic content for a dollar an hour, the diesel generators powering an inference cluster in a water-stressed county, the rare earths pulled from the ground in the eastern Congo. The contract is mute on these. The university, by signing, becomes mute on them too.

This is a category of institutional self-deception worth naming precisely, because higher education has spent the last three years drifting into it at speed. According to the [14], enterprise adoption of generative AI has saturated the academic sector faster than any prior wave of educational technology, with universities concluding framework agreements with the major foundation-model vendors at a pace that outstrips their own governance capacity. The pedagogical conversation has rushed to keep up — syllabi rewrites, integrity policies, AI literacy modules — while the procurement conversation, the one that actually determines whose labor and whose ecosystems are being drawn down to power the seminar room, has barely begun. A recent survey reported by [2] found near-unanimous faculty concern about cognitive impact, but the same institutions whose faculty are alarmed continue to expand vendor footprints because the contracts are negotiated elsewhere, by people who do not teach.

The argument of this essay is that the gap between the pedagogical conversation and the procurement conversation is not an accident of organizational design. It is the operating logic of how universities are absorbing AI: a managerial hand that signs, a pedagogical hand that

[14] PDF hai.stanford.edu

[2] 90% Of Faculty Say AI Is Weakening Student Learning: How ... - Forbes

worries, and a global supply chain whose costs neither hand is asked to see. The institution as such — not the faculty senate alone, not the CIO alone, not the ethics committee alone — needs to be made to see them, because the institution as such is the entity that is signing.

What the Contracts Do Not Say

Begin with the documents themselves. A scan of standard enterprise AI procurement language across North American and European higher education institutions reveals a striking pattern: clauses on data residency, on FERPA compliance, on uptime guarantees, on indemnification for IP infringement — and almost nothing on the conditions of labor under which the model was trained, or on the energy and water intensity of the inference that the contract is purchasing at scale. The French analysis at [8] makes this point sharply for the grandes écoles: governance frameworks have rushed to address admissions algorithms and student-facing risks while leaving the upstream supply chain — the human raters, the data brokers, the compute providers — entirely outside the contractual perimeter. A university that would never knowingly buy coffee uncertified for fair-trade or apparel uncertified for sweatshop-free labor signs annual AI contracts with no analogous instrument.

This is not because the harms are unknown. Kate Crawford’s mapping in [22] has been on the shelf since 2021, and its central claim — that AI is “neither artificial nor intelligent” but a planetary extractive industry running on minerals, water, energy, and underpaid human attention — has had ample time to migrate from the seminar room to the procurement office. It has not. The Iberoamerican mapping published in [15] is even more pointed: institutions in the Global South are signing the same contracts as institutions in the Global North, on the same terms, even though the extractive footprint of those contracts falls disproportionately on their own regions. The asymmetry is structural; the contracts paper over it.

What the contracts also do not say, and what makes the procurement frame so politically thin, is that the labor underpinning a campus chatbot is not abstract. The MIT Press primer at [22] is blunt about this: the “free digital labor” of users producing training data, the wage labor of raters and moderators, and the service labor of those keeping the data centers running are all forms of value extraction that the legal instrument of a SaaS contract is structurally designed not to register. When a university signs, it is not just buying software; it is laundering a labor relation through a license agreement.

[8] IA et grandes écoles : quand un algorithme d’admission ...

[22] The Atlas of AI

[15] PDF La llegada de la IA a la educación superior en Iberoamérica: Un mapa ...

[22] AI Ethics - The MIT Press Essential Knowledge series

A Supply Chain Made Briefly Visible

Every few months, the supply chain surfaces in the press and then is allowed to sink again. We learn — through investigative journalism, through litigation, through whistleblowers — that the safety filters on a flagship model were tuned by workers in Kenya paid less than two dollars an hour to read child sexual abuse material for eight-hour shifts. We learn that an inference data center in Arizona is drawing groundwater from an aquifer the state has classified as overdrawn. We learn that the cobalt for the GPUs is mined under conditions the European Parliament has formally classified as slavery-adjacent. The pattern, traced carefully in [22], is that the visibility of these conditions is episodic by design: each new product cycle resets the narrative to "magic," and the labor and ecological substrate returns to invisibility until the next leak.

Universities have a particular obligation here that distinguishes them from other corporate buyers, and it is worth being precise about why. A bank that licenses an AI service is buying a productivity input. A university that licenses an AI service is integrating that service into a knowledge-producing institution whose social legitimacy depends on the claim that it tells the truth about the world. The piece at [13] makes a version of this argument from inside the NLP community itself: the field's "for social good" framing has tended to obscure rather than illuminate the conditions under which the tools are built, and the institutions that adopt them inherit that obscurity. A university that buys obscurity sells obscurity to its students.

Ruha Benjamin's framework in [22] sharpens the indictment further. The "New Jim Code," she argues, is the way technological systems encode discriminatory outcomes under a veneer of neutrality, and the way institutions adopt those systems precisely because the veneer absolves them of having to make explicit moral choices. A procurement contract is exactly such a veneer. It permits an institution to say "we did not choose to exploit Kenyan workers; we chose a vendor whose terms we accepted." This is, in moral terms, a distinction without a difference, and a university — of all places — should be capable of saying so.

The Pedagogical Alibi

The procurement story is uncomfortable, so it tends to be displaced by a more flattering story about pedagogy. In the displacement narrative, the university is adopting AI because doing so will personalize

[22] *The Atlas of AI*

[13] *Natural language processing for social good: Where we are, what is missing, and where we should go*

[22] *Race After Technology*

learning, scale tutoring, remediate inequality, and prepare graduates for a labor market that will demand AI fluency. The Stanford pilot reported in [16] is the canonical exhibit: human tutors paired with an AI assistant produced measurable learning gains for students of less-experienced tutors, and the result has been cited widely as evidence that AI in education is not merely defensible but morally obligatory. The companion writeup at [24] frames the work explicitly as a scaling-of-expertise intervention.

These results are real. They are also doing rhetorical work that exceeds what they can bear. A tutoring augmentation in a controlled K-12 study does not generalize to a campus-wide license of a general-purpose chatbot, and yet the citation chain runs from Tutor CoPilot to the provost's memo announcing an enterprise deal as if it did. Meanwhile, the contrary evidence has accumulated. The Stanford repository entry at [6] reports that unrestricted access to generative tools during practice phases of learning depresses subsequent unaided performance — students do better with the tool and worse without it, a classic offloading pattern. The BBC's accessible synthesis at [1] catalogues a growing clinical literature on what the Mexican medical-education researchers at [18] have named "metacognitive laziness" — the systematic erosion of the monitoring functions that distinguish learning from completion.

So the pedagogical alibi is on shaky empirical ground even taken on its own terms. But the deeper problem is that the alibi performs a moral substitution. It allows the institution to talk about whether AI helps students learn — a legitimate and important question — while refusing to talk about whose labor and whose water table is being drawn down to make the help available. The substitution is not innocent. It is the reason a faculty senate can debate AI for six hours without anyone mentioning Kenya.

Governance That Looks Inward, Pedagogy That Looks Outward

The imbalance is observable in the structure of the discourse itself. When universities convene AI governance bodies, the work tends to organize around inward-facing risks: academic integrity, FERPA, IP, hallucination liability, reputational exposure. The Quebec-based analysis at [10] notes that the institutional framing of generative AI risk has been almost entirely interior — what the technology does to our students, our integrity, our brand — with scarcely a sentence about what the technology does to the people and places that produce it. The French complement at [11] makes the same observation about

[16] PDF Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise

[24] Tutor CoPilot: A Human-AI Approach for Scaling Real-Time ... - ERIC

[6] Generative AI Can Harm Learning | SCALE Initiative

[1] 'Think outside the bots': How to stop AI from turning your brain ... - BBC

[18] Pereza metacognitiva y descarga cognitiva en la era de la IA generativa ...

[10] Impact de l'IA générative sur la « pensée critique »

[11] L'IA générative dans les études supérieures : entre facilitation ...

French higher education: the policy literature is dense with cheating-and-assessment talk and thin to absent on supply-chain ethics.

The Castlereagh Statement, summarized at [23], is one of the few institutional documents to attempt a wider frame, and even there the operational question — how do we translate principles about labor and ecological responsibility into the actual procurement workflow — is identified as the unfinished business. Castlereagh’s authors are honest that the discourse has reached the manifesto stage and stalled before the contract stage. The contract is where the institution acts. Until governance reaches the contract, governance is performance.

The asymmetry has a familiar shape to anyone who has followed the surveillance-technology debates of the last decade. Shoshana Zuboff’s analysis in [22] traced how institutions adopted data-extractive systems through the same procurement-and-policy split: the policy office issued privacy principles while the procurement office signed contracts that rendered the principles inoperative, and the gap between the two was where the extractive logic lived. The AI procurement story is the same story with the extraction moved one tier upstream — out of the user’s data and into the laborer’s body and the watershed’s aquifer. The institutional response so far has been to repeat the surveillance-era error in a higher-stakes form.

What the Pedagogical Failures Reveal

The failures of implementation, where they are visible, are diagnostically useful, because they reveal what the procurement frame was not designed to catch. The most visible class of failure is the cluster of AI-detection scandals that have produced an emergent body of litigation. The case at [3] is one of dozens now tracked at [4], and the pattern is consistent: an institution licenses a detection tool from a vendor, applies the tool’s output as if it were evidence, and then discovers — typically through a lawyer’s letter — that the tool’s false-positive rate is non-trivial and disproportionately falls on non-native English writers and students whose prose happens to be cleanly structured.

The disparate-impact dimension is documented quantitatively. The [19] reporting, drawing on controlled studies of essay-grading systems, shows that the bias is measurable and persistent across vendor products. The pattern fits exactly the prediction Benjamin makes in [22]: the discriminatory outcome is produced not by anyone’s malice but by the layered choices of training data, optimization target, and vendor incentive — and is then ratified by the institution’s acceptance of the vendor’s claims.

[23] The Castlereagh Statement gives us direction on AI. Now we ...

[22] The Age of Surveillance Capitalism

[3] Adelphi University accused a student of using AI to ... - Newsday

[4] AI Detection Lawsuits: Every Student Case, Outcome, and What the Data ...

[19] PROOF POINTS: Asian American students lose more points in an AI essay ...

[22] Race After Technology

The procurement frame was not designed to catch this, because the procurement frame treats the vendor’s claims as the institution’s working knowledge. If the vendor says the false-positive rate is low, the contract says the rate is low. If the rate turns out to be high and to fall on protected categories of students, the institution discovers this through litigation, not through procurement. The lesson the lawsuits teach is that the contract is an inadequate epistemic instrument for the kinds of tools being purchased — and that lesson generalizes from detection to every other AI category.

A second diagnostic failure is the access gap reported at [7]. The headline finding — that institutions are licensing AI for faculty and administrative use while not extending the licenses to the students whose work is increasingly being evaluated against AI-augmented baselines — reveals the political economy of the procurement decision. AI is being purchased as a productivity tool for the institution before it is being purchased as a learning tool for the students. The supply-chain costs are borne universally; the access is rationed by hierarchy. This is procurement as enclosure.

The third diagnostic failure is the retention-algorithm category analyzed at [20]. Institutions facing enrollment pressure have turned to predictive analytics that score students for risk of attrition and trigger interventions accordingly. The supply chain for these systems is the same supply chain — the same labeled data, the same offshore moderation labor, the same compute substrate — but the application is one in which the student is now the object rather than the user. The Iberoamerican mapping at [15] shows the same retention-algorithm pattern adopted across institutions in the region, often with even less governance scaffolding than in North America. The student becomes, through the contract, a downstream consumer of an extractive supply chain and an upstream input into a predictive scoring system. The procurement frame catches neither end of the loop.

The Awareness That Is Beginning to Travel

There is, however, evidence that the substrate is becoming visible to constituencies the procurement office did not expect. The faculty survey at [2] is one data point; it is best read alongside the more granular qualitative pictures emerging from the French-language sphere at [9], which documents faculty moving past the initial integrity panic into a more structural critique of who benefits from the tools and on whose terms. The UCLA explainer at [5] is part of a small but growing corpus of public-facing scholarship attempting to demystify the systems

[7] Half of Colleges Don’t Grant Students Access to Gen AI Tools

[20] Risk, Retention, and the Algorithmic Institution: Artificial Intelligence as a Policy Response to Higher Education in Crisis

[15] PDF La llegada de la IA a la educación superior en Iberoamérica: Un mapa ...

[2] 90% Of Faculty Say AI Is Weakening Student Learning: How ... - Forbes

[9] IA générative et transformation des pratiques universitaires

[5] Button-pushing explorers: How to grasp that AI agents can ...

for an institutional audience that has been told for two years to be impressed.

Students, too, are beginning to ask different questions. The detection-lawsuit wave is one form of pushback, but a more interesting one is the discourse around the "AI-native graduate" — the marketing concept analyzed at [21] — which positions the student as the end-consumer of an AI-enriched credential. Students are increasingly suspicious of the framing, partly because they can see that the framing is a procurement justification dressed up as a pedagogical philosophy, and partly because the graduates of the first AI-native cohorts are reporting exactly the cognitive-offloading effects that the Stanford and Mexican studies predicted. The pieces at [17] and [12] document a Spanish-language scholarly community that has moved more quickly than its Anglophone counterpart toward connecting the cognitive effects to the structural conditions.

What is still largely missing — and this is the absence that this essay wants to name — is the connecting move that links the labor question to the pedagogical question through the procurement question. The conversation about cognitive offloading happens in one room; the conversation about Kenyan content moderators happens in a second room; the conversation about enterprise licensing happens in a third room; and the institution has not yet been forced to convene the three rooms in the same building. The work of forcing the convening is the next phase of the discourse, and it will not happen unless faculty and students refuse to keep the conversations separate.

Pioneers, Partnership, and the Missing Frame

A handful of institutions are beginning to model what an alternative procurement posture could look like. The pattern is recognizable from the fair-trade and conflict-minerals movements in earlier decades: a small group of universities adopts a more demanding standard, the standard creates a precedent that vendors find easier to meet than to refuse, and over time the standard migrates from outlier to baseline. The contours are visible in the Castlereagh framework discussed at [23], in the procurement-clause experiments referenced in [8], and in the early consortium work that the OEI mapping at [15] describes among Iberoamerican universities pooling negotiating leverage.

The substantive content of an ethical procurement clause is not mysterious. It would require the vendor to disclose the labor conditions under which training data was annotated, including geography, wage rates, and access to mental-health support for workers exposed

[21] The AI-Native Graduate: Redefining What a University ...

[17] Perez metacognitiva y descarga cognitiva en la era de la IA ...

[12] La inteligencia artificial en la educación: potencial transformador ...

[23] The Castlereagh Statement gives us direction on AI. Now we ...

[8] IA et grandes écoles : quand un algorithme d'admission ...

[15] PDF La llegada de la IA a la educación superior en Iberoamérica: Un mapa ...

to traumatic content. It would require the vendor to disclose the energy and water intensity of the inference being purchased, with verification against an independent standard. It would require the vendor to accept audit rights that go beyond the standard SOC-2 perimeter. It would require the vendor to participate in remediation if downstream harms are documented. None of these requirements are technically impossible; all of them are commercially uncomfortable; and an institution with the leverage of a multi-year contract has the standing to impose them.

What is most missing from the discourse, however, is the partnership frame. The procurement frame casts the relationship between university and vendor as a transactional one — the university buys, the vendor sells, the contract governs. The partnership frame would cast it as a relationship in which the university takes responsibility for what the vendor does because the university’s purchase is what makes the vendor’s behavior viable. This is the move Chomsky has made for decades about the relationship between media institutions and the powers they cover — that the institutional buyer is not neutral with respect to what it buys, and that the apparatus of neutrality is itself the apparatus of complicity. The framework is sketched in [22] for the press; the higher-education version is overdue.

[22] Manufacturing Consent

A partnership frame also reopens the question of what the university could be a partner with that is not a vendor. Open-weights models maintained by academic consortia, training corpora assembled and labeled under transparent labor conditions, compute cooperatives that pool resources across institutions — these are not pipe dreams; they are technical possibilities currently underfunded because the procurement money is going elsewhere. The Iberoamerican consortium work is one early model; the European public-AI initiatives are another. The point is not that universities should build their own models; the point is that the choice between “license from a vendor” and “build from scratch” is a false binary that the vendors have an obvious interest in maintaining.

The Institution as a Moral Actor

To return to the procurement officer signing the contract: the question is not whether that officer is a bad person. The officer is doing the job the institution has defined. The question is whether the institution has defined the job correctly. A procurement function that treats labor conditions and ecological footprint as outside its perimeter is a procurement function inherited from an earlier and simpler commercial

environment, and it is being asked to do work in the AI environment that it was not designed to do. The institution has the standing to redefine the perimeter. It has not yet chosen to.

The choice will not be made by the procurement office on its own initiative, because the procurement office is not authorized to make it. It will be made — if it is made — by the senior leadership of the institution, in response to pressure from faculty, students, and the broader public that the institution depends on for its legitimacy. The evidence assembled across the studies cited here — the cognitive harms in [6], the disparate impacts in [19], the access asymmetries in [7], the retention-algorithm enclosures in [20], the supply-chain extraction documented in [22] — is sufficient to make the case. The question is whether anyone in a position of institutional authority will act on it.

There is a version of the next five years in which universities continue to sign the contracts, continue to host the AI literacy workshops, continue to debate the integrity policy, and continue to behave as if the labor and ecological substrate of the tools is somebody else's problem. In that version, the universities will have used their considerable moral capital — the capital they accumulated by claiming to be different from the corporations whose graduates they produce — to subsidize an extractive industry, and they will have done so under the rubric of pedagogy. The damage to the institution will be slow, because legitimacy erodes slowly, but it will be cumulative, and at some point the gap between what the institution claims and what the institution buys will become impossible to paper over.

There is another version in which the procurement conversation and the pedagogical conversation are forced into the same room, and the institution begins to act on the evidence its own scholars have produced. In that version, the contracts get rewritten. The vendors are asked harder questions, and the ones who refuse to answer lose the business. The partnership frame replaces the transactional frame, and the alternatives — consortia, open weights, public compute — begin to receive the funding that currently flows to the incumbents. The students learn AI literacy that includes the supply chain, not just the prompt. The institution recovers the standing to teach ethics by practicing them.

Which version arrives depends on whether the unseen hand stays unseen. Universities are, among other things, in the business of making invisible things visible. They have the tools, the scholars, and the institutional standing to make the supply chain of their own AI contracts visible to themselves and to the publics they serve. The question this week, this semester, and this procurement cycle is whether

[6] Generative AI Can Harm Learning | SCALE Initiative

[19] PROOF POINTS: Asian American students lose more points in an AI essay ...

[7] Half of Colleges Don't Grant Students Access to Gen AI Tools

[20] Risk, Retention, and the Algorithmic Institution: Artificial Intelligence as a Policy Response to Higher Education in Crisis

[22] The Atlas of AI

they will use what they have.

References

1. 'Think outside the bots': How to stop AI from turning your brain ... - BBC
2. 90% Of Faculty Say AI Is Weakening Student Learning: How ... - Forbes
3. Adelphi University accused a student of using AI to ... - Newsday
4. AI Detection Lawsuits: Every Student Case, Outcome, and What the Data ...
5. Button-pushing explorers: How to grasp that AI agents can ...
6. Generative AI Can Harm Learning | SCALE Initiative
7. Half of Colleges Don't Grant Students Access to Gen AI Tools
8. IA et grandes écoles : quand un algorithme d'admission ...
9. IA générative et transformation des pratiques universitaires
10. Impact de l'IA générative sur la « pensée critique »
11. L'IA générative dans les études supérieures : entre facilitation ...
12. La inteligencia artificial en la educación: potencial transformador ...
13. Natural language processing for social good: Where we are, what is missing, and where we should go
14. PDF hai.stanford.edu
15. PDF La llegada de la IA a la educación superior en Iberoamérica: Un mapa ...
16. PDF Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise
17. Pereza metacognitiva y descarga cognitiva en la era de la IA ...
18. Pereza metacognitiva y descarga cognitiva en la era de la IA generativa ...
19. PROOF POINTS: Asian American students lose more points in an AI essay ...

20. Risk, Retention, and the Algorithmic Institution: Artificial Intelligence as a Policy Response to Higher Education in Crisis
21. The AI-Native Graduate: Redefining What a University ...
22. The Atlas of AI
23. The Castlereagh Statement gives us direction on AI. Now we ...
24. Tutor CoPilot: A Human-AI Approach for Scaling Real-Time ... - ERIC