

Generative AI Adoption in Higher Education: Mapping Institutional Policy Responses and Student Usage Trends from Public University Data and National Surveys

Yasmin Alqahtani^{1,*} and Rabee Alqahtani²

¹ Department of Library and Information Sciences, Faculty of Education, University of Bisha, Bisha 61922, Saudi Arabia.

² University of Wollongong, Australia. Email: rabee3_1983@hotmail.com

* Corresponding author: Yasmin Alqahtani. Email: nsqahtani@ub.edu.sa

HOW TO CITE: | ABSTRACT:

Yasmin Alqahtani, Rabee Alqahtani (2026). Generative AI Adoption in Higher Education: Mapping Institutional Policy Responses and Student Usage Trends from Public University Data and National Surveys International Journal of Special Education, 41(25s), 102 - 118

Generative artificial intelligence (GenAI) has moved from novelty to routine study tool in under four years, yet university governance has followed a slower and more uneven path. This study combines two sources of publicly available evidence to map that gap. First, it codes the official GenAI guidance of a purposive sample of 40 highly ranked universities across six world regions (39 analysable) on eight governance dimensions, with a blind second coding of a random subsample (pooled Cohen's kappa = 0.71). Second, it synthesises repeated national and international surveys of students and staff from the United Kingdom, the United States, Australia and multinational samples published between 2024 and 2026. Central guidance was found at 34 of 39 universities (87%), but only 8 used a formal policy instrument. Half of those with guidance (17 of 34) delegated GenAI rules to individual instructors without an institutional default, 7 permitted use by default with acknowledgement and 6 prohibited it by default. Disclosure of use was required at 23 (68%), and 19 (56%) discouraged or disabled AI text detection. A composite governance index (0 to 8) was higher in North America, Europe and Oceania (mean 6.64) than in Asia, Africa, the Middle East and Latin America (mean 4.07; Mann-Whitney $p = .006$). Among institutions that had published guidance, however, the regional difference was small and not significant (6.64 vs 6.11; $p = .25$), so the gap lies mainly in whether guidance is public rather than in its design. Survey data show UK undergraduate use rising from 66% in 2024 to 95% in 2026, and use in assessed work from 53% to 94%, while only 5% (2024) and 25% (2026) of students in global samples were fully aware of comprehensive institutional guidelines. The findings point to a policy-practice gap: institutions regulate mainly through instructor discretion and

COPYRIGHT STATEMENT:

Copyright: © 2026 Authors.
Open access publication under the
terms and conditions
of the Creative Commons Attribution
(CC BY)

license (<http://creativecommons.org/licenses/by/4.0/>). disclosure, while students have already normalised GenAI in assessed work.

Keywords: generative AI; higher education policy; academic integrity; assessment; student survey; policy mapping

1. Introduction

The public release of ChatGPT in November 2022 forced universities to make governance decisions about a technology that students could adopt faster than committees could meet. Early responses ranged from temporary bans, such as the University of Hong Kong's interim prohibition in February 2023, to statements that GenAI could not realistically be banned at all. Within months, international bodies were urging institutions to move from reaction to regulation. UNESCO's guidance called on governments and universities to validate GenAI tools, protect data privacy and build the capacity of teachers and students to use them responsibly (Miao & Holmes, 2023). National quality agencies followed. In Australia, the Tertiary Education Quality and Standards Agency (TEQSA) asked every registered provider in 2024 to submit an action plan addressing the risks GenAI posed to award integrity (TEQSA, 2024).

By 2026 most large universities had published some form of guidance, and many had licensed GenAI tools for their whole community. Universities have nevertheless faced a hard design problem. Rules must be general enough to apply across disciplines, yet specific enough for a student to know whether a given assignment allows a chatbot. They must protect the validity of awards without depending on AI text detectors, whose accuracy and fairness have been widely questioned (Liang et al., 2023; Weber-Wulff et al., 2023). And they must prepare graduates for workplaces where GenAI is becoming routine.

Student behaviour moved faster still. In the United Kingdom, the share of full-time undergraduates using AI in some form rose from two-thirds in late 2023 to near universality two years later, and use for assessed work followed the same path (Higher Education Policy Institute

[HEPI] & Kortext, 2024, 2025, 2026). Multinational surveys report similar levels across very different higher education systems (Digital Education Council [DEC], 2024, 2026). Institutional governance therefore now operates on a population for which GenAI use is the norm rather than the exception.

Prior studies have mapped university GenAI policies at a single point in time, usually early in the policy cycle or within one national system (Moorhouse et al., 2023; Xiao et al., 2023; McDonald et al., 2025). Separate survey literatures describe student adoption (Ravšelj et al., 2025; Stöhr et al., 2024). Few studies put the two side by side, so it remains unclear whether what institutions publish matches what students do and know. Few also code the specific mechanisms through which universities govern use, such as default rules, disclosure requirements and positions on detection, which are the levers students actually encounter.

This paper addresses these gaps using only publicly available data, which makes the approach transparent and repeatable. It makes three contributions. First, it provides a cross-regional map of GenAI governance mechanisms in 39 leading universities at a mature stage of the policy cycle (late 2026), coded with a reliability check. Second, it assembles repeated survey evidence on student and staff use and awareness from 2023 to 2026. Third, it compares the two to identify where institutional responses and student practice align or diverge. The study asks three research questions:

1. **RQ1.** What governance approaches to GenAI in teaching, learning and assessment do leading universities publish, and how do they vary by region?
2. **RQ2.** How have student and staff use of GenAI, and their awareness of

institutional guidance, changed in national and international surveys between 2023 and 2026?

3. **RQ3.** Where do institutional policy responses and reported student practice align or diverge?

The remainder of the paper reviews prior work and the theoretical lens (Section 2), describes the data and coding (Section 3), reports results (Section 4), discusses them (Section 5) and sets out limitations and conclusions (Sections 6 and 7).

2. Literature review and theoretical framing

2.1 Mapping institutional GenAI policy

The first policy-mapping studies captured the initial shock. Xiao et al. (2023) examined the QS top 500 universities in mid-2023 and found that fewer than a third had any ChatGPT policy, although most of those that did chose to embrace rather than ban the tool. Moorhouse et al. (2023) reviewed the 50 top-ranked universities and found that just under half had public guidance on GenAI in assessment. That guidance was typically framed around academic integrity, advice on assessment design and communication with students.

Later studies documented consolidation. McDonald et al. (2025) analysed guidance at 116 US research-intensive (R1) universities and found that 63% encouraged GenAI use, while 41% gave detailed classroom guidance. Jin et al. (2025) compared institutions across six global regions and identified a shared emphasis on academic integrity, teaching enhancement and equity, with weaker attention to data privacy and equal access. Together these studies show a movement from prohibition towards managed adoption.

Three features of this literature limit what it can tell us. Most studies were conducted in 2023 or 2024, when many institutions had only interim statements. Most code the tone or themes of documents (permissive, restrictive, integrity-focused) rather than the specific rules a student meets. And most stop at the document, without

asking whether the people the rules address know about them.

2.2 Assessment frameworks and the detection debate

A parallel literature offers frameworks that policies can adopt. Chan (2023) proposed an AI policy education framework with pedagogical, governance and operational dimensions, drawing on a survey of students and staff in Hong Kong. Perkins et al. (2024) developed the AI Assessment Scale, which lets educators specify permitted GenAI use for each task on a graded scale from no AI to full AI. In Australia, the national regulator's expert discussion paper argued that assessment should prepare students to engage ethically with AI and should rely on multiple, varied tasks to support trustworthy judgements about learning (Lodge et al., 2023).

More recent work questions whether telling students what they may do is enough. Dawson et al. (2024) argue that assessment design should start from validity, meaning whether a task supports sound judgements about what a student can do, rather than from preventing cheating. Corbin et al. (2025) distinguish discursive changes, such as traffic-light labels that tell students how much AI they may use, from structural changes to the task itself. They argue that discursive rules cannot be enforced in unsupervised work and create only an illusion of control. The University of Sydney's two-lane model, which separates secure supervised assessment from open assessment in which any resource may be used, is a prominent structural response (University of Sydney, 2024).

Detection is the other side of this debate. Weber-Wulff et al. (2023) tested 14 AI-text detectors and found them neither accurate nor reliable, especially once text had been paraphrased. Liang et al. (2023) showed that common detectors disproportionately flagged essays by non-native English writers as machine-generated. These findings matter for policy because a rule that cannot be monitored relies on disclosure, trust or assessment design.

2.3 Student and staff adoption

Large surveys consistently report high and rising use. A global study of 23,218 students in 109 countries found ChatGPT was mainly used for brainstorming, summarising and locating literature, and that students wanted clearer institutional guidance (Ravšelj et al., 2025). A Swedish study of 5,894 students found awareness and use were widespread, with attitudes differing by gender and field of study (Stöhr et al., 2024). An Australian study of 1,135 students in 2023 found that confidence was low but rose with hands-on experience (Kelly et al., 2023). Sector surveys by HEPI, the DEC and Tyton Partners have since tracked use annually, making it possible to observe trends rather than single snapshots.

2.4 Theoretical lens: institutional isomorphism

The study draws on institutional theory to interpret policy patterns. DiMaggio and Powell (1983) argue that organisations in the same field become more alike through three pressures: coercive pressure from regulators and funders, mimetic pressure to copy prestigious peers under uncertainty, and normative pressure from professional networks. GenAI created exactly the conditions under which isomorphism is expected. The technology was new, its consequences were uncertain and the costs of an idiosyncratic response were visible. Isomorphism therefore predicts convergence in policy form among universities exposed to the same pressures, together with differences between systems where those pressures differ, for example where a national regulator requires action. It also predicts that convergence may be stronger in documents than in practice, a possibility the comparison with survey data allows us to examine.

2.5 The gap this study addresses

Policy studies describe what institutions publish, and survey studies describe what students do, but the two are rarely analysed together. The present study codes concrete governance mechanisms in a cross-regional sample at a later stage of policy maturity, tests the reliability of that coding, and

sets the results alongside repeated survey evidence on use and awareness.

3. Method

3.1 Design

The study uses a convergent design built entirely on publicly available data. Strand one is a documentary content analysis of university GenAI guidance (RQ1). Strand two is a secondary synthesis of published survey results (RQ2). The strands are compared at the interpretation stage (RQ3), following the logic of triangulation: agreement between independent sources strengthens a conclusion, and disagreement identifies where policy and practice diverge.

3.2 University sample

Forty universities were selected purposively to give regional breadth among globally prominent, research-intensive institutions: eight in North America, six in the UK and Ireland, six in continental Europe, eight in Asia, six in Oceania and six across Africa, the Middle East and Latin America. Institutions were drawn from the upper tiers of international rankings and from each region's leading national universities. Leading institutions were chosen deliberately, because they are the most likely to have published guidance and the most likely to be copied by others, which makes them a useful test of how far governance has developed.

KU Leuven was excluded from analysis because its official guidance pages could not be accessed during data collection, leaving an analytic sample of 39. For analysis, the UK, Ireland and continental Europe were combined (Europe, $n = 11$), as were Africa, the Middle East and Latin America ($n = 6$), because of small cell sizes. For the regional comparison, institutions were grouped into North America, Europe and Oceania ($n = 25$) and Asia, Africa, the Middle East and Latin America ($n = 14$). The full sample is listed in the Appendix.

3.3 Data collection

For each university, the current central guidance on GenAI in teaching, learning and assessment was located through web search and read on official university domains in September and October 2026. Eligible pages included those hosted by provosts, senates, teaching and learning centres, academic integrity offices and IT services, as well as official university news releases where these announced policy. Third-party summaries were not used for coding. Non-English documents (for example, from Tsinghua University, Sorbonne University and Tecnológico

de Monterrey) were translated during reading. Where an institution's guidance was spread across several pages, all relevant central pages were read together. Faculty-level documents were recorded but did not count as central guidance.

3.4 Coding framework

Each institution was coded on the dimensions in Table 1, which were derived from the mechanisms discussed in prior work (Section 2): who sets the rules, what applies when no rule is set, how use is made visible, and how it is monitored. Where a page did not address a dimension, it was coded as not specified rather than inferred.

Table 1. Policy coding framework.

Dimension	Coding categories
Central guidance	University-wide guidance; faculty-level only; none found
Instrument type	Formal policy or regulation; guidance or principles; both
Default student stance	Prohibited by default; instructor discretion with no default; permitted by default with acknowledgement; not specified
Instructors must state rules	Required or explicitly expected; encouraged; not specified
Student disclosure	Required; encouraged; not specified
AI detection	Discouraged or disabled; used; not specified
Staff guidance	Present; absent
Assessment redesign	Explicitly recommended; not recommended
Enterprise GenAI tool	Institution-provided tool identified; none found

The date on which each document was first issued and last updated was also recorded where pages carried dates.

3.5 Coding reliability

A second, independent coding pass was made for a random subsample of 10 of the 34 institutions with central guidance (29%), drawn with a fixed seed. The second coder was blind to the original codes and worked from the official pages. The check covered the four dimensions requiring most interpretive judgement: default stance, instructor obligation, disclosure and detection. Agreement was 75% (30 of 40 decisions), with a pooled

Cohen's kappa of 0.71, indicating substantial agreement. Kappa by dimension was 0.64 for detection, 0.56 for default stance, 0.52 for disclosure and 0.25 for instructor obligation.

Most disagreements arose where documents used hedged language ("should", "may", "expected") or where the relevant rule sat on a page other than the main guidance hub. Disagreements were resolved by re-reading the sources together. This changed eight codes at seven institutions, and the reconciled codes are used throughout. The weaker agreement on instructor obligation reflects a genuine ambiguity in policy language, discussed in Section 5.

3.6 Governance index and analysis

To summarise maturity, a governance index (0 to 8) awarded one point each for: central guidance; a stated position on student use (a default or explicit delegation to instructors); a requirement that instructors state rules; required disclosure; discouraging or disabling detection; staff guidance; recommended assessment redesign; and an institution-provided tool. The index treats discouraging detection as the more mature position, consistent with evidence on detector unreliability (Weber-Wulff et al., 2023). It is descriptive and weights each element equally.

Categorical codes are reported as counts and percentages. Regional differences in the index were tested with a two-sided Mann-Whitney U test, with the rank-biserial correlation as the effect

size. The presence of central guidance was compared with Fisher's exact test. Three sensitivity checks were run: excluding the tool element, excluding the detection element, and restricting the comparison to institutions with central guidance. Analyses used Python (pandas and SciPy).

3.7 Survey data

Survey figures were taken only from the publishers' own reports or official pages (Table 2). Figures were extracted as reported. Because question wording and samples differ across series, trends are interpreted within a series, and cross-series comparisons are treated as indicative only. Where a report gave different figures for the same measure in text and in a chart, the text figure was used.

Table 2. Survey sources.

Series	Publisher	Scope	Editions used	Sample size
Student Generative AI Survey	HEPI and Kortext	UK full-time undergraduates	2024, 2025, 2026	1,250; 1,041; 1,054
Global AI Student and Faculty Surveys; AI in Higher Education Global Survey	Digital Education Council	16 countries (2024); 28 countries (2025 faculty); 35 countries (2026)	2024, 2025, 2026	more than 3,800 students; 1,681 faculty; 27,284 students and 18,114 staff
Time for Class	Tyton Partners	US students, instructors and administrators	2024, 2025, 2026	2025: 1,528 students, 1,572 instructors, 312 administrators
Student Digital Experience Insights	Jisc	UK higher education students	2024/25	15,398
Global Student Survey	Chegg.org	Undergraduates aged 18 to 21, 15 countries	2025	11,706
Student Voices on AI in Higher Education	AI in Higher Education project	Students at four Australian universities	2024	8,028

3.8 Ethics

The study analysed publicly available institutional documents and published aggregate survey results. It involved no human participants and no identifiable personal data, so ethics approval was not required.

4. Results

4.1 Prevalence and form of institutional guidance (RQ1)

Central guidance was found at 34 of the 39 analysed universities (87%). The University of São Paulo had faculty-level guidance only, and no

public university-wide guidance was located for Peking University, KAIST, IIT Bombay or Khalifa University. Most central documents were advisory. Of the 34, 26 (76%) were framed only as guidance or principles, while 8 (24%) included a formal policy or regulation. Formal instruments were found at McGill, Oxford, Trinity College Dublin, the University of Amsterdam, the National University of Singapore, the University of Hong Kong, the University of Sydney and the University of the Witwatersrand. Wits offers a recent example: its Council approved a university AI policy in June 2026, building on teaching-centre guidelines first issued in December 2024.

Guidance was strongly oriented towards teaching staff (Figure 1). Every institution with central guidance offered advice to staff on course and assessment design, and 27 of 34 (79%) explicitly recommended redesigning assessment, for example through secure in-person tasks, oral components or process-based assessment. Most placed a duty on instructors: 25 (74%) required or explicitly expected instructors to state GenAI rules for their course or tasks, and the remaining 9 encouraged it.


For students, the most common integrity mechanism was disclosure. Of the 34 institutions,

23 (68%) required students to acknowledge GenAI use, 7 encouraged it and 4 did not address it. Positions on detection were more divided. Nineteen institutions (56%) discouraged reliance on AI text detectors or had switched them off, often citing accuracy and fairness concerns. UCL, for example, stated that it does not use detectors, and the University of Cape Town discontinued Turnitin's AI score from October 2025. Four institutions, Stanford and three Australian universities (Melbourne, Sydney and UNSW), described detector output as usable, typically as one input to an investigation rather than as proof. The remaining 11 did not state a position.

Institution-provided GenAI tools were identified at 29 of the 34 institutions with guidance (85%), and at 30 of all 39 (77%). They ranged from licensed Microsoft Copilot, Google Gemini and ChatGPT Edu to homegrown platforms such as Cogniti at Sydney, U-M GPT at Michigan, ELM at Edinburgh, dAISy at Imperial and TECgpt at Monterrey. Peking University provided its own study assistants without publishing central assessment rules, showing that provision and rule-making can develop separately.

Guidance is near universal, binding rules are not

Share of the 34 universities with central GenAI guidance that include each element



Blue: rules set at institutional level rather than left to each course

Coded from official university web pages, September to October 2026 · n = 34

Figure 1. Governance elements in central GenAI guidance (n = 34).

4.2 Default rules for students

The default rule for students, meaning what applies when an instructor has said nothing, was

the clearest point of divergence (Table 3). Half of the institutions with guidance (17 of 34) left the decision to each instructor or course without an institutional fallback. Seven (21%) made GenAI

use permissible by default provided it was acknowledged: Cambridge, UCL, Trinity College Dublin, the National University of Singapore, Nanyang Technological University, Sydney (for open assessments) and Auckland. Six (18%)

treated use as prohibited unless explicitly permitted: Stanford, Toronto, Oxford, Imperial, Amsterdam and Cape Town. Four did not address the question.

Table 3. Default student stance by region, universities with central guidance (n = 34).

Region	Instructor discretion	Permitted with acknowledgement	Prohibited by default	Not specified	Total
Europe	4	3	3	1	11
North America	5	0	2	1	8
Oceania	4	2	0	0	6
Asia	3	2	0	0	5
Africa, Middle East and Latin America	1	0	1	2	4
All	17	7	6	4	34

The categories hide variation in how defaults are built. Several permissive institutions attach the default to a type of task. NUS allows acknowledged use in unsupervised take-home work, and Sydney and Auckland apply their open lane only to assessments that are not secured. Several delegating institutions supply a menu of standard positions for instructors to choose from: UCL uses three categories, UNSW six and Monash four, while Arizona State offers three model syllabus statements. Even where no fallback is stated, these menus narrow the range of rules a student is likely to meet.

4.3 Regional variation

Governance maturity differed sharply by region (Figure 2). The mean governance index was 6.83 in Oceania, 6.73 in Europe and 6.38 in North America, compared with 4.00 in Asia and 4.17


across Africa, the Middle East and Latin America. Taken together, North America, Europe and Oceania scored 6.64 against 4.07 for the other regions (Mann-Whitney $U = 266$, $p = .006$; rank-biserial $r = 0.52$). All 25 institutions in the first group had central guidance, against 9 of 14 (64%) in the second (Fisher's exact $p = .003$).

The difference survived the removal of the tool element ($p = .016$) and the detection element ($p = .002$). However, when the comparison was restricted to institutions that had published central guidance, the gap narrowed to 6.64 against 6.11 and was not statistically significant ($U = 141$, $p = .25$). Regional differences therefore arise mainly from whether public central guidance exists, not from how institutions that publish guidance design it. Variation within the second group was wide: the University of Tokyo, the University of Cape Town and Tecnológico de Monterrey each scored 7, while five institutions scored

0 or 1.

North America, Europe and Oceania score higher on GenAI governance

Governance index (0 to 8): one dot per university; blue line marks the regional mean



Coded from official university web pages, September to October 2026 · n = 39

Figure 2. Governance index by region (n = 39).

4.4 Timing and revision

Guidance has been revised repeatedly rather than settled once. Where pages carried dates, 13 of the 34 institutions documented a first version in 2023, and at least 22 showed a version issued or substantially updated in 2025 or 2026. Revisions often changed the substance of the rules. Sydney moved in 2025 from allowing AI unless a coordinator opted out to its two-lane model; Oxford adopted a formal policy on AI in summative assessment in 2025; Tsinghua issued guiding principles in November 2025; Trinity College Dublin published version 2 of its College statement in June 2026; and Auckland announced lane labels for all courses from 2027. Three institutions (Harvard, Michigan and the Technical University of Munich) published undated guidance, which makes it harder for students to judge whether rules are current.

4.5 Student and staff usage trends (RQ2)

Student use rose steeply in every repeated series (Figure 3). Among UK undergraduates, any AI use rose from 66% (2024) to 92% (2025) and 95%

(2026), and use for assessed work from 53% to 88% and then 94% (HEPI & Kortext, 2024, 2025, 2026). In the United States, frequent (daily or weekly) student use rose from 14% in spring 2023 to 36% in 2024 and 42% in 2025, while frequent instructor use rose from 4% to 30% over the same period (Tyton Partners, 2025). By 2026, 32% of US students and 25% of instructors reported daily use (Tyton Partners, 2026).


Multinational surveys report similar levels. The DEC found that 86% of students across 16 countries regularly used AI in their studies in 2024, with 54% using it weekly (DEC, 2024). Its 2026 survey of 27,284 students in 35 countries found 88% used AI in their learning, though the figure was lower in the United States and Canada (65%), and 77% of 18,114 staff used AI in their teaching (DEC, 2026). Chegg's 2025 survey of 11,706 undergraduates in 15 countries found 80% had used GenAI to support their studies (Chegg.org, 2025). In Australia, 83% of 8,028 students at four universities had used AI for study and 44% did so at least weekly (AI in Higher Education, 2024).

Not every series reports such high levels. Jisc's 2024/25 survey of 15,398 UK higher education students found only 34% said they used AI tools as part of their learning (Jisc, 2025), less than half the HEPI figure from the same year. The

difference probably reflects question wording and the survey context, since the Jisc item sits within a broad digital experience survey, and it shows why trends should be read within a series.

Student GenAI use has far outpaced awareness of guidance

Per cent of students, by survey year



Series differ in wording and sample; compare trends within a series. Survey year = edition year.

HEPI & Kortext (2024, 2025, 2026); Tyton Partners (2025); Digital Education Council (2024, 2026)

Figure 3. Student GenAI use and awareness of institutional guidance, 2023 to 2026.

4.6 Awareness, support and institutional adoption

Awareness and support lagged behind use. In 2024 only 5% of students globally said they were fully aware of their institution's AI guidelines and found them comprehensive, rising to 25% in 2026; the 2026 figure for staff was 21% (DEC, 2024, 2026). Faculty reported limited institutional preparation: in 2025, only 6% fully agreed that their institution had provided sufficient resources to develop their AI literacy (DEC, 2025). Only 23% of students in 2026 felt adequately prepared to use AI professionally (DEC, 2026).

UK series show improvement from a low base. The share of students agreeing their institution's policy was clear rose from 63% in 2024 to 80% in 2025, and the share saying staff were well equipped to support them rose from 18% to 42%. Even so, only 36% had received AI skills support

from their institution in 2025 (HEPI & Kortext, 2024, 2025). In 2026, 48% said teaching staff were helping them develop AI skills for their careers, 37% said their institution encouraged AI use and 38% said it provided them with AI tools (HEPI & Kortext, 2026). Jisc found that 23% of UK students had been given access to AI platforms by their institution (Jisc, 2025), and Chegg found that 69% of students internationally believed their university should provide GenAI tools (Chegg.org, 2025).

Australian data show the tension between rules and behaviour most directly. Only 32% of students felt they had received enough guidance to use AI effectively, 91% worried about breaking university rules, and 40% reported having used AI in assessment when they were not supposed to (AI in Higher Education, 2024).

At the institutional level, US administrators reporting a formal GenAI policy rose from 3% in 2023 to 24% in 2024, 28% in 2025 and 32% in 2026 (Tyton Partners, 2024, 2025, 2026). These figures cover a broad cross-section of US colleges and are far below the 100% observed among the eight North American institutions in this sample, underlining that the sample represents the leading edge of the sector. In Australia, the national regulator required every provider to submit a GenAI action plan in 2024 (TEQSA, 2024).

5. Discussion (RQ3)

5.1 From prohibition to delegated governance

The sample confirms the shift from early bans towards managed adoption reported by Xiao et al. (2023) and McDonald et al. (2025). By late 2026 only six of 34 institutions with guidance kept a prohibition default, and even these allowed instructors to permit specified uses. The dominant model, however, is neither institutional permission nor prohibition but delegation: half of the universities leave the rule to each course, and three-quarters place an explicit duty on instructors to state it.

Delegation has clear advantages. It respects disciplinary differences, since appropriate use in a programming course differs from appropriate use in a creative writing seminar. Frameworks such as the AI Assessment Scale (Perkins et al., 2024) and institutional category menus give instructors a shared vocabulary for setting rules. But delegation also moves the burden of policy-making to individual teaching staff, many of whom report limited preparation. Only 6% of faculty in a global 2025 survey fully agreed that their institution had resourced their AI literacy (DEC, 2025). The weak intercoder agreement on instructor obligation ($\kappa = 0.25$) is itself revealing. Institutional documents often say instructors "should" or are "expected to" state rules without saying what happens if they do not, so the obligation is unclear even to a careful reader.

The pattern fits the isomorphism lens. The near-universal presence of staff guidance, assessment

redesign advice and disclosure rules among leading institutions suggests strong mimetic and normative pressure. Universities copied recognisable models, such as Sydney's lanes or UCL's categories, and drew on shared professional networks and frameworks. Where a regulator acted, as in Australia, coercive pressure produced especially uniform results: every Oceania institution in the sample had guidance, required instructors to state rules and provided an institutional tool.

5.2 The policy-practice gap

Comparing the two strands reveals three areas of misalignment.

First, **use outpaces awareness**. Near-universal student use sits beside low awareness of comprehensive guidance (25% globally in 2026). When rules vary by course, students must work out the regime for each task, and a student taking four courses may meet four different rules in a single semester. This is consistent with Australian findings that most students worry about breaking rules while two in five admit doing so. Delegated governance appears to produce documents that are clear to their authors but much less clear to the students they govern.

Second, **disclosure is the main integrity lever, but detection is being abandoned**. Two-thirds of institutions require students to acknowledge GenAI use, and more than half explicitly discourage detectors, consistent with evidence that detectors are unreliable and biased (Liang et al., 2023; Weber-Wulff et al., 2023). Disclosure without reliable detection depends on trust and on assessment design. Corbin et al. (2025) argue that rules of this kind are discursive and cannot be enforced in unsupervised work. The sample suggests institutions partly recognise this: four in five recommend redesigning assessment, and structural approaches that separate secured from open assessment are spreading from Sydney to Auckland and beyond. Yet redesign is usually recommended rather than required, and the default for most unsupervised work remains an instructor's statement.

Third, **provision lags permission**. Although 77% of sampled universities offer an institutional GenAI tool, only 38% of UK students in 2026 and 23% in the Jisc 2024/25 survey reported being given access by their institution. Part of this gap reflects the sample, since leading institutions are likely to be ahead of the sector. Part may also reflect low awareness of tools that exist. Either way, where provision is not universal or not known, students who pay for premium tools may gain an advantage. This equity concern was identified by Jin et al. (2025) and by UNESCO (Miao & Holmes, 2023) as underdeveloped in policy.

5.3 Regional asymmetry

The regional comparison yields a more precise conclusion than a simple league table. The governance index was much lower in Asia, Africa, the Middle East and Latin America overall, but among institutions that had published central guidance, the regions scored similarly. The University of Tokyo, Cape Town and Tecnológico de Monterrey all designed guidance comparable to the strongest in Europe and North America. The regional difference therefore lies mainly in whether central guidance is public, not in its design.

This matters for interpretation. Low scores at Peking University, KAIST, IIT Bombay and Khalifa University do not show that these institutions are permissive or prohibitive. They may govern GenAI through faculty rules, internal systems or national frameworks not visible on public pages, and some invest heavily in their own AI platforms. Public transparency is nonetheless a governance quality in its own right: students, prospective students and partner institutions cannot follow rules they cannot see.

5.4 Implications

For institutions, the results support five practical steps:

1. Set an explicit institutional default for when an instructor is silent, ideally tied to assessment type (secured or open), so that silence is not left to interpretation.

2. Replace hedged wording about instructor duties with a clear requirement and a short menu of standard positions.
3. Pair disclosure rules with structural assessment redesign, rather than relying on statements that cannot be enforced.
4. Make approved tools and their availability visible to students, so that provision closes rather than widens access gaps.
5. Date and version all guidance, and publish it on a single stable page.

For regulators, the Australian experience suggests that a light requirement to plan can lift baseline governance across a system. For researchers, the governance index offers a repeatable instrument for tracking change in public policy texts over time.

6. Limitations and future research

Six limitations apply. First, the sample is purposive and weighted towards highly ranked, well-resourced universities, so prevalence figures should not be generalised to whole systems; the Tyton administrator data suggest that the wider sector is less advanced. Second, coding relied on publicly accessible pages. Guidance behind staff or student logins (for example, the full HKU policy) and faculty-level rules were not captured, so the absence of public guidance does not prove the absence of governance. Third, policies change quickly, and the coding is a snapshot of September and October 2026. Fourth, only 10 of the 34 institutions with guidance were double-coded, and agreement on instructor obligation was modest, so codes for the remaining institutions carry more uncertainty. Fifth, the governance index weights its elements equally and treats discouraging detection as more mature, which is a normative choice, although the regional result held when individual elements were removed. Sixth, the survey series use different questions, populations and sampling methods, and some measures changed between editions, so cross-series comparisons are indicative.

Future research could repeat the coding annually to build a panel, extend it to a stratified sample

that includes teaching-focused institutions, and double-code the full sample. Linking institutional documents to student surveys run within the same universities would allow a direct test of whether clearer defaults raise student awareness and reduce unauthorised use. Studies of how instructors actually choose among delegated options would also help explain the gap between policy text and student experience.

7. Conclusion

Nearly four years after ChatGPT's release, leading universities have largely settled on a governance model of delegated rules, mandatory disclosure, scepticism towards detectors and recommended assessment redesign. Students, meanwhile, have made GenAI a routine part of their study and

assessment. The central weakness is not the absence of policy but its legibility: rules set course by course remain unclear to most students and staff, and access to approved tools trails permission to use them. Regional differences in governance lie mainly in whether guidance is published at all. Making rules clear, visible and enforceable through assessment design, rather than simply more numerous, is the next task for institutional policy.

Acknowledgements

The authors are thankful to the Deanship of Graduate Studies and Scientific Research at University of Bisha for supporting this work through the Fast-Track Research Support Program.

References

- AI in Higher Education. (2024). Student voices on AI in higher education: Survey insights. The University of Queensland, Monash University, Deakin University and University of Technology Sydney. https://aiinhe.org/wp-content/uploads/2024/10/aiinhe_surveyinsights.pdf
- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, Article 38. <https://doi.org/10.1186/s41239-023-00408-3>
- Chegg.org. (2025). Global student survey 2025. <https://www.chegg.org/global-student-survey-2025>
- Corbin, T., Dawson, P., & Liu, D. (2025). Talk is cheap: Why structural assessment changes are needed for a time of GenAI. *Assessment & Evaluation in Higher Education*, 50(7), 1087-1097. <https://doi.org/10.1080/02602938.2025.2503964>
- Dawson, P., Bearman, M., Dollinger, M., & Boud, D. (2024). Validity matters more than cheating. *Assessment & Evaluation in Higher Education*, 49(7), 1005-1016. <https://doi.org/10.1080/02602938.2024.2386662>
- Digital Education Council. (2024). What students want: Key results from DEC global AI student survey 2024. <https://www.digitaleducationcouncil.com/post/what-students-want-key-results-from-dec-global-ai-student-survey-2024>
- Digital Education Council. (2025). What faculty want: Key results from the global AI faculty survey 2025. <https://www.digitaleducationcouncil.com/post/what-faculty-want-key-results-from-the-global-ai-faculty-survey-2025>
- Digital Education Council. (2026). AI in higher education global survey 2026. https://actualidad.udla.cl/wp-content/uploads/sites/15/2026/06/6a3b9fb2f4e0472d0160d771_Digital-Education-Council-AI-in-Higher-Education-Global-Survey.pdf
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160. <https://doi.org/10.2307/2095101>
- Higher Education Policy Institute & Kortext. (2024). Provide or punish? Students' views on generative AI in higher education (HEPI Policy Note 51). <https://www.hepi.ac.uk/wp-content/uploads/2024/01/HEPI-Policy-Note-51.pdf>

-
- Higher Education Policy Institute & Kortext. (2025). Student generative AI survey 2025. <https://www.hepi.ac.uk/wp-content/uploads/2025/02/HEPI-Kortext-Student-Generative-AI-Survey-2025.pdf>
- Higher Education Policy Institute & Kortext. (2026). Student generative AI survey 2026 (HEPI Report 199). <https://www.hepi.ac.uk/wp-content/uploads/2026/03/HEPI-Report-199-Gen-AI-Survey-2026.pdf>
- Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence*, 8, Article 100348. <https://doi.org/10.1016/j.caeai.2024.100348>
- Jisc. (2025). 2024/25 UK higher education students' digital experience insights survey findings. <https://digitalinsights.jisc.ac.uk/reports-and-briefings/our-reports/2024-25-uk-higher-education-students-digital-experience-insights-survey-findings/>
- Kelly, A., Sullivan, M., & Strampel, K. (2023). Generative artificial intelligence: University student awareness, experience, and confidence in use across disciplines. *Journal of University Teaching and Learning Practice*, 20(6), Article 12. <https://doi.org/10.53761/1.20.6.12>
- Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns*, 4(7), Article 100779. <https://doi.org/10.1016/j.patter.2023.100779>
- Lodge, J. M., Howard, S., Bearman, M., Dawson, P., & Associates. (2023). Assessment reform for the age of artificial intelligence. Tertiary Education Quality and Standards Agency. <https://www.teqsa.gov.au/sites/default/files/2023-09/assessment-reform-age-artificial-intelligence-discussion-paper.pdf>
- McDonald, N., Johri, A., Ali, A., & Hingle, A. (2025). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. *Computers in Human Behavior: Artificial Humans*, 3, Article 100121. <https://doi.org/10.1016/j.chbah.2025.100121>
- Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. <https://doi.org/10.54675/EWZM9535>
- Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open*, 5, Article 100151. <https://doi.org/10.1016/j.caeo.2023.100151>
- Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6). <https://doi.org/10.53761/q3azde36>
- Ravšelj, D., Keržič, D., Tomažević, N., Umek, L., Brezovar, N., Iahad, N. A., Abdulla, A. A., Akopyan, A., Aldana Segura, M. W., AlHumaid, J., Allam, M. F., Alló, M., Andoh, R. P. K., Andronic, O., Arthur, Y. D., Aydin, F., Badran, A., Balbontín-Alvarado, R., Ben Saad, H., . . . Aristovnik, A. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLOS ONE*, 20(2), e0315011. <https://doi.org/10.1371/journal.pone.0315011>
- Stöhr, C., Ou, A. W., & Malmström, H. (2024). Perceptions and usage of AI chatbots among students in higher education across genders, academic levels and fields of study. *Computers and Education: Artificial Intelligence*, 7, Article 100259. <https://doi.org/10.1016/j.caeai.2024.100259>
- Tertiary Education Quality and Standards Agency. (2024). Gen AI strategies for Australian higher education: Emerging practice. <https://www.teqsa.gov.au/about-us/news-and-events/latest-news/gen-ai-strategies-australian-higher-education-emerging-practice>
- Tyton Partners. (2024). Time for class 2024: Racing forward, bridging the gap between generative AI proficiency and educational practice. <https://tytonpartners.com/racing-forward-bridging-the-gap-between-generative-ai-proficiency-and-educational-practice/>
- Tyton Partners. (2025). Time for class 2025. https://4213961.fs1.hubspotusercontent-na1.net/hubfs/4213961/Publications/Time%20for%20Class/Tyton%20Partners_Time%20for%20Class%202025.pdf
- Tyton Partners. (2026). Time for class 2026. <https://tytonpartners.com/time-for-class-2026/>
-

University of Sydney. (2024, November 27). University of Sydney's AI assessment policy: Protecting integrity and empowering students [News release]. <https://www.sydney.edu.au/news-opinion/news/2024/11/27/university-of-sydney-ai-assessment-policy.html>

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19, Article 26. <https://doi.org/10.1007/s40979-023-00146-z>

Xiao, P., Chen, Y., & Bao, W. (2023). Waiting, banning, and embracing: An empirical analysis of adapting policies for generative AI in higher education [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2305.18617>

Appendix: Coded university sample

Each university name links to the main official page coded. Rows are grouped by region (North America, Europe, Oceania, Asia, then Africa, the Middle East and Latin America) and sorted by governance index. KU Leuven was sampled but excluded because its guidance pages could not be accessed.

Table A1. Coded university sample.

University	Country	Guidance	Default stance	Disclosure	Detection	Tool	Index (0-8)
McGill University	Canada	Central	Instructor discretion	Required	Discouraged	Yes	8
Stanford University	United States	Central	Prohibited by default	Required	Used	Yes	7
University of California, Berkeley	United States	Central	Instructor discretion	Encouraged	Discouraged	Yes	7
Arizona State University	United States	Central	Instructor discretion	Required	Not specified	Yes	6
Harvard University	United States	Central	Instructor discretion	Not specified	Not specified	Yes	6
Massachusetts Institute of Technology	United States	Central	Instructor discretion	Not specified	Not specified	Yes	6
University of Toronto	Canada	Central	Prohibited by default	Encouraged	Discouraged	Yes	6
University of Michigan	United States	Central	Not specified	Not specified	Discouraged	Yes	5
Imperial College London	United Kingdom	Central	Prohibited by default	Required	Discouraged	Yes	8
University College London	United Kingdom	Central	Permitted with acknowledgement	Required	Discouraged	Yes	8
University of Amsterdam	Netherlands	Central	Prohibited by default	Required	Discouraged	Yes	8
University of Oxford	United Kingdom	Central	Prohibited by default	Required	Discouraged	Yes	8
ETH Zurich	Switzerland	Central	Instructor discretion	Required	Discouraged	Yes	7
Trinity College Dublin	Ireland	Central	Permitted with acknowledgement	Required	Discouraged	Yes	7

University	Country	Guidance	Default stance	Disclosure	Detection	Tool	Index (0-8)
Technical University of Munich	Germany	Central	Instructor discretion	Not specified	Not specified	Yes	6
University of Cambridge	United Kingdom	Central	Permitted with acknowledgement	Required	Discouraged	None found	6
University of Edinburgh	United Kingdom	Central	Instructor discretion	Required	Discouraged	Yes	6
Delft University of Technology	Netherlands	Central	Not specified	Encouraged	Discouraged	Yes	5
Sorbonne University	France	Central	Instructor discretion	Required	Not specified	None found	5
Australian National University	Australia	Central	Instructor discretion	Required	Discouraged	Yes	8
The University of Melbourne	Australia	Central	Instructor discretion	Required	Used	Yes	7
The University of Sydney	Australia	Central	Permitted with acknowledgement	Required	Used	Yes	7
UNSW Sydney	Australia	Central	Instructor discretion	Required	Used	Yes	7
Monash University	Australia	Central	Instructor discretion	Required	Not specified	Yes	6
University of Auckland	New Zealand	Central	Permitted with acknowledgement	Encouraged	Not specified	Yes	6
The University of Tokyo	Japan	Central	Instructor discretion	Encouraged	Discouraged	Yes	7
Nanyang Technological University	Singapore	Central	Permitted with acknowledgement	Required	Not specified	Yes	6
National University of Singapore	Singapore	Central	Permitted with acknowledgement	Required	Discouraged	None found	6
The University of Hong Kong	Hong Kong SAR	Central	Instructor discretion	Encouraged	Not specified	Yes	6
Tsinghua University	China	Central	Instructor discretion	Required	Not specified	Yes	6
Peking University	China	None found	Not specified	Not specified	Not specified	Yes	1
Indian Institute of Technology Bombay	India	None found	Not specified	Not specified	Not specified	None found	0
KAIST	South Korea	None found	Not specified	Not specified	Not specified	None found	0
Tecnológico de Monterrey	Mexico	Central	Instructor discretion	Required	Discouraged	Yes	7

University	Country	Guidance	Default stance	Disclosure	Detection	Tool	Index (0-8)
University of Cape Town	South Africa	Central	Prohibited by default	Encouraged	Discouraged	Yes	7
University of the Witwatersrand	South Africa	Central	Not specified	Required	Discouraged	None found	6
American University of Beirut	Lebanon	Central	Not specified	Required	Not specified	None found	4
Universidade de São Paulo	Brazil	Faculty only	Not specified	Encouraged	Not specified	None found	1
Khalifa University	United Arab Emirates	None found	Not specified	Not specified	Not specified	None found	0
