

# Delegation Is Not Assistance: Agentic AI and Assessment Governance in Higher Education - A Comparative Documentary Analysis of 12 Universities

Wong, Kaiyan<sup>1\*</sup> & Sun, Bojian<sup>1</sup>

<sup>1</sup>SEGi University, Malaysia

\*Corresponding author email: [wongkaiyan@segi.edu.my](mailto:wongkaiyan@segi.edu.my)

Received 05 September 2026, Revised 06 October 2026, Accepted 09 October 2026, Available online 11 October 2026

To link to this article: <https://doi.org/10.53797/ujssh.v5i2.19.2026>

**Abstract:** Generative AI (GenAI) assessment policies in higher education have centred on originality, disclosure, and permission to use AI. Agentic AI, which can plan, use tools, access platforms, and act with limited human intervention, challenges this logic because it can take over the whole action chain of an assessed task rather than only generate content. This study examines how assessment governance changes from GenAI to agentic AI. Using exploratory comparative documentary analysis, public central guidance from 12 universities was theoretically sampled: six agentic-aware cases and six GenAI-dominant comparators, with a retrieval cut-off of 3 October 2026. Documents were coded on six dimensions using a directed framework informed by Bacchi's "What's the Problem Represented to Be?" approach. All agentic-aware cases distinguished agentic AI from GenAI, and most set explicit delegation and access boundaries, whereas no comparator explicitly addressed agents operating institutional systems on a student's behalf. The comparators, however, were strong on evidence of learning and task-level transparency, and programme-level assurance was sparse in both groups. The study proposes the Delegation-Evidence-Access (DEA) framework, which moves governance from tool permission towards responsibility allocation, and derives four assessment contexts and a task-, course-, and programme-level implementation model. Implications for Malaysian higher education are discussed. Findings are descriptive, based on selected public documents, and not representative; no inter-rater reliability was assessed.

**Keywords:** agentic AI, generative AI, higher education, assessment, governance

## 1. Introduction

Since late 2022, large-language-model-based generative artificial intelligence (GenAI) has rapidly changed how university students read, write, programme, research, and complete assessment. Early institutional governance accordingly centred on a comparatively clear question: if students use tools such as ChatGPT to generate text, code, or images, to what extent can the submitted product still be treated as their own work?

A substantial research base has emerged around this problem. Luo (2024), analysing policies from 20 leading universities, found that GenAI risk was mainly represented as the possibility that students would no longer submit original work. GenAI was positioned as external assistance, outside independent intellectual labour. Wang, Dang, Wu, and Mac (2024) showed that US universities were moving from a narrow academic-integrity response towards teaching support, ethics, privacy, and differentiated-use guidance. An, Yu, and James (2025) found that 94% of faculty-facing guidance among the top 50 US universities emphasised establishing and clearly communicating course-level GenAI policies. In Asia, Dai, Lai, Lim, and Liu (2025) identified text generation, student management, and academic integrity as recurring policy priorities.

This governance paradigm, however, contains an implicit technical assumption: AI mainly intervenes by producing content, while the student remains the principal actor who decides when to invoke the system, what to prompt, how to evaluate the response, and what to submit. Agentic AI weakens that assumption.

Educational research published in 2026 describes agentic AI as systems that can pursue goals over time, plan multiple steps, use external tools, maintain state or memory, and act with less continuous human intervention. In a

systematic review of 33 empirical studies, Liu, Tang, Wu, Chen, and Wang (2026) found that GenAI agents were already used for cognitive support, self-regulation, affective support, and human-AI interaction. They identified trust, governance, and contextual fit as major issues for the next phase of adoption. Baradziej (2026) likewise concluded that the higher-education evidence base on agentic AI was expanding quickly but remained at an early empirical stage.

This technological change has already entered university policy. Queen's University (2026) noted that agentic AI can navigate websites, log into systems, execute multi-step tasks, and make decisions on a user's behalf; some tools can interact with learning-management systems and potentially submit coursework. The University of Kansas (2026) warned that AI agents may log into learning systems and complete academic work autonomously. This makes it harder to establish whether submitted outputs reflect a student's own knowledge and effort.

A more demanding governance question follows. When AI can plan, operate, execute, and submit as well as generate an answer, do policy frameworks organised mainly around originality, citation, disclosure, and permission to use AI still have sufficient conceptual resolution?

Research has begun to address the educational roles of agentic AI, learner agency, and self-regulation. Salah, Al-Busaidi, Al-Ghawi, and Shawky (2026) propose bounded autonomy. They argue that when AI assumes planning, monitoring, and judgement, education must distinguish support for self-regulation from replacement of the regulatory capacities students are meant to develop. Pauzi, Vizgirda, Duran, and Mavrikis (2026) conceptualise student agency in AI-mediated learning as a property shaped by interfaces, decision rights, and the allocation of action between humans and systems. The targeted literature and policy search undertaken for this study, however, identified no peer-reviewed comparative analysis that examined agentic-AI assessment governance across multiple universities while systematically contrasting it with mature GenAI assessment policy. The gap claimed here is thus an emerging comparative-policy gap; it is not a claim that agentic AI has never been discussed.

The study addresses three research questions:

RQ1. How does the core problem represented in university assessment policy change from GenAI to agentic AI?

RQ2. What new governance boundaries emerge when AI can act on behalf of students beyond those already present in GenAI policy?

RQ3. How can assessment governance protect learner agency and credible evidence of learning while still permitting educationally valuable forms of agentic-AI use?

The central argument of this article is that the most important change introduced by agentic AI is a change in how agency is configured within assessment, not simply an increase in AI capability. Governance therefore needs to move beyond a tool-permission logic (whether AI was used) towards a responsibility-allocation logic. Such a logic asks what may be delegated, what evidence must remain attributable to the student, and what institutional actions may be authorised to an AI agent.

## **2. Literature Review and Theoretical Framework**

### **2.1 From Originality Governance to Assessment Design**

Early GenAI policies commonly focused on whether assessed work had been completed by the student. This concern is defensible because university assessment depends on an inferential chain: student performance produces observable work or behaviour, from which claims are made about knowledge, skill, or competence. When much of the observable product is generated externally, the inference from product to student capability becomes less secure.

Luo's (2024) WPR analysis shows that originality became a dominant way of representing the GenAI problem. It also argues that this representation may oversimplify distributed knowledge production in digitally mediated work. Policy and practice have since begun to move from blanket prohibition towards differentiated governance.

University College London (2026), for example, uses three categories of GenAI use in assessment: AI not permitted, AI used as an assistive tool, and AI integrated into the assessment activity. In the assistive category, AI may support the student but should not complete the assessed work on the student's behalf. In integrated assessment, by contrast, students may be expected to demonstrate critical and responsible use of AI.

Harvard University guidance reframes the key question from detecting AI to establishing whether work remains credible evidence of learning. Suggested mechanisms include oral examination, process checkpoints, in-class activity, and other forms of verification that expose a student's understanding rather than relying exclusively on the final artefact (Harvard University, 2026).

Monash University Malaysia (2026) articulates a closely related principle: restrictions on GenAI should rest on educational reasons, the nature of the assessment task, and the function of that task in producing evidence of student learning. On this view, permission or prohibition is a means rather than the final policy objective. The objective is to ensure that assessment continues to serve its evidentiary function.

UNSW Sydney has gone further. It describes a shift from permissions to purpose and connects AI rules to programme-level assurance of learning and secure assessment (Wilcox, Jones, & Beavis, 2026). This trajectory suggests that, even before agentic AI, stronger GenAI governance was already moving from protecting AI-free

products towards protecting credible evidence of learning. That principle remains relevant in an agentic-AI environment.

## 2.2 Agentic AI Changes Agency, Not Merely Output

Conventional GenAI can generate a paragraph, image, or code fragment. Agentic AI can instead accept a higher-level goal and then search, select information, invoke tools, access platforms, edit files, run code, send messages, or complete a workflow across multiple steps.

For assessment, the significance of this difference lies in how decision-making is distributed. A student who asks GenAI for three possible essay topics, and then independently chooses, researches, evaluates, and writes, retains substantial intellectual responsibility. Consider instead a student who instructs an agent to complete the assignment while the system enters the course platform, reads the instructions, searches for sources, drafts an answer, and submits it. That student has done something qualitatively different: they have delegated the learning task.

Queen's University (2026) makes this distinction explicit. Its guidance contrasts ordinary GenAI, which still depends on the user to prompt, revise, and decide what to submit, with agentic AI capable of carrying out multi-step work. When an autonomous agent performs coursework for a student, the policy problem resembles outsourcing academic work more than the use of an unauthorised content generator. The same guidance notes that unauthorised agentic-AI use may be difficult to infer from the final product alone. Visible or verifiable learning processes are therefore more sustainable than post-hoc detection.

This change can be conceptualised as a transfer of agency, in which some combination of goal setting, strategy selection, information acquisition, execution, and monitoring passes from the student to an AI system.

## 2.3 Self-Regulated Learning and Regulatory Substitution

Zimmerman (2002) conceptualises self-regulated learning as a cyclical process that includes goal setting, strategy selection, performance monitoring, and reflection. Learning thus involves more than producing a correct answer: planning, monitoring, error detection, and correction may themselves be educational outcomes.

Agentic AI creates what Salah et al. (2026) call a regulation paradox. The more capable a system becomes at autonomous planning, monitoring, and correction, the more easily it can substitute for the regulatory activities students need to practise. Their bounded-autonomy approach proposes limiting AI involvement according to learner development, timing, transparency, contestability, and duration.

This has direct implications for assessment governance. The relevant question extends beyond who generated a piece of text to which decisions and actions required by the learning task still belonged to the student. In this article, learner agency is treated as the observable allocation of decision rights and action responsibility within assessment, rather than as an abstract personal trait.

## 2.4 From WPR to Three Governance Boundaries

The study also draws on Bacchi's (2009) What's the Problem Represented to Be? approach. Policies do more than respond to a fixed external problem. By permitting, prohibiting, categorising, or requiring disclosure, they also construct what counts as the problem to be governed.

GenAI policy commonly represents the problem as originality, unauthorised assistance, or the authenticity of submitted work. Agentic AI exposes three related but non-substitutable governance questions:

1. Delegation boundary. Which cognitive or operational responsibilities may a student transfer to AI?
2. Evidence boundary. What evidence is required for a university to make a defensible judgement that learning belongs to the student?
3. Access boundary. Which systems, credentials, data, files, and institutional actions may an AI agent access or perform on the student's behalf?

These three dimensions form the Delegation-Evidence-Access (DEA) framework developed in this article.

## 3. Method

### 3.1 Research Design

The study uses exploratory comparative documentary analysis rather than a statistically representative survey. It does not aim to estimate the proportion of universities worldwide that already have agentic-AI policies. Instead, it identifies the conceptual dimensions that become visible when institutions begin to govern AI systems capable of autonomous action.

The study used theoretical sampling and a maximum-difference comparison. The first group, termed agentic-aware, includes universities whose public central guidance explicitly uses concepts such as agentic AI or AI agents and connects them to teaching, learning, assessment, academic integrity, or institutional systems. The second group, termed GenAI-dominant, comprises institutions with comparatively mature assessment governance whose selected documents

remain organised mainly around generative-AI policy. The second group is not a control group in a causal sense. It is used to identify which existing governance principles remain applicable when the technical object changes.

The document retrieval cut-off was 3 October 2026.

### 3.2 Cases and Materials

Twelve universities were included. The six agentic-aware cases were Queen's University, the University of Edinburgh, the University of Kansas, Saint Xavier University, Touro University, and Case Western Reserve University. The six GenAI-dominant comparators were the University of Cambridge, University College London, Stanford University, Harvard University, Monash University Malaysia, and UNSW Sydney. Table 1 summarises the institutional cases and the principal governance orientation identified in the selected public documents.

**Table 1.** Institutional cases and principal policy orientation

| Group          | University                      | Country/region | Principal orientation in selected document                                                                                                                     |
|----------------|---------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agentic-aware  | Queen's University              | Canada         | Distinguishes GenAI from agentic AI; frames delegated coursework as delegation; emphasises verifiable process, programme-level assurance, and system security. |
| Agentic-aware  | University of Edinburgh         | United Kingdom | Explicitly addresses generative and agentic AI and restricts agent access to login-protected learning and assessment environments.                             |
| Agentic-aware  | University of Kansas            | United States  | Notes that agents may log into learning systems and complete work autonomously; emphasises human oversight and institutionally supported tools.                |
| Agentic-aware  | Saint Xavier University         | United States  | Defines agentic AI at university level; establishes least-privilege, data/access control, human responsibility, and transparency principles.                   |
| Agentic-aware  | Touro University                | United States  | Defaults to prohibiting agentic AI from autonomously producing, solving, or deciding on behalf of students in graded assessment.                               |
| Agentic-aware  | Case Western Reserve University | United States  | Governance covers predictive, generative, and agentic AI; stresses that AI should not replace human judgement, original work, or critical thinking.            |
| GenAI-dominant | University of Cambridge         | United Kingdom | Centres on unacknowledged AI-generated content, academic misconduct, and assessment-specific requirements.                                                     |
| GenAI-dominant | University College London       | United Kingdom | Uses prohibited, assistive, and integrated categories and requires explicit assessment-level communication.                                                    |
| GenAI-dominant | Stanford University             | United States  | Treats unauthorised GenAI as comparable to external assistance; restricts substantive completion of work and allows course-level rules.                        |
| GenAI-dominant | Harvard University              | United States  | Emphasises verification of understanding, process evidence, oral examination, and independent performance.                                                     |
| GenAI-dominant | Monash University Malaysia      | Malaysia       | Requires AI restrictions to have educational justification and to reflect the learning evidence an assessment is intended to produce.                          |
| GenAI-dominant | UNSW Sydney                     | Australia      | Moves from tool permission towards educational purpose and programme-level assurance of learning.                                                              |

In this study, "not identified" means only that a provision was not identified in the selected public central document. It does not imply that no faculty, course, information-technology rule, or non-public institutional procedure addresses the issue.

### 3.3 Coding Framework

A six-dimensional directed coding framework was developed from the literature review and preliminary document reading:

A1. Explicit agentic-AI distinction. Whether the document explicitly distinguishes autonomously acting AI from conventional generative AI.

A2. Delegation boundary. Whether it explicitly regulates AI completing assessed cognitive or operational work on behalf of a student.

A3. Access boundary. Whether it explicitly regulates agent access to university accounts, credentials, learning-management systems, data, files, or institutional systems.

A4. Evidence of learning. Whether policy or assessment design is explicitly connected to credible evidence of what the student personally knows or can do.

A5. Task-level transparency. Whether the course, teacher, or assessment task is expected to communicate what AI use is permitted, restricted, or prohibited.

A6. Programme-level assurance of learning. Whether the governance model extends beyond individual tasks or courses to the independent demonstration of programme or degree-level outcomes.

Each dimension was coded E (explicit), P (partial or implicit), or - (not identified in the selected document). Analysis proceeded in two passes. The first applied the directed coding framework; the second re-examined negative cases and documents that did not fit the emerging explanation. Counts are descriptive only and are not treated as population estimates. Because the study did not employ two independent human coders for a complete blind recoding, it does not report Cohen's kappa or another inter-rater coefficient; this limitation is discussed below.

## 4. Results

### 4.1 Six-Dimensional Policy Coding

**Table 2.** Six-dimensional coding of assessment-governance documents

| University                      | Agentic distinction | Delegation | Access | Learning evidence | Task transparency | Programme assurance |
|---------------------------------|---------------------|------------|--------|-------------------|-------------------|---------------------|
| Queen's University              | E                   | E          | E      | E                 | E                 | E                   |
| University of Edinburgh         | E                   | E          | E      | P                 | E                 | P                   |
| University of Kansas            | E                   | E          | E      | P                 | -                 | -                   |
| Saint Xavier University         | E                   | E          | E      | P                 | E                 | -                   |
| Touro University                | E                   | E          | P      | E                 | E                 | -                   |
| Case Western Reserve University | E                   | P          | P      | P                 | E                 | -                   |
| University of Cambridge         | -                   | P          | -      | E                 | E                 | -                   |
| University College London       | -                   | P          | -      | E                 | E                 | P                   |
| Stanford University             | -                   | P          | -      | P                 | E                 | -                   |
| Harvard University              | -                   | P          | -      | E                 | E                 | -                   |
| Monash University Malaysia      | -                   | P          | -      | E                 | E                 | -                   |
| UNSW Sydney                     | -                   | P          | -      | E                 | E                 | E                   |

Note. E = explicit; P = partial or implicit; - = not identified in the selected document.

All six agentic-aware cases explicitly distinguished agentic AI. Five established an explicit delegation boundary and one addressed it partially. Four established explicit access boundaries and two addressed them partially. Two explicitly connected governance to evidence of learning and four did so partially. Five explicitly addressed task-level transparency, and programme-level assurance was explicit in one case and partial in one.

None of the six GenAI-dominant comparator cases explicitly distinguished agentic AI in the selected document. All six contained partial principles against AI replacing the student's assessed work, but none reached the agentic-specific delegation definition used in this study. No selected comparator document explicitly addressed an AI agent autonomously logging into an LMS, holding user credentials, or operating institutional systems on the student's behalf. On the other hand, five of the six comparators explicitly linked assessment to credible evidence of learning, and all six

addressed task-level transparency. Programme-level assurance was again relatively uncommon, with one explicit and one partial case.

In these cases, then, the principal structural difference lies in the addition of delegation and access as explicit governance variables, rather than in AI use itself.

## **4.2 Shift One: From Who Wrote the Product to Who Completed the Action Chain**

GenAI policies frequently state that students must submit their own work, that AI must not substantially complete an assessment, or that AI use must be disclosed. These principles already touch on delegation, but their practical object is usually the final product.

Stanford University, for example, treats unauthorised GenAI as analogous to assistance from another person and generally prohibits using it to complete substantive portions of assignments unless the instructor specifies otherwise. University College London states that, in its assistive category, AI can support a student but should not complete the assessed work on that student's behalf. The principle that assessed work should not simply be outsourced is thus older than agentic AI.

What changes is the scope of outsourcing. An autonomous agent can potentially cover the whole action chain, not only the final text. Queen's University provides the clearest example, locating its policy distinction in delegation rather than in the form of content. Touro University takes a more restrictive approach: by default, agentic systems may not autonomously generate, solve, or decide on behalf of students in graded assessment. Despite these different policy choices, both cases shift the object of governance from what AI generated to which responsibilities the student transferred to AI.

## **4.3 Shift Two: System Access Becomes Part of Assessment Governance**

Access produced the clearest structural difference between the two groups. In a conventional GenAI workflow, a student may copy an assessment question into an external interface, receive content, and return to the university system to submit it. Assessment governance and digital-infrastructure governance can thus remain partly separate.

Agentic AI collapses that separation. The University of Kansas (2026) warns that an agent may log into a learning system and complete work autonomously, which raises both academic-integrity and persistent-system-access concerns. Saint Xavier University (2026) defines agentic AI in terms that include browsing websites, logging into systems, sending messages, retrieving files, invoking APIs, executing workflows, and making multi-step decisions. It applies least-privilege principles to data and access rights. Queen's University (2026) links assessment design with credential and system security and explicitly recognises potential interaction with learning-management systems.

Once an AI agent can use a student identity to read assessment instructions, access course resources, modify files, or submit work, identity authentication, system permission, and assessment validity become different facets of the same educational-governance problem.

## **4.4 Shift Three: The Evidence-of-Learning Logic Remains Central**

A comparison based purely on terminology might suggest that mature GenAI policy becomes obsolete once institutions begin to discuss agentic AI. The findings do not support that conclusion. The GenAI-dominant comparator group was particularly strong on learning evidence and task-level transparency.

Harvard University (2026) explicitly encourages assessment design that verifies a student's understanding through oral examination, process explanation, in-class work, or other checks beyond a large take-home product. Monash University Malaysia (2026) places the educational legitimacy of an AI restriction in the function of the assessment as evidence of learning. University College London (2026) differentiates learning outcomes that must be demonstrated independently from those that can be supported by AI and those in which responsible AI use is itself part of the assessed capability. UNSW Sydney links secure assessment to programme learning outcomes (UNSW Sydney, 2026; Wilcox et al., 2026).

The durable principle, then, is not the preservation of AI-free work. It is that assessment must generate evidence strong enough to support the capability claim made about the student.

## **4.5 Shift Four: Task-Level Rules Are More Developed Than Programme-Level Assurance**

Programme-level assurance was sparse in both groups, with one explicit and one partial case in each. This matters because a set of clear task-level AI rules does not, by itself, show that a graduate independently possesses the capabilities the degree programme claims.

Queen's University (2026) explicitly calls for assurance of learning at programme level. It argues that decisions about which outcomes require strong individual evidence should not be left entirely to individual course instructors. UNSW Sydney's assessment policy connects secure assessment with programme learning outcomes and places independent capability, professional practice, and AI capability within a programme architecture (UNSW Sydney, 2026; Wilcox et al., 2026).

This leaves a governance gap. Universities are becoming more capable of specifying how AI may be used in a single assignment. Many, however, have less explicit answers to where and how a graduating student independently demonstrates the critical capabilities the programme promises.

## 5. Discussion

### 5.1 Agentic AI Creates an Ownership-of-Action Problem

GenAI assessment debates have largely asked who wrote the product. Agentic AI adds a second question: who performed the chain of consequential judgements and actions required to complete the learning task? The two questions are not equivalent.

A student may personally press the final submit button after delegating problem decomposition, evidence search, source evaluation, argument planning, code execution, debugging, and conclusion formation to an agent. Conversely, a student may let an agent execute many mechanical operations while retaining problem definition, evaluation criteria, audit responsibility, and final judgement. A single axis of use versus non-use cannot adequately capture the educational meaning of AI involvement.

Assessment governance should instead examine the allocation of learning responsibility.

### 5.2 The Delegation-Evidence-Access Framework

The comparative analysis supports a three-boundary DEA framework for assessment governance.

#### 5.2.1 Delegation Boundary

The first question is which cognitive and operational responsibilities may be transferred to AI. Delegation should not be classified solely by whether AI is used, because its educational meaning depends on the learning outcome. If a course is designed to develop literature-evaluation capability, an agent that autonomously searches, filters, evaluates, and judges evidence quality may replace the target capability. If the learning outcome is the audit and governance of an AI research workflow, configuring an agent, examining its trace, and critiquing its errors may be an appropriate assessment activity. The same technical behaviour can thus carry different educational meanings under different learning objectives.

#### 5.2.2 Evidence Boundary

The second question is what justifies believing that the capability belongs to the student. In an agentic-AI environment, the evidentiary value of a final product alone is likely to decline further. Assessment can strengthen inference through three complementary forms of evidence. Product evidence includes the final essay, programme, design, report, or solution. Process evidence includes drafts, version histories, decision rationales, interaction logs, error corrections, or research pathways. Performance evidence includes oral defence, live demonstration, a resource-constrained independent task, live explanation, or a transfer task.

This does not imply that every course should return to invigilated examinations. Evidence requirements should be proportionate to the importance and risk of the learning outcome. High-stakes programme capabilities require stronger independent evidence, whereas low-stakes formative tasks can permit substantially more AI-supported experimentation.

#### 5.2.3 Access Boundary

The third question is the one conventional assessment policy most easily misses: where may an agent go, and what may it do on a student's behalf? Governance may need to specify whether an agent may hold or use university credentials, access an LMS autonomously, read non-public course materials, or access information about peers or staff. It may also need to specify whether an agent may modify platform data, submit work, run persistently in the background, invoke third-party services or APIs, execute high-impact actions without human confirmation, or act without an auditable record.

In an agentic-AI environment, access control is more than an information-security detail. It affects whether assessed activity can still be defensibly attributed to the student.

### 5.3 From AI Traffic Lights to Four Educational Contexts

The DEA framework can also organise assessment into four educational contexts, so that institutions need not write new rules for every product release.

#### 5.3.1 Independent Capability Demonstration

AI and external agents do not perform the critical parts of the task. This context suits foundational capabilities, professional or licensure-related competencies, high-stakes programme outcomes, and other situations in which the institution must establish independent mastery. Evidence may include live performance, oral examination, controlled

practice, practical demonstration, or another credible method. The purpose is to obtain independent evidence, not to oppose AI.

### 5.3.2 AI-Assisted Learning

AI is permitted for bounded support such as feedback, language editing, brainstorming, or other activities that do not replace the core learning outcome. The student retains responsibility for problem definition, judgement, reasoning, and final decisions. This context is broadly compatible with mature assistive GenAI policies.

### 5.3.3 Bounded Agentic Collaboration

An agent may execute specified subtasks, but its autonomy is constrained. It might search public material without making the final judgement on evidence quality, execute code that the student must explain, organise data without using university credentials to enter protected systems, or prepare actions that require human confirmation. Process traces may be retained where they are educationally necessary. Of the four contexts, this one most directly operationalises the bounded-autonomy principle proposed by Salah et al. (2026).

### 5.3.4 Agentic AI as an Assessed Capability

In programmes where professional work will require students to design, deploy, supervise, or audit autonomous AI, a comprehensive prohibition may reduce authenticity. The assessed capability can instead include setting appropriate goals, defining permission boundaries, identifying errors and risks, evaluating agent actions, taking over after failure, explaining final judgements, and accepting responsibility for outcomes. Queen's University (2026) already recognises that supervised agentic-AI use can have educational value when configuring, deploying, and critically evaluating agents is itself the learning objective.

The relevant policy question is therefore what role agentic AI may play in relation to the outcome being assessed, rather than whether it is permitted in the abstract.

## 6. Repositioning Academic Integrity

### 6.1 From Authorship to a Chain of Responsibility

Traditional academic-integrity rules often focus on authorship: who wrote this sentence or produced this answer? In an autonomous-agent environment, questions about responsibility across the action chain have more explanatory power. Who set the goal, chose the strategy, obtained information, evaluated evidence, executed the action, checked for errors, decided to submit, and accepted responsibility for the result?

If most of these roles are delegated to AI, disclosure alone may not restore assessment validity. Disclosure can solve a transparency problem, but it does not by itself establish learning. Transparency is necessary in many contexts, yet it is not equivalent to evidence of competence.

### 6.2 Why AI Detection Becomes Less Central

When AI operates as an autonomous or semi-autonomous agent, the final artefact may reflect multiple tool calls, human revisions, system interactions, and operational steps. Inferring the degree of AI involvement from the final text alone becomes even less defensible.

Queen's University (2026) notes that unauthorised agentic-AI use may be difficult to detect reliably from a final submission. It therefore emphasises assessment processes that make student learning visible or verifiable. This direction is consistent with broader GenAI assessment research. Lye and Lim (2024) argue that institutional responses should move beyond detection and prohibition towards assessment redesign, and Kofinas, Tsay, and Pike (2025) show that even so-called authentic assessment cannot simply be assumed to be immune to GenAI intervention.

A more sustainable question than how to catch students who used AI is whether the assessment still produces credible evidence of human learning.

## 7. Policy Implementation Model: Task, Course, and Programme

### 7.1 Task Level

Each significant assessment should specify the learning outcome and which AI activities are permitted. It should also state which cognitive responsibilities may not be delegated, whether agentic AI may execute defined subtasks, whether an agent may access any university system, whether disclosure or process evidence is required, and how the student will demonstrate understanding of the final result. This moves policy beyond the phrase "AI permitted" towards an explicit allocation of responsibility.

## 7.2 Course Level

Course leaders need to examine the portfolio of tasks rather than each assignment in isolation. A clear AI rule for every assignment does not automatically show that a course achieves its learning outcomes. The course should offer opportunities to learn responsible AI use and to develop core capabilities independently. It should include at least one credible means of verifying critical outcomes and give educational rather than arbitrary reasons for restrictions or permissions. This approach aligns closely with the learning-evidence principle articulated by Monash University Malaysia (2026).

## 7.3 Programme Level

Programme governance must answer a question that individual courses cannot fully resolve: where does a student who receives the degree independently demonstrate the critical capabilities the programme claims to certify? Institutions can map each programme learning outcome to at least one assessment node capable of producing strong individual evidence. At the same time, they can ensure that students learn professionally relevant forms of AI collaboration.

UNSW Sydney (2026) has already moved part of assessment assurance to the level of programme learning outcomes by requiring secure assessment of programme or student learning outcomes. Agentic AI makes this programme-level architecture more important, because autonomous systems can operate across tasks and courses rather than within a single assignment.

## 8. Significance for Malaysian Higher Education

The Malaysian higher-education context is not peripheral to this transition. The Ministry of Higher Education Malaysia (2024) has already issued guidance on generative AI in higher-education teaching and learning, covering governance, teaching and learning, academic integrity and ethics, and professional development. The Higher Education Digitalisation Policy 2025-2030 further requires institutions to update guidance, planning, and strategy as AI capability and ethical expectations develop (Ministry of Higher Education Malaysia, 2025).

Agentic AI therefore does not necessarily require a wholly separate governance architecture. A more realistic approach treats autonomous agency as a new risk, and a new educational object, within existing AI governance.

From a Malaysian and ASEAN perspective, Teoh and Tan (2026) identify persistent gaps between human-centred values and operational oversight, institutional autonomy and accountability, technological adoption and organisational capacity, and regional principles and implementation readiness. They argue that universities need to translate high-level AI principles into governance that institutions can implement. The DEA framework contributes to that translation for assessment.

Five priorities follow for Malaysian institutions, including private universities. First, central AI policy should distinguish content-generating systems from systems that can act on behalf of users. Second, university credentials, passwords, LMS access, student records, and automated submission should be incorporated explicitly into AI assessment governance under least-privilege principles. Third, assessment briefs should specify which learning responsibilities may not be delegated, in addition to whether AI is allowed. Fourth, degree programmes should identify the outcomes that require independent individual evidence instead of leaving all such decisions to individual course instructors. Fifth, programmes that genuinely require agentic-AI capability should provide controlled or sandboxed environments in which students can learn to configure, supervise, audit, and correct agents.

This approach addresses two opposing risks. One is allowing agentic AI to replace capabilities students must develop; the other is restricting AI so heavily that graduates are unprepared for professional forms of human-AI collaboration.

## 9. Theoretical Contributions

### 9.1 From Ownership of Content to Ownership of Action

Existing GenAI assessment research has shown the difficulty of treating originality as a simple binary property in AI-mediated knowledge production (Luo, 2024). The present study extends that insight: once AI systems can act autonomously, content ownership alone is insufficient. Assessment governance must also consider ownership of action, that is, who performed the critical cognitive and institutional actions required to produce the claimed learning outcome. This concept helps explain why two artefacts that both involved AI can have very different educational meanings.

### 9.2 Bringing Access Rights into Assessment Theory

Information-security governance normally addresses accounts, data, APIs, and system permissions, whereas assessment theory focuses on validity, reliability, and learning outcomes. Agentic AI connects these domains directly. When an agent can enter a course platform under a student identity and complete actions, access rights affect whose performance

is actually being assessed. Access thus becomes a condition of assessment validity as well as a technical-security concern.

### 9.3 Moving Assessment Governance from Assignments to Programme Assurance

Because agentic systems can persist and act across tasks, fragmented assignment-level rules are less able to assure graduate capability. This study therefore distinguishes the governance function of course rules from that of programme assurance. Course rules govern AI use within learning activities, whereas programme assurance protects the credibility of the qualification. Future research on AI and assessment should examine programme architectures as well as individual assignments.

## 10. Practical Contribution: A DEA Decision Table

**Table 3.** DEA-based assessment-governance decision table

| Decision dimension   | Question for educators/programme leaders                                        | Illustrative governance response                                                                                                                   |
|----------------------|---------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Delegation (D)       | Which reasoning, judgement, and actions must the student perform personally?    | Specify non-delegable activities; define agent-executable subtasks; require explanation of consequential decisions.                                |
| Evidence (E)         | What evidence is sufficient to show that the capability belongs to the student? | Combine product, process, and performance evidence; oral defence; live demonstration; version history; transfer task.                              |
| Access (A)           | Which systems may an agent enter and what actions may it perform?               | No credential sharing; least privilege; prohibit autonomous submission where appropriate; controlled APIs; human confirmation of critical actions. |
| Transparency         | Can students clearly determine the rules?                                       | State the authorised AI role, limits, and disclosure requirements in the assessment brief.                                                         |
| Programme assurance  | Where does the degree independently evidence its core capabilities?             | Map programme outcomes to secure assessment nodes and to a programme-level AI-capability map.                                                      |
| Equity and inclusion | Do the rules create unjustified barriers for particular students?               | Review accessibility; provide equivalent tools or routes; preserve reasonable adjustments.                                                         |
| Data governance      | What sensitive information can an agent encounter?                              | Use institutionally approved tools; classify data; restrict sensitive inputs; maintain audit trails where justified.                               |
| Human responsibility | Who remains accountable for consequential decisions?                            | Require human-in-the-loop review and make the student responsible for checking and explaining the final submission.                                |

The main advantage of the table is that it is technologically neutral. Even if students stop using today's chat-oriented products, the DEA questions remain applicable whenever a system can execute tasks and act as an agent.

## 11. Relationship to Existing University Policies

The findings do not imply that every university should immediately create a separate document titled Agentic AI Policy. The cases show four institutional pathways.

The first is explicit extension, in which an existing GenAI framework is expanded with agentic-AI definitions and rules, as at Queen's University and Saint Xavier University. The second is strong restriction, in which autonomous agency is prohibited by default in assessed activities, as at Touro University. The third is principle absorption: agentic AI is included in broader human-centred AI governance, but assessment-specific operationalisation remains less detailed, as at Case Western Reserve University. The fourth is technologically neutral assurance. Here, policies may not yet use agentic terminology but have comparatively strong structures for evidence of learning, secure assessment, and programme assurance, as at UNSW Sydney.

Policy modernisation should not be reduced to adding the term agentic. A policy may use the term yet remain weak if it does not define what may be delegated, which systems may be accessed, and how student capability will be verified. Conversely, a policy without agentic terminology may still be comparatively resilient if it is built around robust learning assurance.

## 12. Limitations

Five limitations qualify the interpretation of this study. First, the 12 cases were selected theoretically and are not a statistically representative sample of universities worldwide. Counts such as five of six describe the selected cases and must not be treated as estimates of global prevalence.

Second, the agentic-aware group was sampled partly because explicit agentic-AI guidance existed. The group difference in explicit agentic terminology is therefore partly a consequence of the design. The analytically more important dimensions are the downstream differences in delegation, access, evidence, and programme assurance.

Third, the analysis focused mainly on public, central, English-language documents. Faculty guidance, course outlines, internal information-security controls, and non-public procedures may contain additional rules. "Not identified" cannot therefore be read as "institutionally absent".

Fourth, policy text is not policy implementation. Tsao (2025) shows that even where formal AI rules exist, ambiguity, resources, and contextual interpretation can lead students and teachers to enact them differently. Further research should examine how delegation and access boundaries are negotiated in practice across teaching, student-support, governance, and information-technology roles.

Fifth, the coding involved two analytic passes and a negative-case review, but not full blind duplicate coding by two independent human coders, so no inter-rater reliability statistic is reported. A larger follow-up study should pre-register the coding manual, expand the cross-national sample, and conduct independent double coding.

Finally, agentic-AI capabilities are changing quickly. The analysis is a policy snapshot up to 3 October 2026, and longitudinal policy research would be more informative than a one-time classification.

## 13. Future Research Agenda

The findings support at least five empirical extensions. First, a cross-national longitudinal database could track movement from content-generation governance towards delegated-action governance between 2023 and 2028. Second, assessment experiments could compare how well final-product evidence, process evidence, oral defence, and transfer tasks establish student capability in agentic-AI environments. Third, studies of different levels of agent autonomy could test whether planning, monitoring, error detection, and transfer are supported or displaced, directly examining the regulatory-substitution hypothesis. Fourth, Malaysian and ASEAN research could compare the readiness of public and private institutions with respect to credentials, LMS access, agentic AI, and programme-level learning assurance. Fifth, equity research should examine cases in which AI agents function as reasonable adjustments for students with disabilities. It should also determine how policy can preserve core learning outcomes without misclassifying necessary assistive technology as academic misconduct.

These questions bear more directly on the long-term problem of assessment validity than continued reliance on improvements in AI-writing detection.

## 14. Conclusion

Agentic AI does more than add a more capable writing tool to higher education; it changes the distribution of agency within assessment. In conventional GenAI contexts, institutions were mainly concerned with AI-generated content entering a student's work. In agentic-AI contexts, systems can also take on planning, search, judgement, operation, system access, and potentially task submission. Permission to use AI no longer provides sufficient conceptual resolution.

In the 12 public policy cases compared here, the clearest new governance structures concern the delegation of learning responsibility and access to digital systems. At the same time, mature GenAI assessment reforms, represented by cases such as Harvard, UCL, Monash University Malaysia, and UNSW, suggest that agentic AI does not require institutions to discard existing policy learning. The principle that assessment must generate credible evidence of student capability remains the central bridge between the two phases.

The DEA framework therefore asks three linked questions. Delegation: what may not be transferred to an agent? Evidence: what is sufficient to show that learning belongs to the student? Access: where may AI act on the student's behalf, and what may it do? All three are required. Delegation rules without evidence specify prohibited behaviour but do not establish competence. Evidence rules without access governance may redesign assessment while allowing autonomous agents to control institutional systems under a student identity. Access security without delegation boundaries may protect infrastructure without protecting the learning process.

In the age of agentic AI, the objective of assessment governance should not be to preserve an AI-free learning environment. It should be to maintain a defensible distinction among human learning, AI-supported learning, and delegated AI action. For future universities, the key question is no longer simply whether a student used AI. It is which capabilities, judgements, and responsibilities can credibly be attributed to the student in the performance on which credit and qualifications are awarded.

### Author Contributions (CRediT)

Wong Kaiyan and Sun Bojian contributed equally (50% each) to Conceptualization, Methodology, Investigation, Formal analysis, Data curation, Writing - original draft, and Writing - review & editing. Both authors approved the final manuscript.

### Funding

This research received no external funding.

### Ethics Statement

This study analysed publicly available institutional and policy documents and did not involve recruitment of human participants, intervention, or collection of private personal data.

### Data Availability Statement

The institutional policy documents analysed in this study are publicly available through the sources listed in the References. Because university AI policies are dynamic documents, the analysis should be interpreted as a policy snapshot with a retrieval cut-off of 3 October 2026.

### Conflict of Interest

The authors declare no conflicts of interest.

### Appendix A: Coding Framework

A1 Explicit agentic-AI distinction: explicit separation of autonomously acting AI from conventional generative AI. A2 Delegation boundary: explicit regulation of AI completing assessed cognitive or operational work for a student. A3 Access boundary: explicit regulation of agent access to accounts, credentials, learning-management systems, data, files, or institutional systems. A4 Evidence of learning: explicit linkage between policy or assessment design and credible evidence of what the student personally knows or can do. A5 Task-level transparency: an expectation that permitted, restricted, or prohibited AI use is communicated at course or assessment level. A6 Programme-level assurance: governance extending beyond individual tasks to independent demonstration of programme or degree-level learning outcomes. Codes: E = explicit; P = partial or implicit; - = not identified in the selected document.

### References

- An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22, 10. <https://doi.org/10.1186/s41239-025-00507-3>
- Bacchi, C. (2009). *Analysing policy: What's the problem represented to be?*. Pearson Australia.
- Baradziej, S. (2026). A systematic review of the roles of agentic AI in higher education. *Social Sciences & Humanities Open*, 14, 103418. <https://doi.org/10.1016/j.ssaho.2026.103418>
- Case Western Reserve University. (2026). Governance & policies: AI at CWRU. Retrieved October 3, 2026, from <https://case.edu/ai/governance-policies>
- Dai, Y., Lai, S., Lim, C. P., & Liu, A. (2025). University policies on generative AI in Asia: Promising practices, gaps, and future directions. *Journal of Asian Public Policy*, 18(2), 260-281. <https://doi.org/10.1080/17516234.2024.2379070>
- Harvard University, Derek Bok Center for Teaching and Learning. (2026). Guidance for faculty on addressing AI-related academic integrity issues. Retrieved October 3, 2026, from <https://bokcenter.harvard.edu/guidance-faculty-addressing-ai-related-academic-integrity-issues>
- Kofinas, A. K., Tsay, C. H. H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*, 56(6), 2522-2549. <https://doi.org/10.1111/bjet.13585>
- Liu, X., Tang, Q., Wu, W., Chen, X., & Wang, C. (2026). GenAI-agents in education: A systematic review of macro-level trends, educational roles, theoretical foundations and future agendas. *Interactive Learning Environments*. <https://doi.org/10.1080/10494820.2026.2672562>

- Luo, J. J. (2024). A critical review of GenAI policies in higher education assessment: A call to reconsider the originality of students' work. *Assessment & Evaluation in Higher Education*, 49(5), 651-664. <https://doi.org/10.1080/02602938.2024.2309963>
- Lye, C. Y., & Lim, L. (2024). Generative artificial intelligence in tertiary education: Assessment redesign principles and considerations. *Education Sciences*, 14(6), 569. <https://doi.org/10.3390/educsci14060569>
- Ministry of Higher Education Malaysia. (2024). Garis Panduan Penggunaan Teknologi Kecerdasan Buatan Generatif (KBG) Dalam Pengajaran dan Pembelajaran (PdP) Pendidikan Tinggi. <https://jpt.mohe.gov.my/portal/index.php/ms/penerbitan/175-garis-panduan>
- Ministry of Higher Education Malaysia. (2025). Higher Education Digitalisation Policy 2025-2030. <https://www.mohe.gov.my/en/download/publications-journals-and-reports/dasar-pendigitalan-pendidikan-tinggi-2025-2030>
- Monash University Malaysia. (2026). Education policy and procedure: Academic integrity and generative AI. Retrieved October 3, 2026, from <https://www.monash.edu.my/education-excellence/Education-Policy-and-Procedure/academic-integrity>
- Pauzi, Z., Vizgirda, V., Duran, L. D. D., & Mavrikis, M. (2026). Towards evaluating student agency in AI-mediated learning. In Companion Publication of the 2026 ACM Designing Interactive Systems Conference, 221-225. <https://doi.org/10.1145/3802974.3809419>
- Queen's University. (2026). Statement from the Vice-Provost (Teaching and Learning) on agentic artificial intelligence. Retrieved October 3, 2026, from <https://www.queensu.ca/provost/teaching-and-learning/teaching-and-learning-communications/statements>
- Salah, M., Al-Busaidi, R., Al-Ghawi, A., & Shawky, W. (2026). The regulation paradox: Agentic AI, bounded autonomy, and self-regulated learning in higher education. *Frontiers in Education*, 11, 1902833. <https://doi.org/10.3389/feduc.2026.1902833>
- Saint Xavier University. (2026). Policy on artificial intelligence and agentic AI. Retrieved October 3, 2026, from <https://www.sxu.edu/about/offices/information-technology/ai-guidelines-policy.html>
- Stanford University Teaching Commons. (2026). Artificial intelligence teaching guide: Analyzing the implications of AI for your course. Retrieved October 3, 2026, from <https://teachingcommons.stanford.edu/teaching-guides/artificial-intelligence-teaching-guide/analyzing-implications-ai-your-course>
- Teoh, K. B., & Tan, W. L. (2026). Policy alignment and trustworthy AI governance in higher education: A conceptual analysis of Malaysia's reform agenda and ASEAN governance imperatives. *Journal of Interdisciplinary Studies in Education*, 15(4), 1-22. <https://doi.org/10.32674/eqphhh11>
- Touro University. (2026). Policy for appropriate use of artificial intelligence in coursework. Retrieved October 3, 2026, from <https://www.touro.edu/students/policies/academic-integrity/artificial-intelligence-ai-addendum/>
- Tsao, J. (2025). Trajectories of AI policy in higher education: Interpretations, discourses, and enactments of students and teachers. *Computers and Education: Artificial Intelligence*, 9, 100496. <https://doi.org/10.1016/j.caeai.2025.100496>
- University College London. (2026). Three categories of GenAI use in assessment. Retrieved October 3, 2026, from <https://www.ucl.ac.uk/teaching-learning/generative-ai-hub/three-categories-genai-use-assessment>
- University of Cambridge. (2026). Generative AI and assessment. Retrieved October 3, 2026, from <https://blendedlearning.cam.ac.uk/artificial-intelligence-and-education/generative-ai-and-assessment>
- University of Edinburgh. (2026). Using generative AI in your studies: Guidelines for students. Retrieved October 3, 2026, from <https://information-services.ed.ac.uk/computing/comms-and-collab/elm/generative-ai-guidance-for-students/using-generative>
- University of Kansas. (2026). AI agents. Retrieved October 3, 2026, from <https://ai.ku.edu/ai-agents>
- UNSW Sydney. (2026). Assessment and progression policy. Retrieved October 3, 2026, from <https://www.unsw.edu.au/governance/policy/policy-hub/assessmentandprogressionpolicy>
- Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in higher education: Seeing ChatGPT through universities' policies, resources, and guidelines. *Computers and Education: Artificial Intelligence*, 7, 100326. <https://doi.org/10.1016/j.caeai.2024.100326>

- Wilcox, T., Jones, M. I., & Beavis, A. (2026, August 12). From permissions to purpose: A program level approach to assessment in an AI-integrated university. UNSW Sydney. <https://www.education.unsw.edu.au/news-events/news/permissions-purpose-program-level-approach-assessment-ai-integrated-university>
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64-70. [https://doi.org/10.1207/S15430421TIP4102\\_2](https://doi.org/10.1207/S15430421TIP4102_2)