Externalising tacit institutional knowledge:
An analytic autoethnography of an AI-augmented compliance reporting workflow
Working Paper [Preprint]

Dr Brigid Freeman
University of Melbourne, Melbourne, Australia
July 2026

Abstract

Scholarship on generative artificial intelligence (AI) in higher education has concentrated on student-facing applications and prominent corporate services functions, with limited attention to AI-augmented professional work in specialised compliance domains. This paper addresses that gap through an analytic autoethnography of a report governance system developed for multi-jurisdictional government compliance reporting at an Australian public university. Drawing on a documentary audit trail, the analysis maps five dimensions of the System across three analytical layers: the inherited reporting context, the AI capabilities deployed and the governance design developed. The paper is grounded in organisational knowledge codification, boundary objects, infrastructure and articulation work. The case shows how the requirements of AI-augmented work prompted tacit institutional knowledge to be articulated, structured and verified in explicit governance artefacts. It identifies partial AI-mediated externalisation as the central theoretical contribution and presents the two-surface AI workflow and AI model-agnostic design as potentially transferable design logics. The analysis also demonstrates how conventional outputs obscure the sophisticated design work supporting them, with implications for institutional recognition, professional development, system ownership and maintenance. The findings are bounded by a single-practitioner, single-institution case, with successor operation and accountability infrastructure identified as priorities for further inquiry.

Keywords: generative AI, higher education compliance, AI-augmented governance, knowledge codification, analytic autoethnography, multi-jurisdictional reporting

Introduction

Generative AI's rapid integration into higher education has prompted extensive scholarly attention. Research has concentrated on student-facing applications and institutional gatekeeping and legitimacy concerns, including assessment and academic integrity (Adamakis & Rachiotis, 2025; Bittle & El-Gayar, 2025; Ncube et al., 2026), pedagogy and cognitive offloading (Lodge & Loble, 2026), student learning outcomes (Liu et al., 2025) and the foundational need to develop AI literacy (Almatrafi et al., 2024). Few studies have examined broader institutional domains, including administration (Buame et al., 2025). AI-augmented professional work in specialised compliance settings, where recurring deliverables to external authorities are substantial and complex, remains largely unexamined.

This paper presents a documented case in such a setting: the introduction of a report governance system within a discrete organisational unit (Faculty) at an Australian public university (hereafter, the University). The System was designed and documented by the system developer to manage multi-jurisdictional compliance reporting to the Commonwealth Government and two state government departments for a suite of funded university coursework programs. The paper traces the codification of tacit institutional knowledge into structured artefacts, the infrastructure through which those artefacts are organised and retained, and the coordination labour required to develop and sustain that infrastructure. Its central contribution is to show how AI mediation can support a practitioner to partially externalise tacit institutional knowledge into governance infrastructure designed for durability, auditability and transfer across staff and technological change.

Methodologically, the paper draws on the practitioner reflection tradition, which uses reflective and autoethnographic methods to examine knowledge developed through sustained engagement with generative AI (Creely & Blannin, 2025; Freeman, 2026; Panke, 2025). It extends this tradition from individual creative and academic practice into institutional governance work. Adjacent literature on cognitive and competency dynamics in higher education teaching and learning frames the broader discussion but is not central to the analysis, which concerns the design and establishment of a governance system.

Reflective accounts of AI-augmented governance system design are not yet a developed body of scholarship. Practitioners positioned to produce such accounts may have limited outlets or incentives for documented reflection, while the necessary combination of compliance expertise, AI capability literacy and reflective discipline is unevenly distributed across institutional roles. This paper therefore treats the case as both a design account and an analytic window onto AI-augmented institutional knowledge work. It pursues one descriptive and one interpretive question. Descriptively, it asks what design choices, deployed across the dimensions of multi-jurisdictional compliance reporting, characterise an AI-augmented governance system designed for continuity across staff transitions and AI platform change. Interpretively, it asks what the design of such a system reveals about knowledge codification, infrastructure and articulation work in compliance practice.

The contributions are bounded by the single-developer and single-institution character of the case and are offered for further examination rather than as generalisable findings. Empirically, the paper documents an operational system designed to strengthen continuity in a high-stakes compliance domain. Methodologically, it offers a framework for disaggregating inherited reporting complexity, deployed AI capability and governance design. Theoretically, it develops an account of partial AI-mediated externalisation, showing how the requirements of AI-augmented work prompt tacit institutional knowledge to be articulated, structured and verified in explicit artefacts. Boundary-object theory is used to interpret how shared governance artefacts mediate between human work functions within the AI-augmented workflow. Practically, the paper identifies the two-surface AI workflow and AI model-agnostic design as potentially transferable design logics for compliance settings characterised by recurring obligations, staff transition and technological change.

The case

The reporting environment

The system was developed to govern compliance reporting for multiple funding instruments held by the University with three Australian governments, comprising contracts, memorandums of understanding (MOUs) and contract variations extending over an eight-year period. These instruments fund a suite of accredited programs ranging from diplomas and bachelor's degrees to master's coursework degrees. The resulting reporting environment featured considerable complexity. Reporting cycles overlap but are not identical across funding authorities, with individual agreements establishing their own milestones, deadlines, content requirements and submission channels. Furthermore, accredited programs differ in terms of cohort intake patterns and required data. Accountability expectations require submission-grade documentation with traceable sourcing for all quantitative claims. The executed contracts and their Deeds of Variation accumulate over time, so that the operative requirements for any single report sit dispersed across multiple documents. In the twelve-month development period analysed, approximately 20 reports were identified as falling due across the three jurisdictions, each typically requiring multiple components including narrative reports, data attachments, financial declarations, child safety statements and statutory letters.

Prior to the system's development, reporting knowledge was held substantially in individual staff members' dispersed working files, and each cycle proceeded as a near-bespoke effort. Tacit knowledge about government report structure, language conventions, source hierarchies and acceptable framing was carried by the system developer, with limited codification into reusable artefacts. Institutional exposure to staff transition was correspondingly high. This vulnerability reflects a wider sectoral challenge, in which turnover and recurrent restructuring in professional roles disrupt continuity and contribute to the loss of institutional memory and knowledge amid the proliferation of compliance and accountability requirements (Woelert et al., 2026). The absence of durable, shared artefacts therefore created a material risk that reporting capability would be weakened or lost when responsibility moved between staff.

The Report Governance System

The new AI-augmented report governance system (hereafter, the *System*) is the centralised digital repository and operating framework designed by the system developer to address this environment. The System is located within the organisational unit's Document Library, itself reconstructed during the project's development phase. While a SharePoint-based document repository existed at the commencement of the project, its structure and interface were no longer considered fit-for-purpose, and a decision was taken to establish a new Library. Approximately 350 individual files were downloaded, consistent nomenclature applied and a new folder hierarchy designed. The files were uploaded into a newly established Library located within the University's SharePoint and Teams environment. The Library is therefore substantially reconstructed and constitutes the foundational layer on which the System rests.

The new System consolidates within the Library the funding instruments, report templates, data files, working drafts and supporting evidence required to meet the reporting obligations of all funders, organised across multiple folders. A master timeframes and requirements document carries an embedded action tracker and jurisdictional schedule tables identifying report obligations and timeframes. AI-ready Word templates exist for multiple recurring narrative report types, each with an embedded data dictionary. An all-fields data dictionary workbook covers approximately 150 unique data fields, and a consolidated all-data archive carries an Index sheet listing the source files. The data dictionary workbook and all-data archive are AI-enabled (i.e., Excel spreadsheets). The System's pre-created working subfolders in SharePoint exist for several forthcoming reports, each holding a populated reference subfolder with the most recent submitted version of the same report type and a blank working subfolder for the new draft. The reference copies support style continuity and AI context-loading. The (internal) Operational Guide developed alongside the System describes each component, the source-of-truth principles, the forward timeline and the operational walkthrough. This reflection draws on the Guide as the most complete description of the System.

The role and developmental trajectory

The project was undertaken by a practitioner with prior experience in higher education academic and professional roles, and expertise in program management and compliance reporting. Prior expertise did not extend to digitised compliance system design, technology-enabled template engineering, university AI-augmented workflow design or infrastructure development. Familiarity with AI-augmented practice developed through self-directed creative projects (Freeman, 2026), continuous professional development, and research. Those activities established competencies in multi-model orchestration, prompt design and generative AI as a diagnostic mirror for surfacing tacit criteria. This paper's case translates those competencies into the governance domain. Where the earlier project treated AI as a generative source refined against tacit aesthetic criteria, the System treats AI as a codification mechanism externalising aspects of tacit institutional knowledge into durable governance artefacts, in a domain with different stakes and verification standards.

Investigation and development of an AI-augmented design was deliberate. The alternative was to continue the predominantly manual reporting practice previously pursued at the Faculty, which met contractual obligations but was variable in quality and labour-intensive. The case is therefore one of voluntary AI adoption by an early-adopter practitioner, a framing taken up in the limitations.

Institutional governance context

The System operates within the University's institutional governance framework. The most directly relevant instrument is the University's responsible AI adoption policy, which commits the University to the Australian Government's AI Ethics Principles and establishes a senior committee with institutional oversight. The policy operates within a university-wide AI strategy deployed in partnership with a major technology vendor. During development, the AI applications available under institutional agreements comprised an enterprise productivity assistant with embedded AI capabilities, and an enterprise deployment of a general-purpose generative AI model (Microsoft Copilot). The policy is reinforced by the University's code of conduct, data governance policy, information security policy and acceptable use policy.[1] These policies shape the System's design materially, such that operational deployment is enabled for the AIs available under institutional agreements, and the workflow must accommodate data classification and data-flow expectations. The AI model-agnostic design principle developed in the discussion is, in significant part, a response to this governance environment.

Theoretical grounding

This reflection draws on a body of theory that connects the codification of knowledge to the artefacts through which institutions hold it, and to the labour through which those artefacts are maintained. Polanyi (1966) distinguished tacit knowledge, held in practice and not fully reducible to articulable rules, from codified, transmissible explicit knowledge, positing that practitioners know more than they are able to articulate. Nonaka and Takeuchi (1995) extended the distinction to the organisational level through the Socialisation, Externalisation, Combination and Internalisation (SECI) model of knowledge creation, in which externalisation (i.e., the conversion of tacit into explicit knowledge) operates as the bridge between individual knowing and shared organisational knowledge. The SECI model was subsequently elaborated with attention to the enabling contexts within which conversion occurs (Nonaka et al., 2000). The stakes of externalisation are institutional as much as conceptual. The literature on knowledge loss through staff transition (DeLong, 2004; Hofer-Alfeis, 2008) establishes that knowledge remaining tacit departs with its holder, connecting codification directly to continuity concerns in higher education. This continuity issue can also be read through organisational memory theory, which treats memory as distributed across multiple retention sites, including individuals, organisational structures, procedures, culture, physical locations and external archives (Walsh & Ungson, 1991).

Codification produces artefacts, and artefacts that serve several parties at once raise analytical questions of their own. Star and Griesemer (1989) described boundary objects as artefacts flexible enough to accommodate the local purposes of the parties employing them while stable enough to retain a common identity across sites. They distinguished four types, including *repositories*, defined as "ordered 'piles' of objects which are indexed in a standardized fashion" (p. 410), and *standardized forms*, devised as "methods of common communication across dispersed work groups" (p. 411). Star (2010) later cautioned that we should not only think about boundary objects as flexible across different interpretations. We should also

[1] The institution's policy instruments are described without citation to preserve the anonymity of the case institution.

consider what information such objects must carry for each task, what practical work they support, and how they shift from loosely defined shared forms into more precise, structured tools in specific work contexts. When boundary objects become interconnected within organised systems, standardised and embedded in routine work, they may develop into enduring boundary infrastructures (Bowker & Star, 1999; Star, 2010).

Infrastructure studies (Edwards, 2010; Star, 2010; Star & Ruhleder, 1996) characterise *infrastructures*, including knowledge, digital and informational infrastructures, as deeply embedded in organisational policy frameworks, processes, technologies, norms and relationships. Because infrastructures develop from existing (inherited) technologies, routines and conventions, they carry forward constraints that make rapid, substantial change difficult. They typically support work invisibly until disruption or failure makes them visible. They also tend to change incrementally rather than all at once. Edwards (2010) further argues that robust infrastructures enable knowledge to be produced, shared, maintained and verified over time; without them, claims are more difficult to substantiate, and knowledge is vulnerable to deterioration or loss.

Knowledge and digital infrastructures require continual support, adaptation and coordination. Strauss (1985, 1988) termed the coordination of distributed tasks, efforts and actors *articulation work*, a form of work inherent in any division of labour. Suchman (1995) demonstrated that well-performed service and support work often disappears from view and that its representation shapes its value, with visibility bringing both recognition and exposure. Together, these theoretical traditions trace a progression from tacit knowledge to codified, structured artefacts, from artefacts to infrastructure and from infrastructure to the labour sustaining it. The concepts of articulation work and the visibility and invisibility of work provide a lens for analysing how the System's artefacts mediate across funding authority requirements, internal university compliance-reporting functions and successive role-holders, connect work across AI-enabled platforms and render design labour visible or invisible.

Methodology

Methodological approach

The paper sits within the tradition of reflective practice research (Schön, 1983), examining one practitioner's experience of designing and establishing an AI-augmented System, without advancing generalisable claims. Specifically, it adopts an analytic autoethnographic approach (Anderson, 2006), in which personal experience is used to develop broader understandings of professional practice through systematic, layered self-observation. Anderson (2006) identifies defining features of this approach as complete member researcher status, analytic reflexivity, narrative visibility of the researcher's self, dialogue with informants beyond the self and commitment to theoretical analysis. The researcher holds complete member researcher status, having designed, developed and operated the System as system developer. In this paper, the terms 'system developer' and 'practitioner' are used interchangeably to describe that role. Analytic reflexivity proceeds through documented engagement with project artefacts and prior scholarship, and positionality is rendered visible in this paper. Dialogue beyond the self is partially carried by sustained engagement with colleagues throughout the System's development, although not enacted as formal interview-based fieldwork. The documentary residue of that engagement (i.e., successive artefact versions revised in response to feedback and recorded sign-off points) forms part of the audit trail, and the interactions themselves are not analysed as research data. Theoretical analysis proceeds through the traditions established above.

The single-case focus is deliberate. The System spans the full reporting cycle, from obligation identification through to submission and post-cycle archive, and its development period required design choices across multiple domains by a practitioner whose prior expertise lay outside compliance system design. The distance between existing expertise and the demands of the case makes the trajectory of design development more visible (Freeman, 2026).

Evidence base and analytic method

The evidence base comprises the completed System as established (as at May 2026) together with the audit trail of artefacts retained during its development. The trail includes the multiple folders of the System and their contents, the master timeframes document and the suite of AI-ready Word templates, each with an embedded appendix mapping placeholders to data fields, data types, source hints and example values. It further includes the all-fields data dictionary workbook, the consolidated all-data archive with its Index sheet and the Operational Guide. Chat transcripts from the multi-model AI sessions through which the principal artefacts were developed, together with iterative template versions retained at key milestones, complete the trail and provide a record of design choices. These artefacts are both the empirical material for the reflection and the operational output of the project.

The analytic method involved reflection-on-action (Schön, 1983) on this evidence base. A systematic audit of the artefacts identified five dimensions of the System, each representing a design response to a distinct feature of the inherited reporting environment. Each dimension was mapped across three analytical layers: the inherited reporting context, the AI capabilities deployed and the governance design developed. The audit proceeded abductively, moving between the artefact set and candidate system dimensions. A dimension was retained where an inherited system feature, a deployed AI capability and a corresponding design response could each be traced in the audit trail. Candidate dimensions were combined where their design responses overlapped. The resulting three-column synthesis, presented in the findings, provides the analytical structure for the discussion. While the grid condenses approximately twelve months of recursive design cycles, each row represents the cumulative outcome of multiple iterations. For example, the data-governance dimension was retained because the audit trail linked an inherited reporting problem, namely quantitative data extracted from diverse datasets through a combination of manual and digital methods for report templates with differing fields and requirements, to AI-supported consolidation and codification. The resulting governance response comprised the all-fields data dictionary, source-of-truth principles and consolidated all-data archive. This trace from heterogeneous extraction practices, through AI-mediated structuring, to durable governance artefacts exemplifies the procedure applied across the five dimensions.

AI collaboration and data governance

This paper was produced through multi-model AI collaboration. The models supported structural organisation, theoretical synthesis, source review and critique, while the researcher verified all sources and retained responsibility for the analysis, editorial decisions and scholarly claims. The System was established to operate within the University's information environment using institutionally approved AI applications. Institutional data were handled within that environment. This paper examines the architecture of the design artefacts rather than their contents.

Researcher positionality

The analysis is based on documentary artefacts of the researcher's own design work, including template schemas, placeholder conventions, data dictionary structures and operating principles, together with published institutional policy and prior scholarship. The study did not involve human research participants, the collection of personal data or the analysis of identifiable personal information. Formal ethics committee approval was therefore not required. Researcher positionality (Holmes, 2020; Jackson, 2013), or "the position that [the researcher has] chosen to adopt within a given research study" (Savin-Baden & Major, 2013, p. 71), nonetheless influenced the study's conception, case selection, analytic method and claims. I developed the System while employed by the Faculty and therefore occupy an insider position (Le Gallais, 2008; Schutz, 1976) in relation to the case and its institutional context. After completing its development and handover, I transitioned to another organisation. Proceeding self-reflexively (Bourke, 2014), and cognisant of my stake in the context examined, I have sought to make my position visible and ground the claims in the documented audit trail.

Findings

The findings are presented in Table 1. A note of proportion is warranted, as the third column documents the System as established, in advance of independent deployment.

Table 1

Dimensions of the Report Governance System by analytical layer

Dimension	Inherited reporting context	AI capabilities deployed	Governance design developed
Contract and obligation tracking	Obligations dispersed across approximately fifteen contracts, variations, MOUs and correspondence spanning three jurisdictions. No single consolidated statement of operative requirements. Calendar awareness, submission channels and recipients held by the system developer.	LLMs and source-synthesis tools (i.e., AI applications synthesising across a defined document corpus) consolidating obligations into contract schedules, a priority view and a submission channels table.	A master timeframes and requirements document with multiple jurisdictional schedule tables, each obligation traced to its current source clause. An embedded action tracker as the first stop in any cycle. A consolidated submission channels table.
Report templating and reuse	Each cycle a near-bespoke effort. Variable data interleaved with fixed prose. Structural and language conventions variable and held tacitly. All values, including stable historical values, re-extracted and represented each cycle.	LLMs for template construction from prior submitted reports, placeholder schema design, field definition and annotation, and identification of stable versus changing values.	A suite of AI-ready Word templates with uniquely keyed placeholders for cycle-specific values. An embedded appendix mapping every placeholder to its field, type, source and example, bridging the workflow's two AI surfaces (spreadsheet and LLM). Progressive hardcoding of completed historical values, reducing the active placeholder set each cycle.
Data governance	No consolidated definition of data fields collected. Discrepancies between source sheets resolved ad hoc. Cohort and program data dispersed across files, with no cross-program view for longitudinal or audit use.	LLMs for articulating data field definitions and cross-program annotations, codifying observed conflicts into operating principles and consolidating source files into a single workbook with a manifest.	An all-fields data dictionary workbook (approximately 150 data fields across multiple sections) annotated by source, report, funder and program. Codified source-of-truth principles (i.e., executed agreements authoritative for scope, the master timeframes document for deadlines and channels, candidate-level sheets over conflicting summary tabs, calculated data fields recomputed). A consolidated all-data archive with an Index sheet, updated each report cycle.

Dimension	Inherited reporting context	AI capabilities deployed	Governance design developed
Operational workflow	Reports prepared in ad hoc locations by multiple, changing parties. Carry-forward content embedded in individual and corporate memory. Data extraction and narrative drafting performed in one undifferentiated working process. Operational procedures held tacitly.	LLMs for designing naming and folder conventions within existing institutional infrastructure and for narrative report drafting against templates. AI-in-spreadsheet capabilities for data extraction within source workbooks. Articulation of the workflow into discrete, AI model-agnostic steps.	A per-report folder convention with two internal subfolders (i.e., most recent submitted report and new working draft), operationalising contractual content continuity. A two-surface AI workflow separating structured data extraction (spreadsheet surface) from generative composition (LLM surface), with practitioner verification between them. A thirteen-step AI model-agnostic walkthrough documented in the Operational Guide.
System durability and continuity	High institutional exposure to staff transition. Forward deadline awareness held informally. Artefacts updated reactively. Multi-year retrospective work undertaken from scratch. Previous practice did not deploy AI systematically.	LLMs for projecting forward timelines from contractual milestones, externalising aspects of tacit institutional knowledge into durable explicit artefacts and articulating cycle-aligned maintenance routines.	AI model-agnostic design (i.e., standard Word format, markdown-compatible content, placeholder and appendix conventions readable by any capable model, no platform-specific dependencies). Pre-created working subfolders for multiple pending reports. An internal Operational Guide enabling a successor without prior context to commence work. Documented maintenance routines per folder. A prepared evidence base for retrospective reporting.

The reflection reveals five discrete dimensions: contract and obligation tracking; report templating and reuse; data governance; operational workflow; and system durability and continuity. These dimensions are elaborated across three analytical layers: inherited reporting context; AI capabilities deployed; and governance design developed. Reading across the table, several patterns are visible at a findings level. A consistent pattern of partial externalisation runs through each of the five dimensions. Aspects of institutional knowledge previously held substantially tacitly across multiple staff and dispersed files are rendered explicit through AI-mediated codification (Nonaka & Takeuchi, 1995; Polanyi, 1966) in the report conventions, the data dictionary, the source-of-truth principles and the operational walkthrough. The System also exhibits design under conditions of human and AI capability change. As the project progressed over a development period, several dimensions were redesigned, particularly the placeholder convention and the two-surface workflow. Incorporation of these capabilities only became feasible once AI functions arrived within spreadsheet surfaces, while the AI model-agnostic principle operates as a commitment to operability across platforms. A recurrent two-surface workflow pattern appears across dimensions, with the embedded data dictionary as the bridge, naming data points consistently so that the AI on each surface (i.e., LLM and spreadsheet) can engage without intervening translation.

The System also exhibits dynamics of visibility and invisibility (Star & Ruhleder, 1996; Suchman, 1995), as its outputs are visually indistinguishable (i.e., to management, organisational unit, institution, and funder) from artefacts of the same kind produced manually. The artefacts are visible while the design work behind them largely is not (Suchman, 1995). The organisational components (i.e., the Document Library, the master timeframes document and the folder structure) were in active use by colleagues at the point of handover, while the AI-augmented components (i.e., the templates with embedded dictionaries and the two-surface workflow) had been operated principally by the developer for reporting cycles completed for the period March to May 2026. Deployment subsequent to the system developer's transition to another organisation is beyond the scope of this paper.

Discussion

Partial externalisation of tacit institutional knowledge through AI mediation

The partial externalisation pattern extends organisational accounts of knowledge conversion into AI-mediated conditions. This case presents a distinctive form of organisational externalisation: it is undertaken by a single practitioner occupying complete member researcher status, without an antecedent process of collective socialisation. The knowledge externalised is partly the system developer's own and partly institutional knowledge accessible through the role, held by subject-matter colleagues, embedded in contractual documents and dispersed across operational files. AI mediation also places a structural premium on explicit specification. Reliable and auditable generative AI use requires judgement to be articulated in reviewable form. The system developer must specify requirements sufficiently to inform and subsequently verify generated responses, thereby externalising aspects of tacit institutional knowledge into explicit artefacts that can be reviewed, refined and transferred.

Faraj et al. (2026) describe a related dynamic as dialogical vetting: the recursive questioning, evaluation, reframing and re-prompting through which organisational knowers supply contextual judgement that machine outputs lack. In considering "the entangled dynamics of knowing-with-large language models", they argue that such models blur the distinction between tacit and explicit knowledge (Faraj et al., 2026, p. 343). This blurring raises questions about the knowledge-conversion processes (socialisation, externalisation, combination and internalisation) represented by Nonaka and Takeuchi's (1995) SECI model. The present case moves this dynamic from ad hoc individual human-AI interaction towards institutionalised workflow. Agreed templates, data fields and workflow rules embed aspects of dialogical vetting within the System, reducing reliance on moment-by-moment prompting while retaining human contextual judgement in system design, verification and governance.

The implications extend to scholarship on knowledge codification and knowledge loss through staff transition (DeLong, 2004; Hofer-Alfeis, 2008). In Walsh and Ungson's (1991) terms, the System redistributes organisational memory from the individual system developer towards more durable retention structures. In this case, these include documented procedures, folder architecture, templates, dictionaries and archived source materials. This redistribution extends beyond artefacts. Three structured training sessions were conducted, including with a colleague responsible for structured data extraction and program coordination, supporting shared operation of the System across roles.

Design under disruption and AI model-agnostic design as institutional alignment

The System was progressively redesigned as the surrounding AI capability environment changed. The placeholder convention was revised, and the two-surface workflow became feasible once AI functions were available in spreadsheet services. Infrastructure studies demonstrate that infrastructure changes incrementally and requires ongoing maintenance (Edwards, 2010; Star, 2010). Building AI-augmented workflows and systems therefore requires continuous adaptation rather than single-pass implementation, as specifications evolve with expanding human and AI capabilities. Architectures that preserve adaptability, including formats decoupled from any single AI surface and machine-readable placeholder schemas transferable across surfaces, may prove more durable than premature commitment to a fixed configuration.

The System's purposeful AI model-agnostic design principle is the most consequential of these choices, addressing pressures beyond capability churn. Institutional governance and enterprise arrangements shaped the operating environment without fixing it to a single application. During development, a second enterprise AI application was introduced alongside the existing deployment, and further change remained likely. The System also needed to remain usable by people with varying levels of AI proficiency rather than depend on the developer's accumulated expertise. AI model-agnostic design therefore provides adaptable infrastructure capable of accommodating changes in institutional applications, available capabilities and user expertise.

Institutional compliance settings are often shaped by vendor change, capability change and evolving AI governance. In such settings, AI model-agnostic artefacts designed for machine legibility are more likely to remain operable and adaptable across changes in applications, capabilities and institutional arrangements. The present case supplies a deliberately preliminary data point. Within a single development period, both the available human and AI capabilities and the vendor environment shifted, expanding the range of institutionally supported AI applications available to staff. Accommodating these changes required re-grounding the operational walkthrough and revising aspects of the templates and data dictionaries.

The two-surface AI workflow as a design pattern for compliance work

The AI workflow separates structured data extraction (spreadsheet surface) from generative narrative composition (LLM surface), with the system developer mediating between the two and the embedded data dictionary as the bridge. Read through Star and Griesemer (1989) and Star (2010), the embedded data dictionary operates as a boundary object whose mediating work now spans AI surfaces as well as human parties. The narrative-drafting LLM and the data-extraction capability in the spreadsheet recognise the same data points by name and definition, without intervening translation, while the dictionary remains readable and usable by the human system developer.

The pattern carries potential transferability to other compliance domains, as such work characteristically requires both quantitative extraction from structured sources and qualitative composition against contractual requirements. The two-surface design preserves the system developer's verification role at the most error-prone step, between extracted data and generative composition. The spreadsheet surface also extended the reach of that role. In the present case, data sheets extracted from university systems arrived carrying errors, and AI capabilities within the spreadsheet made it practicable to interrogate them for integrity at the point of receipt, surfacing discrepancies that could otherwise have passed into submitted reports.

Boundary-object theory provides a useful but necessarily limited interpretive lens for this workflow. Star (2010) was explicit that boundary objects arise in cooperative work between communities of practice and that their interpretive flexibility is grounded in human interpretation. The AI-enabled spreadsheet and LLM are not communities of practice and do not interpret in Star's sense. Rather, the data dictionary mediates between the human activity systems of data stewardship and narrative reporting, both of which are AI-augmented. The artefact must therefore remain meaningful to human users while also satisfying machine-legibility requirements across the workflow.

The visibility and invisibility of AI-augmented work

The System's outputs are visually indistinguishable from equivalent artefacts produced without AI augmentation, a finding with consequences that the articulation work tradition helps to interpret. Strauss (1985, 1988) named such coordination labour articulation work, and Suchman (1995) showed that the visibility of work is a design outcome, consequential for recognition and exposure alike. This case suggests that AI augmentation intensifies this dynamic instead of resolving it. The indistinguishability of outputs supports adoption, as artefacts are recognisable and usable regardless of users' AI familiarity, consistent with the System's commitment to operability regardless of high-level AI proficiency. The same indistinguishability obscures the design work, with consequences for institutional recognition (design choices are difficult to perceive in outputs), professional development (what is learned through AI-augmented practice is not visible in what is produced) and knowledge transfer (the artefacts can be inherited as documents and folders, while the working practice that animates them requires distinct support to transfer). Institutions seeking to support sophisticated AI-augmented workflow, rather than simple AI adoption, may accordingly need visibility mechanisms for such work, distinct from those for its outputs.

Implications for higher education management

The case carries implications for higher education management beyond immediate efficiency gains. Dependence on individual staff makes compliance reporting vulnerable to disruption, particularly amid institutional fiscal constraint and insecure employment. The System illustrates one mitigation strategy by making AI-augmented workflows less dependent on individual knowledge. However, locally developed systems also raise questions of management, resourcing, ownership and recognition. To outlast their designers, they require an institutional owner, maintenance resources and an appropriate place within enterprise architecture. Continuing professional development is also necessary to distribute the increasingly sophisticated capabilities on which such innovation depends. The differential adoption pattern suggests that the distribution of AI capability among staff influences what can be transferred. The accountability concerns identified in the limitations further indicate that auditability and data lineage warrant management attention as agentic capabilities mature. Encouraging comparable transformations therefore requires institutional recognition and reward for the often invisible design work through which AI-augmented workflows and infrastructure are developed.

Limitations

Several limitations warrant acknowledgement. This is a single-practitioner account in which design assessments and interpretations are self-reported. That is consistent with analytic autoethnography (Anderson, 2006); however, it does not carry the evidential weight of external evaluation. AI adoption was voluntary rather than crisis-driven, so the case concerns voluntary adoption and individual change agency, not forced implementation.

The final form of the System was tested during the last reporting round completed before the developer's departure. Its succession readiness, however, remains unverified in operation by successor staff. Although the handover included the Operational Guide and three training sessions, the AI-augmented layer had been operated only by its designer. The System therefore reduced dependence on individual knowledge while retaining some dependence on the practice required to animate its codified artefacts. The single-institution design further limits transferability. Benchmarking drew on prior professional experience rather than formal comparison, partly because comparable institutional practices are not publicly documented.

Efficiency gains may also create additional demand rather than reduced workload (Blanco, 2026), although this lies outside the present evidence. A further governance concern arises as AI models gain capacity to modify institutional documentation. Government reporting requires defensible data provenance, yet accountability becomes challenging where lineage is mediated through an AI-augmented workflow. The System's all-data archive and Index sheet provide a basis for reconciliation and future versioning controls; however, the broader tension between delegation and accountability remains unresolved (Blanco, 2026) in this and many other cases.

Conclusion

This paper has examined the design and establishment of an AI-augmented Report Governance System for higher education compliance reporting, mapping five dimensions across three analytical layers and interpreting the resulting patterns. The case documents an infrastructure designed and built to remain operable after its designer's departure and through changes in the platforms on which it runs, in a domain where errors have external consequences and reporting is subject to audit. Its analytical framework, use of boundary-object theory to interpret shared governance artefacts and account of AI model-agnostic design are offered for testing elsewhere.

The argument advanced is one of transferable design logic rather than generalisability. The design logic appears most likely to transfer where the work responds to formal obligations with specified deadlines, content requirements and submission channels, requires both qualitative narrative composition and quantitative data extraction, and deploys multiple AI capabilities across multiple models and programs. Transfer further appears plausible where the work involves multiple people contributing at different times, including across role transitions, and is housed within institutional collaborative workspaces. Under analogous conditions, the design pattern, operating principles and methodological approach are offered as adaptable in other institutions and compliance domains.

Avenues for further inquiry follow from the analysis. Accountability infrastructure for AI-augmented compliance work, including versioning protocols and data lineage, warrants attention as agentic AI capabilities mature. Comparative case studies, using the conditions above for case selection, could consider which elements transfer. The dynamics this case makes visible are offered as a contribution to an emerging conversation about AI-augmented governance systems in higher education compliance practice.

Statements and declarations

Competing interests. The author declares no competing financial or non-financial interests. The author designed and operated the system examined, an insider position disclosed and addressed in the researcher positionality statement.
Funding. No funding was received for conducting this study or for the preparation of this manuscript.

References

Adamakis, M., & Rachiotis, T. (2025). Artificial intelligence in higher education: A state-of-the-art overview of pedagogical integrity, artificial intelligence literacy, and policy integration. *Encyclopedia, 5*(4), 180. https://doi.org/10.3390/encyclopedia5040180

Almatrafi, O., Johri, A., & Lee, H. (2024). A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019-2023). *Computers and Education Open, 6*, 100173. https://doi.org/10.1016/j.caeo.2024.100173

Anderson, L. (2006). Analytic autoethnography. *Journal of Contemporary Ethnography, 35*(4), 373-395. https://doi.org/10.1177/0891241605280449

Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: A systematic review and research agenda. *Information, 16*(4), 296. https://doi.org/10.3390/info16040296

Blanco, C. (2026, May 20). *Will AI deliver freedom from admin or greater mental load?* University World News. https://www.universityworldnews.com/post.php?story=20260519152453263

Bourke, B. (2014). Positionality: Reflecting on the research process. *The Qualitative Report, 19*(33), 1-9. https://doi.org/10.46743/2160-3715/2014.1026

Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

Buame, J. A., Tiika, B. J., & Lotsu, S. A. (2025). The role of AI in shaping the discharge of duties of university administrators: A systematic review. *Pan-African Journal of Education and Social Sciences, 6*(2), 13-31. https://doi.org/10.56893/pajes2025v06i02.02

Creely, E., & Blannin, J. (2025). Creative partnerships with generative AI. Possibilities for education and beyond. *Thinking Skills and Creativity, 56*, 101727. https://doi.org/10.1016/j.tsc.2024.101727

DeLong, D. W. (2004). *Lost knowledge: Confronting the threat of an aging workforce*. Oxford University Press.

Edwards, P. N. (2010). *A vast machine: Computer models, climate data, and the politics of global warming*. MIT Press.

Faraj, S., Perez-Torrents, J., Mantere, S., & Bhardwaj, A. (2026). A time for monsters: Organizational knowing after large language models. *Strategic Organization, 24*(2), 343-356. https://doi.org/10.1177/14761270251410675

Freeman, B. (2026). *I'll be your mirror: Disaggregating technological affordance from developed competence in AI-augmented practice* [Working paper]. SSRN. https://doi.org/10.2139/ssrn.6527560

Hofer-Alfeis, J. (2008). Knowledge management solutions for the leaving expert issue. *Journal of Knowledge Management, 12*(4), 44-54. https://doi.org/10.1108/13673270810884246

Holmes, A. G. D. (2020). Researcher positionality: A consideration of its influence and place in qualitative research. A new researcher guide. *Shanlax International Journal of Education, 8*(4), 1-10. https://doi.org/10.34293/education.v8i4.3232

Jackson, E. (2013). Choosing a methodology: Philosophical underpinning. *Practitioner Research in Higher Education, 7*(1), 49-62.

Le Gallais, T. (2008). Wherever I go there I am: Reflections on reflexivity and the research stance. *Reflective Practice, 9*(2), 145-155. https://doi.org/10.1080/14623940802005475

Liu, X., Guo, B., He, W., & Hu, X. (2025). Effects of generative artificial intelligence on K-12 and higher education students' learning outcomes: A meta-analysis. *Journal of Educational Computing Research, 63*(5), 1249-1291. https://doi.org/10.1177/07356331251329185

Lodge, J. M., & Loble, L. (2026). *Artificial intelligence, cognitive offloading and implications for education*. University of Technology Sydney. https://doi.org/10.71741/4pyxmbnjaq.31302475

Ncube, P. D. N., Dzvapatsva, G. P., Matobobo, C., & Ranga, M. M. (2026). Redefining student assessment in AI-infused learning environments: A systematic review of challenges and strategies for academic integrity. *AI and Ethics, 6*, Article 68. https://doi.org/10.1007/s43681-025-00871-w

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.

Nonaka, I., Toyama, R., & Konno, N. (2000). SECI, ba and leadership: A unified model of dynamic knowledge creation. *Long Range Planning, 33*(1), 5-34. https://doi.org/10.1016/S0024-6301(99)00115-6

Panke, S. (2025). How can (A)I research this? An autoethnographic exploration of generative AI in research, teaching and instructional design. *Journal of Teacher Education, 76*(3), 230-244. https://doi.org/10.1177/00224871251325065

Polanyi, M. (1966). *The tacit dimension*. Doubleday.

Savin-Baden, M., & Major, C. H. (2013). *Qualitative research: The essential guide to theory and practice*. Routledge.

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Schutz, A. (1976). *Collected papers II: Studies in social theory*. Martinus Nijhoff.

Star, S. L. (2010). This is not a boundary object: Reflections on the origin of a concept. *Science, Technology, & Human Values, 35*(5), 601-617. https://doi.org/10.1177/0162243910377624

Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science, 19*(3), 387-420. https://doi.org/10.1177/030631289019003001

Star, S. L., & Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research, 7*(1), 111-134. https://doi.org/10.1287/isre.7.1.111

Strauss, A. (1985). Work and the division of labor. *The Sociological Quarterly, 26*(1), 1-19. https://doi.org/10.1111/j.1533-8525.1985.tb00212.x

Strauss, A. (1988). The articulation of project work: An organizational process. *The Sociological Quarterly, 29*(2), 163-178. https://doi.org/10.1111/j.1533-8525.1988.tb01249.x

Suchman, L. (1995). Making work visible. *Communications of the ACM, 38*(9), 56-64. https://doi.org/10.1145/223248.223263

Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *The Academy of Management Review, 16*(1), 57–91. https://doi.org/10.5465/amr.1991.4278992

Woelert, P., Chesters, J., Martinussen, M., & Gannaway, J. (2026). Administrative burden in Australian universities: Insights into dimensions and drivers from a nationwide survey. *Science and Public Policy, 53*(2), 171–183. https://doi.org/10.1093/scipol/scaf029