

# Contextual Barriers to AI Governance Adaptation in Zambian Private Higher Education

Innocent Kumwenda\*, Joe Likando., Misheck Malama., Happy Siame., Davy Mainde

Africa Research University, Lusaka, Zambia

\*Correspondence Author

DOI: <https://doi.org/10.47772/IJRISS.2026.100800836>

Received: 30 August 2026; Accepted: 04 September 2026; Published: 18 September 2026

## ABSTRACT

Strategic leaders who recognise artificial intelligence's importance do not necessarily have the freedom to act on that recognition. This paper investigates the contextual conditions that constrain strategic leadership adaptation to AI-driven transformation across Zambia's twelve Tier 1 private higher education institutions, using a sequential exploratory mixed-methods design combining a structured questionnaire completed by one hundred governance-role staff with semi-structured interviews conducted with twelve strategic leaders purposively drawn from five institutions. The paper finds that constraint operates from two directions: structural conditions bearing down from above, in the form of financial resource scarcity, infrastructure and connectivity limitations, and regulatory ambiguity, and capability deficits operating from within, in the form of uneven staff digital competence and institutional cultures oriented toward compliance rather than proactive governance. Financial and infrastructural constraint are the most acutely felt of these conditions, but regulatory ambiguity produces a counterintuitive dynamic: rather than creating latitude for institutions to act as they see fit, the absence of a sector-specific mandate produces governance paralysis in institutional cultures built around compliance with external accreditation and regulatory bodies. Financial and infrastructural constraints are compounded by a familiar sustainability failure: externally funded capacity-building erodes once donor support ends, leaving institutions no more resilient than before. Read together, these conditions are interpreted, on the qualitative evidence rather than a tested statistical interaction, as a pincer constraint that would plausibly limit institutional response even among leaders with considerably stronger individual AI literacy than the sector currently possesses; composite-level testing did not find contextual constraint significantly associated with strategic leadership adaptation. The paper's originality lies in documenting, for the first time empirically, how these structural and capability constraints interact and compound within Africa's private higher education sector, with implications for institutional strategy, sector regulation, and national infrastructure and funding policy.

**Keywords:** AI governance; contextual barriers; higher education; resource constraints; regulatory ambiguity

## INTRODUCTION

An institution's leaders can recognise artificial intelligence's strategic importance, want to govern its use responsibly, and still find themselves unable to act, not for lack of will but because the conditions surrounding them make action difficult or impossible. This distinction, between a leadership deficit and a contextual constraint, is easy to collapse in practice but consequential to keep separate: interventions aimed at leaders who are willing but structurally constrained look very different from interventions aimed at leaders who are simply reluctant. Companion research on this same population of Zambian private higher education leaders has established that individual AI literacy is markedly underdeveloped and that formal institutional AI governance is largely absent, and that leaders nonetheless recognise AI's strategic importance nearly universally. That combination raises an obvious further question, one this paper addresses directly: what stands between recognition and action.

The distinction matters most acutely in a private higher education sector, where institutions carry the full financial and regulatory burden of AI governance without the subsidy, national research infrastructure, or

ministry guidance that public universities in many systems can draw on. Private institutions in Zambia compete for the same fee-paying student base while bearing sole responsibility for the pedagogical and ethical consequences of technologies their students and staff have already adopted, a structural position that shapes not only whether governance happens but what kind of governance is even feasible to attempt. A financially secure, well-connected, publicly subsidised university and a fee-dependent private university operating in a connectivity-constrained region face the same nominal governance task, evaluating and overseeing AI use responsibly, under radically different resource conditions, and a body of evidence that does not distinguish between these conditions risks recommending governance responses that only the better-resourced can actually implement. Much of the existing international guidance on AI governance in higher education, produced largely by and for well-resourced systems, implicitly assumes a baseline of infrastructure, funding discretion, and regulatory clarity that this paper's evidence suggests cannot be taken for granted in Zambia's private HEI sector, or plausibly in comparable private higher education systems elsewhere on the continent; adapting that guidance responsibly therefore requires first understanding, empirically, which of its assumptions do and do not hold locally, which is precisely the descriptive task this paper undertakes.

Two bodies of literature bear on this question without fully answering it for this setting. Research on technology adoption in resource-constrained higher education systems has documented financial and infrastructural barriers extensively, typically as a single generic category of resource scarcity. A separate and more recent literature on AI governance specifically has begun to document regulatory ambiguity as a distinct governance challenge, particularly in jurisdictions where a national AI strategy exists in principle but has not been translated into sector-specific operational guidance. Neither literature, however, has examined how financial, infrastructural, regulatory, and cultural constraints interact simultaneously within a single institutional governance context, nor whether they compound each other in ways that a single-factor account would miss.

This paper addresses that gap using evidence gathered for the second objective of a wider doctoral programme examining how strategic leaders across Zambia's private higher education sector are adapting to AI-driven change. It asks: which contextual conditions are associated with constrained strategic leadership adaptation to AI governance across Zambia's Tier 1 private universities, and how do these conditions interact. Addressing this question mixed-methods matters because a purely quantitative account can establish which constraints are most acutely felt across the sector but cannot explain the mechanisms through which a given constraint actually operates inside an institution, while a purely qualitative account, however rich, cannot establish whether a pattern observed in a handful of institutions reflects a sector-wide condition or an idiosyncratic local circumstance.

The paper proceeds as follows. Section 2 describes the mixed-methods design used to investigate contextual constraint. Section 3 presents the findings, moving from the overall pattern of constraint across the six measured conditions to the specific mechanisms of financial scarcity, regulatory ambiguity, competence and culture, and infrastructure readiness in turn. Section 4 concludes by characterising these conditions as a pincer constraint operating from above and within simultaneously, and sets out what this implies for institutional strategy, sector regulation, and national policy.

## **METHODOLOGY**

### **Research Design**

This paper draws on data collected for a wider sequential exploratory mixed-methods investigation into how strategic leaders across Zambia's private higher education sector are adapting to AI-driven change, in which an initial qualitative phase shaped the development of a subsequent quantitative instrument, with the two strands integrated at the interpretation stage.

### **Population and Sampling**

The target population comprised senior academic and administrative staff holding governance responsibility across the twelve Tier 1 private universities registered by the Higher Education Authority of Zambia, an estimated population of 156 individuals. A stratified random sample of 113 was calculated using Slovin's formula at a 95 percent confidence level and drawn across all twelve institutions and five governance role

categories, producing 100 usable questionnaire responses, an 88.5 percent response rate. A purposively selected group of twelve strategic leaders, spread across five of the twelve institutions using maximum variation sampling to capture diversity in ownership type, size, and operating history, took part in semi-structured interviews lasting between 46 and 68 minutes, audio-recorded and later transcribed, with each transcript sent back to its participant to confirm accuracy.

## Measures and Analysis

Contextual constraint was measured using six Likert-scale items: financial resource constraint, regulatory limitation on leadership flexibility, staff digital competence, institutional culture toward technological change, infrastructure and connectivity constraint, and the lack of any AI-specific standard from the sector regulator. These items returned strong internal consistency (Cronbach's  $\alpha = 0.890$ ), exceeding the conventional 0.70 threshold for social science instruments. The corresponding section of the interview guide explored the same conditions in depth, generating the accounts used throughout Section 3 to explain the mechanisms behind the quantitative pattern. Quantitative data were analysed descriptively; the composite-level inferential testing referenced below is reported in Section 2.5, not in the Results, for the reason explained there. Qualitative data were analysed using Braun and Clarke's reflexive six-phase thematic analysis procedure, with independent coding of a transcript subsample by a member of the supervision team to support analytical dependability.

## Methodological Justification

The mixed-methods design was chosen deliberately over either single-method alternative. A quantitative-only design could establish which of the six contextual conditions were most acutely felt across the sector but could not explain why a condition like regulatory ambiguity, on its face a source of institutional freedom rather than constraint, was nonetheless experienced by respondents as limiting; that explanatory work required the interview accounts reported in Section 3. A qualitative-only design, conversely, could generate rich institutional accounts but could not establish whether those accounts reflected a sector-wide condition or the particular circumstances of the five institutions where interviews were conducted; the stratified probability sample underlying the quantitative phase, drawn from all twelve Tier 1 institutions, supports that broader claim in a way five qualitative case accounts alone could not. Where the two strands agreed, quantitative items and independently volunteered interview accounts describing the same underlying condition, that convergence is treated as corroborating evidence rather than as two separate findings, strengthening confidence in the pattern beyond what either strand could establish on its own.

## What the Composite-Level Analysis Does and Does Not Show

Composite-level testing conducted as part of the broader study found no statistically significant bivariate correlation between the Contextual Factors composite and Strategic Leadership Adaptation ( $r = -0.088, p = .385$ ), and no significant regression effect for Contextual Factors once AI Literacy and Institutional AI Response were included alongside it in the same model ( $F(3, 96) = 0.51, p = .680$ ). These null results are reported here, in the Methodology, rather than in Section 3, because they set the interpretive boundary for everything that follows rather than adding a further descriptive finding to it. A sample of 100 respondents supports bivariate and group-level testing but is too blunt an instrument to detect a threshold-style relationship, one where a given constraint only starts to bind once it passes some severity tipping point, of exactly the kind the qualitative accounts in Section 3 describe; formal moderation testing, appropriately powered, would be required to test that relationship directly and is identified as a priority for future research rather than a claim this paper makes. Consequently, the six contextual conditions examined in Section 3 are reported as descriptively associated with constrained adaptation, and their mechanisms are explained through the qualitative accounts; they are not statistically demonstrated causes. Where Sections 3 and 4 describe these conditions as interacting, compounding, or jointly constraining adaptation, including among hypothetically more AI-literate leaders, that description is an interpretation grounded in the qualitative accounts, not a tested statistical interaction or threshold effect. Interview participants are identified throughout by anonymised codes running from P1 through P12, tagged only with a role label; no institution is named alongside any individual account.

## RESULTS

### Constraint From Above, Deficit From Within

The six contextual items divide cleanly into two groups. Four items describe structural conditions external to any individual leader's control, financial scarcity, infrastructure and connectivity limitation, regulatory ambiguity, and the absence of sector-specific AI standards, and all four returned mean scores well above the neutral midpoint of the five-point scale, indicating that respondents experienced these as active constraints rather than background conditions. The remaining two items describe capabilities internal to the institution, staff digital competence and an institutional culture receptive to technological change, and both fell below the neutral midpoint, indicating that respondents did not regard their institutions as well equipped internally to adapt even where external constraint were absent. Financial resource constraint returned the highest mean of the section, closely followed by infrastructure and connectivity constraint; regulatory limitation and the lack of a sector-specific AI standard followed at a similar, only slightly lower level. The overall pattern is one of structural barriers bearing down from outside the institution and capability deficits operating from within it, a genuine pincer rather than a single dominant constraint, and it is this combination, rather than any one condition in isolation, that the remainder of this section unpacks.

The four structural items clustered relatively closely together in mean score, with financial and infrastructural constraint marginally ahead of the two regulation-related items, but the gap between the highest- and lowest-scoring structural item was modest compared with the gap separating the structural cluster as a whole from the two capability items. That clustering is itself informative: it suggests respondents did not experience one overwhelming constraint against a background of otherwise manageable conditions, but rather a broadly uniform pressure across several structural fronts at once, exactly the configuration that makes a single-factor policy response, addressing funding alone, or infrastructure alone, unlikely to unlock adaptation on its own.

It is worth being explicit about what this descriptive pattern does and does not establish. The items reported in Table 1 measure respondents' own perceptions of the conditions constraining their institution, gathered through a single cross-sectional questionnaire; they were not designed to, and are not used here to, demonstrate a causal chain running from any one contextual condition to any specific governance outcome. What the pattern does establish, with reasonable confidence given a stratified sample spanning all twelve Tier 1 institutions and all five governance role categories, is the prevalence and consistency of the perceived constraint itself: agreement on the four structural items was not concentrated among a handful of outlying respondents or institutions but distributed broadly across the sample, a consistency that argues for a sector-wide condition rather than a set of scattered, institution-specific complaints.

**Table 1**

Table 1. Mean scores and reliability for Contextual Factors Associated With Strategic Leadership Adaptation (Section D) items (n = 100)

| Item                                                                              | Mean | SD   | Agreement % |
|-----------------------------------------------------------------------------------|------|------|-------------|
| Financial resources constrain AI adoption                                         | 3.87 | 1.03 | 66%         |
| Infrastructure and connectivity constraints limit AI deployment                   | 3.83 | 1.04 | 63%         |
| Absence of Higher Education Authority AI standards creates governance uncertainty | 3.68 | 1.00 | 57%         |
| Regulatory frameworks limit leadership flexibility                                | 3.66 | 1.08 | 58%         |
| Institutional culture encourages technological change                             | 2.77 | 1.05 | 25%         |

|                                                  |      |      |     |
|--------------------------------------------------|------|------|-----|
| Staff digital competence supports AI integration | 2.61 | 1.08 | 22% |
|--------------------------------------------------|------|------|-----|

*Note. Five-point Likert scale (1 = Strongly Disagree to 5 = Strongly Agree). Section D Cronbach's  $\alpha = 0.890$ . Items ordered by descending mean. Agreement % combines Agree and Strongly Agree responses.*

### Financial Resource Constraints and the Donor-Funding Sustainability Trap

Financial resource constraint was the single most acutely felt condition in the section, and the consistency of agreement across respondents, reflected in a comparatively tight standard deviation, suggests a sector-wide structural condition rather than a circumstance affecting only a handful of institutions. A Vice Chancellor participant laid out what taking AI governance seriously would actually cost: dedicated IT capacity, staff development, outside technical advice, the licences a tool itself requires, and a process for auditing equity outcomes, together a sum that would test even a comfortably resourced university, and for an institution living almost entirely off student fees, with neither government subsidy nor an endowment to fall back on, that sum read less like a line item to be approved than a wish list competing against salaries, facilities upkeep, and the library budget for the same limited pot of money.

A Deputy Vice Chancellor participant identified a further, compounding dimension of the financial constraint: the sustainability failure characteristic of externally funded technology initiatives across the sector. An AI awareness programme for academic staff, well received and well evaluated, had run at his institution on donor funding two years earlier; once that funding ended, the institution had no budget of its own to continue or build on it, and the capacity the programme had created began to atrophy. He was clear that this was not a one-off failure specific to his institution but a pattern he recognised across the sector: capacity built during a funding window, then lost once the funding window closes. This pattern, of donor-funded capacity eroding once external support ends and leaving institutions no more resilient than before the intervention, echoes the wider literature documenting how technology initiatives across African universities tend to fare once outside funding runs out, and the evidence from this study's five participating institutions extends that pattern specifically into AI governance capacity-building. What makes the pattern especially costly in the AI governance context, rather than merely regrettable, is the pace at which the underlying technology continues to change: a training programme that lapses for two years in a stable technology domain leaves an institution roughly where it started, but a governance capability that lapses for two years against a fast-moving technology leaves an institution further behind than before the intervention began, since the tools, vendors, and risks the lapsed training addressed have themselves moved on. Financial constraint in this specific domain therefore compounds over time in a way that a generic resource-scarcity account would understate.

### Regulatory Ambiguity as Governance Paralysis, Not Governance Freedom

A clear majority of respondents said existing regulatory frameworks constrain how much flexibility they have as leaders, a finding that seems counterintuitive on first reading, given that no AI-specific governance requirement for higher education currently exists on Zambia's statute books, and one might expect a regulatory vacuum to widen institutional discretion rather than narrow it. The interview data resolve the apparent contradiction. In institutions whose governance culture has been built through years of engagement with accreditation bodies and regulatory compliance, the absence of an explicit external mandate is not experienced as room to manoeuvre; it is experienced as an absence of the institutional justification that governance action in such cultures typically requires. A Deputy Vice Chancellor participant, from an institution whose own identity was built around technology, described trying to start a formal AI governance process, only to be met with two blocking questions: which regulation demanded it, and against which published standard would the institution be judged; with neither question answerable, the effort was in danger of becoming a personal crusade rather than something the institution had actually committed to. In his account, silence from the regulator did not translate into freedom to manoeuvre; it translated directly into paralysis.

A Dean participant with legal professional training sharpened the point into a concrete governance risk rather than a general institutional discomfort: Zambia's data protection legislation clearly has implications for the data students generate through AI tool use, but translating those general legal obligations into specific guidance for

the AI platforms actually in use at her institution was not something she could responsibly do without regulatory clarification that, at present, does not exist. The law exists in principle; its application to AI in her institution's particular context has not been operationalised by any regulator, leaving her unable to advise governance decisions with the confidence the underlying legal obligation would seem to demand. This account fits with recent Zambian scholarship arguing that the country's biggest AI governance problem is not a missing national strategy on paper, one already exists, but the gap between that strategy's stated commitments and the sector-specific operational guidance institutions would need before they could actually act on it. It is also consistent with broader continental evidence that African higher education systems tend to govern AI reactively, responding to specific incidents as they arise rather than through frameworks developed proactively in advance of governance need; the pattern documented here in Zambia's private HEI sector is not, on this evidence, a local anomaly but an instance of a wider regional condition.

### **Staff Digital Competence, Institutional Culture, and the AI-Specific Skills Gap**

Institutional culture and staff digital competence, the two capability items in Section D, both scored below the neutral midpoint, a pattern broadly consistent with the wider literature on faculty digital readiness in resource-constrained higher education systems. A Dean of School participant offered a distinction that matters directly for how any capacity-building response ought to be designed: it would be wrong to call academic staff digitally incompetent as a group, since most of her colleagues handled technology confidently and well in their core teaching work; where the real shortfall sat was narrower and newer, in knowing how to weave AI tools into teaching without compromising academic standards, how to steer students toward productive rather than integrity-eroding use of those tools, and how to judge whether a given AI tool actually suited a given learning objective. Almost none of her colleagues, she said, had ever received formal preparation in that specific, comparatively recent skill set.

That distinction carries a direct policy implication: a skills gap that is specific to AI pedagogy will not be closed by generic information and communications technology training aimed at digital literacy in general. The institutional culture item pointed to a related, broader condition compounding the competence deficit. A Dean participant, drawing on the same professional background referenced in Section 3.3, identified the compliance orientation typical of professionally accredited academic programmes as a particularly entrenched cultural obstacle to proactive AI governance: in his experience, the first question raised about any significant institutional change was invariably what the relevant professional accrediting body required, a habit of mind shaped by years of external compliance culture; when he tried to raise AI governance as something the institution needed to address proactively, colleagues asked what the accreditor had said about it, an orientation that, given how far AI's pace of change outstrips any accreditor's ability to respond, leaves the institution perpetually behind the very technology it is meant to be governing. Organisational culture formed through sustained engagement with external accreditation and compliance systems tends, on this account, to default toward stability rather than proactive adaptation whenever it meets a disruption the regulatory environment has not yet caught up with, precisely the dynamic these accounts describe.

The competence and culture findings should not be read as a criticism of the staff or leaders concerned. Both accounts quoted above come from participants describing the conditions their colleagues and institutions operate under, not passing judgement on individual capability or motivation; the Dean of School participant was explicit that the academic staff she referred to were competent and effective in their core teaching roles, and the compliance-orientation account describes a rational institutional adaptation to decades of accreditation-driven governance, not an irrational aversion to change. Read this way, the competence and culture findings point toward the same conclusion as the structural findings in Sections 3.2 and 3.3: the constraint is systemic and historically produced, not a matter of individual attitude that motivational messaging or generic encouragement could resolve.

### **Infrastructure Readiness and the Digital Divide Within**

Infrastructure and connectivity constraint returned the second-highest mean in the section, close behind financial constraint, with a clear majority of respondents agreeing that connectivity limitations constrain AI deployment at their institution. A Deputy Vice Chancellor participant gave this finding a concrete institutional history: an

AI-based learning tool piloted at his institution had performed reliably for students with stable internet access but had been slow or entirely unusable for the considerable proportion of students relying on mobile data, in effect creating a digital divide inside the institution through a tool that had been introduced specifically to improve learning. The lesson he took from the experience was that readiness has to be checked before a tool goes live, not left to surface on its own once the students least able to use it start falling behind.

The item on missing regulator standards scored close to the regulatory-ambiguity item covered in Section 3.3, adding breadth to that finding rather than simply repeating it: Section 3.3's interview accounts described institutional cultures that will not commit governance resources without an external mandate behind them, and the survey data show that this particular gap, no published standard from the sector regulator, registers widely across the whole sample, not merely with the handful of participants who happened to raise it unprompted in interview. The infrastructure result also lines up with evidence from outside this study: a recent cross-country assessment built around an international readiness index ranks Zambia toward the back of the sub-Saharan African field specifically on infrastructure, alongside a small cluster of similarly placed neighbours and well behind the region's better-connected leaders. The connectivity constraint documented here is accordingly not an artefact of this study's five participating institutions or of Zambia's private HEI sector considered in isolation, but a sector-level expression of a national infrastructure position that independent external measurement corroborates.

The digital-divide account in this section and the sustainability-failure account in Section 3.2 are, on inspection, two instances of the same underlying dynamic: a well-intentioned AI intervention, whether a training programme or a learning tool, deployed without first verifying that the surrounding structural conditions could actually sustain it, and consequently either lapsing once initial support ended or excluding the very students it was meant to reach. Both point toward the same governance discipline: verification of structural readiness, financial, infrastructural, or otherwise, as a formal precondition for deployment, rather than an afterthought addressed only once a programme or tool has already failed some part of the population it targeted.

## CONCLUSION

This paper set out to identify the contextual conditions constraining strategic leadership adaptation to AI-driven transformation across Zambia's Tier 1 private higher education sector, and to characterise how those conditions interact. The evidence points to a pincer constraint operating simultaneously from two directions: structural conditions imposed from outside the institution's control, financial scarcity, infrastructure and connectivity limitation, and regulatory ambiguity, compounded by capability deficits operating from within it, uneven AI-specific staff competence and an institutional culture oriented toward compliance rather than proactive adaptation. The qualitative accounts suggest, though this study did not statistically test it (Section 2.5), that none of these conditions operates in isolation from the others: financial scarcity plausibly limits what infrastructure and training an institution can afford; infrastructure limitation plausibly undermines the value of whatever training and tools scarce finance does secure; regulatory ambiguity plausibly removes the compliance-driven impetus that might otherwise justify committing scarce finance to governance infrastructure in the first place; and a compliance-oriented institutional culture, formed by years of engagement with external accreditation, plausibly makes that regulatory silence disabling rather than liberating. Read together, and interpreted through the qualitative evidence rather than a tested statistical interaction, these conditions are consistent with a pincer that would constrain governance adaptation even among a hypothetical cohort of leaders substantially more AI-literate than the leaders surveyed for this study actually are; as Section 2.5 explains, composite-level testing did not find contextual constraint significantly associated with strategic leadership adaptation, and this specific claim should accordingly be read as an interpretation awaiting direct statistical test, not a demonstrated finding.

This configuration carries a specific and, in the authors' view, under-appreciated implication: contextual barriers of this kind will not be resolved by leadership development or literacy training alone, however well designed, because the barriers documented here sit outside any individual leader's capacity to remove through personal effort. A companion study of this same leader population found governance-level AI literacy critically underdeveloped and formal institutional AI governance largely absent; the findings reported here suggest, on the qualitative evidence rather than a statistically tested interaction, an important qualification to that companion finding: that even leaders who acquired strong AI literacy tomorrow would still meet financial, infrastructural,

regulatory, and cultural conditions that constrain what they could do with it. On this interpretation, literacy and contextual constraint are better understood as separate governance levers requiring separate interventions than as a single problem literacy training alone could resolve, though this specific claim, like the interaction it rests on, has not itself been statistically demonstrated. Financially, the recurring pattern of donor-funded capacity eroding once external support ends points toward the need for sustainable, recurring funding mechanisms, whether through public-private partnership, sector-wide cost-sharing, or dedicated higher education technology funds, rather than further one-off capacity-building grants that this evidence suggests are unlikely to outlast their funding period. On infrastructure, the digital-divide account documented in Section 3.5 argues for treating verified connectivity and device-access readiness, checked across every channel students actually use to get online, mobile data included, as a formal precondition for any AI deployment decision, rather than a consideration addressed only after a tool has already failed some of the students it was meant to serve. On regulation, the evidence argues against both extremes of waiting passively for a mandate that this sector's own accounts suggest may not arrive in time, and of institutions each independently improvising standards in isolation; the more productive posture documented by at least one participating institution is to treat the regulator as a governance partner, developing draft frameworks calibrated to institutional reality and submitting them for regulatory engagement rather than regulatory instruction. On competence and culture, the finding that the real gap is AI-specific pedagogical competence rather than general digital literacy argues against generic information and communications technology training as an adequate response, in favour of training purpose-built for AI governance and pedagogy specifically.

These four response areas are not independent policy options to be selected among; the qualitative evidence indicates that addressing any one in isolation is likely to yield limited returns, a pincer-configuration interpretation offered descriptively rather than as a statistically tested interaction (Section 2.5). A sustainably funded institution may still struggle to deploy an AI tool responsibly across a student body with uneven connectivity; a well-connected institution may still be unable to govern responsibly if regulatory ambiguity blocks the case for committing governance resources in the first place; and an institution with both adequate funding and infrastructure may still be unable to translate that advantage into responsible practice if its academic staff lack AI-specific pedagogical training and its culture treats proactive governance as someone else's responsibility until an accreditor asks. The practical implication for institutional leaders is to resist treating any single constraint as the constraint, and for Zambia's sector regulator and national policymakers alike, to design interventions, whether standards, funding mechanisms, or infrastructure investment, that anticipate and address this compounding relationship rather than treating each contextual condition as a separate line item to be resolved independently of the others.

These findings should be read alongside several limitations. As explained in Section 2.5, the null bivariate and regression results obtained at composite level should be read as a limitation of statistical power at the available sample size rather than as evidence against the contextual mechanisms the qualitative data describe in detail; appropriately powered moderation testing remains a priority for future research before any causal or even confirmed-moderating claim, including the interaction and threshold effects discussed interpretively in Sections 3 and 4, could be made with confidence. The qualitative interview accounts, while purposively varied by institution type, size, and ownership, were collected at five of the sector's twelve Tier 1 universities, and the specific institutional mechanisms described here should not be assumed, without further verification, to generalise in every particular to the remaining seven. Future research extending the qualitative inquiry sector-wide, and testing the specific moderating and interaction effects among the four structural constraints identified here, would substantially strengthen the evidence base this paper has begun to establish. A further limitation concerns the cross-sectional nature of the design: Zambia's regulatory environment, funding landscape, and infrastructure provision are all plausibly moving targets rather than fixed conditions, and the specific severity of each constraint documented here should be understood as a snapshot of conditions at the time of data collection rather than a permanent feature of the sector.

## **Ethical Considerations**

This study received ethical clearance from the Africa Research University Research Ethics Committee (Certificate Reference ARU/2026/05/005; approved 29 May 2025; expiry 29 May 2026). All data collection took place within the approval period. All respondents and interview participants received a written participant

information sheet describing the study's purpose, the voluntary basis of participation, how data would be used, and the right to withdraw at any time without consequence, and written informed consent was obtained before any data collection began; no coercion or inducement was used in recruitment. Institutional access was obtained through formal written requests to each participating university. Questionnaire responses were collected anonymously with no identifying information recorded. Interview participants were assigned pseudonymous codes (P1 to P12) and are identified throughout only by role category, with no institutional affiliation attached to any individual account, reflecting the small number of senior office-holders per institution and the consequent risk of indirect identification. Audio recordings and transcripts were stored in password-protected files on an encrypted device, accessible only to the research team, with signed consent documentation stored separately from research data. All data will be retained securely for five years following publication, in accordance with Africa Research University's data management policy, and thereafter securely destroyed. Interview questions elicited professional governance perspectives rather than personal information, and no participant was exposed to foreseeable physical, psychological, professional, or reputational risk through participation.

## ACKNOWLEDGEMENTS

The authors thank the strategic leaders and institutional staff across Zambia's Tier 1 private universities who gave their time to complete the questionnaire and participate in interviews for this study, and Africa Research University for institutional support during the research.

## Competing Interests

The authors declare no competing interests. The corresponding author's institutional affiliation, as Principal and Registrar at one of the twelve Tier 1 institutions constituting the study's target population, is disclosed as a matter of transparency; this insider position is addressed directly as part of the study's methodology and reflexivity practice.

## Author Contributions

I. Kumwenda: conceptualisation, methodology, data collection, formal analysis, writing – original draft. J. Likando: supervision, methodology review, writing – review and editing. M. Malama: data curation, validation, writing – review and editing. H. Siame: investigation, validation, writing – review and editing. D. Mainde: resources, writing – review and editing. All authors approved the final version of the manuscript and agree to be accountable for its content.

## Funding

This research received no external or institutional grant funding. Data collection costs were met by the corresponding author.

## Data Availability

The questionnaire and interview data supporting this study are not publicly available due to the confidentiality undertakings given to participants and the institutions involved but are available from the corresponding author on reasonable request, subject to ethical approval and institutional consent.

## Disclaimer

The views expressed in this article are those of the authors and do not necessarily represent the official position of Africa Research University or any of the institutions whose staff participated in this study.

## REFERENCES

1. Bhandari, P. (2023) 'Data collection methods: step-by-step guide', Scribbr. Available at: <https://www.scribbr.com>

2. Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101.
3. Braun, V. and Clarke, V. (2022) *Thematic Analysis: A Practical Guide*. London: Sage Publications.
4. Cohen, L., Manion, L. and Morrison, K. (2018) *Research Methods in Education*. 8th edn. London: Routledge.
5. Creswell, J.W. and Plano Clark, V.L. (2023) *Designing and Conducting Mixed Methods Research*. 4th edn. Thousand Oaks, CA: Sage Publications.
6. Hambulo, F., Simataa, L. and Musonda, F. (2025a) 'Navigating the governance framework for AI in Zambian higher education: a critical examination of the National AI Strategy 2024-2026', *International Journal of Literacy and Education*.
7. Harle, J., Lamptey, R.B., Mwangi, A., Nzegwu, F. and Okere, O. (2021) *Creating Digital Content and Delivering Digital Learning in African Universities: Stories of Innovation from Ghana, Kenya and Nigeria*. Oxford: INASP.
8. Hennessy, S., Jordan, K. and Wagner, D.A. (2021) *Problem Analysis and Focus of EdTech Hub's Work: Technology in Education in Low- and Middle-Income Countries*. EdTech Hub Working Paper. Available at: <https://doi.org/10.53832/edtechhub.0059>
9. Higher Education Authority of Zambia (2024) *The State of Higher Education in Zambia 2023: Quality Assurance in Higher Education*. Lusaka: Higher Education Authority.
10. Higher Education Authority of Zambia (2026a) *Gazette Notice No. 665 of 2026: Classification of Universities*. Government Gazette of the Republic of Zambia, Vol. LXII, No. 46, 17 April 2026. Lusaka: Government Printer.
11. Malatji, M. (2026) 'Bridging the AI divide in sub-Saharan Africa: challenges and opportunities for inclusivity'. Available at: <https://arxiv.org/abs/2601.06145>
12. Mbigi, L. (2005) *The Spirit of African Leadership*. Randburg: Knowledge Resources.
13. Saunders, M., Lewis, P. and Thornhill, A. (2019) *Research Methods for Business Students*. 8th edn. Harlow: Pearson.
14. Sangwa, S., Ngobi, D., Ekosse, E. and Mutabazi, P. (2026) 'AI governance in African higher education: status, challenges, and a future-proof policy framework', *Artificial Intelligence and Education*, 1(1), Article 2054.
15. Schein, E.H. (2010) *Organizational Culture and Leadership*. 4th edn. San Francisco: Jossey-Bass.
16. Zambia, Ministry of Technology and Science (2025) *National Artificial Intelligence Strategy 2025–2027*. Lusaka: Government of the Republic of Zambia.