

Trust, Because You Can't Verify: Privacy and Security Hurdles in Education Technology Acquisition Practices

Easton Kelso
eakelso@asu.edu
Arizona State University
Tempe, Arizona, USA

Ananta Soneji
asoneji@asu.edu
Arizona State University
Tempe, Arizona, USA

Sazzadur Rahaman
sazz@cs.arizona.edu
University of Arizona
Tucson, Arizona, USA

Yan Shoshitaishvili
yans@asu.edu
Arizona State University
Tempe, Arizona, USA

Rakibul Hasan
rakibul.hasan@asu.edu
Arizona State University
Tempe, Arizona, USA

Abstract

The education technology (EdTech) landscape is expanding rapidly in higher education institutes (HEIs). This growth brings enormous complexity. Protecting the extensive data collected by these tools is crucial for HEIs as data breaches and misuses can have dire security and privacy consequences for the data subjects, particularly students, who are often compelled to use these tools. This urges an in-depth understanding of HEI and EdTech vendor dynamics, which is largely understudied.

To address this gap, we conducted a semi-structured interview study with 13 participants who are in EdTech leadership roles at seven HEIs. Our study uncovers the EdTech acquisition process in the HEI context, the consideration of security and privacy issues throughout that process, the pain points of HEI personnel in establishing adequate protection mechanisms in service contracts, and their struggle in holding vendors accountable due to a lack of visibility into their system and power-asymmetry, among other reasons. We discuss certain observations about the status quo and conclude with recommendations for HEIs, researchers, and regulatory bodies to improve the situation.

CCS Concepts

• **Security and privacy** → **Human and societal aspects of security and privacy** ; Privacy protections; Social aspects of security and privacy .

Keywords

education technology, data privacy, security, higher education institutes, acquisition process, vendors, contracts

ACM Reference Format:

Easton Kelso, Ananta Soneji, Sazzadur Rahaman, Yan Shoshitaishvili, and Rakibul Hasan. 2024. Trust, Because You Can't Verify: Privacy and Security

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '24, October 14–18, 2024, Salt Lake City, UT, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0636-3/24/10

<https://doi.org/10.1145/3658644.3690353>

Hurdles in Education Technology Acquisition Practices. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*, October 14–18, 2024, Salt Lake City, UT, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3658644.3690353>

1 Introduction

Digital technologies have become pervasive at higher education institutes (HEIs), integrated into every step of pedagogical and institutional processes[12]. The huge amounts of data these education technologies (EdTech) collect, encompassing academic records, demographics, behavioral data, mobility and campus life, as well as financial, health, and employment records, have made HEIs lucrative targets for cyber attacks: in 2023, data breaches at HEIs have cost an average of 3M USD [37]. Simultaneously, reports on the abuse of collected data by service providers and their affiliates for profiling, tracking, advertising, as well as direct selling of the data to data brokers, are on the rise [2, 11, 36]. Such practices have deep societal consequences, ranging from privacy violations to potential discrimination for employment and other benefits [7, 12, 22, 35]. In the long run, it can also create chilling effects from constant surveillance, undermining students' freedom of expression [20, 46].

Prior research has uncovered privacy and safety issues of EdTech and concerns about those issues from students as well as instructors [8, 16, 17, 39], as they have no choice but to adopt institutionalized technologies [1, 43]. The lack of comprehensiveness in federal data protection laws (such as FERPA [41]) and their existing loopholes [35, 48] can be exploited by vendors to evade accountability [5, 31]. This motivates a need for in-depth scrutiny of the EdTech acquisition practices of HEIs and how they ensure security and privacy during procurement and throughout the life cycle of EdTech, which is largely unexplored. This paper addresses this gap through a semi-structured interview study ($n = 13$) involving people in leadership positions (including Chief Information Security Officers (CISOs), directors, and other senior individuals) at seven HEIs across the US. Specifically, we aim to answer the following research questions:

- (1) **Privacy Posture (RQ1)**: What regulatory frameworks and internal processes do HEIs use to inform their security and privacy policy?
- (2) **Acquisition (RQ2)**: How do HEIs navigate the EdTech acquisition process safeguarding security and privacy? How are data ownership and liabilities negotiated?

- (3) **Post-acquisition (RQ3):** How do HEIs review and perceive data usage, sharing, and accountability concerns throughout the EdTech life cycle?

By exploring the above questions, we provide novel insights into the EdTech acquisition process among HEIs. We highlight our key findings here:

Regulatory leverage: The only leverage that HEIs have over EdTech vendors during the acquisition process is the law- and regulation-backed data policies (§ 5) such as FERPA (Family Educational Rights and Privacy Act [41]) and HIPAA (Health Insurance Portability and Accountability Act [45]). While data breaches pose significant financial risks for HEIs, vendors often get away with only minor reputational hits (§ 5.1). Additional regulation, where it exists (e.g., at the state level), helps HEIs enforce additional security and privacy (and as a bonus, accessibility) requirements to EdTech vendors (§ 5.2). Though regulations help HEIs, ambiguities surrounding data breach liability limit HEIs' ability to hold vendors accountable (§ 7.1).

Security concerns: HEIs have almost no visibility into the security posture and practices of EdTech providers, resigning themselves to rely on questionnaire-based self-reported assessments (e.g., HEC-VAT [10]) to evaluate vendor's risk profile and data management policies (§ 6.3). Few HEIs perform internal audits or penetration tests for the HEI system when EdTech is directly integrated into the system (§ 7.3). Notably, visibility into sub-vendors is even worse, as HEIs lack authority over sub-vendors to undergo assessments, exposing them to potential security and privacy risks downstream (§ 6.4). Some small vendors and individual departmental EdTech acquisitions bypass institution-level assessments altogether, leaving HEIs blind and powerless to how their data is used (§ 6.3).

Contractual oddities: Contracts can only go so far in helping HEIs ensure data ownership and protection. Data mining, especially relating to AI, is an increasing source of uncertainty and concern in HEI contracts with EdTech vendors. While contracts can limit vendors' use of HEI data, the lack of visibility leaves HEIs unaware of data misuse until after the fact (§ 7.1). EdTech's off-boarding process is fraught with uncertainty, as HEIs lack the power and means to verify data deletion by vendors. Thus, even when contracts require the vendors to delete institutional data after their service is discontinued, HEIs cannot verify whether data has been actually deleted; at best they can get a written confirmation (e.g., an email or other document) from the vendors stating that they have deleted all data (§ 7.4). We discuss the implications of these findings and conclude with suggestions to improve the status quo in Section 8.

Contributions. Overall, our contributions are as follows:

- We conduct the first qualitative study with *HEI technology leadership* to uncover the EdTech acquisition process and how institutional security and privacy posture dictates that process.
- We provide insights into the nature of EdTech-related institutional policies of HEIs, including how these policies are enforced and monitored during EdTech acquisition and throughout their lifetime.

- We highlight regulatory and technical limitations that hinder security and privacy compliance, especially in the post-EdTech acquisition phase.

2 Background and Literature Review

Education technologies. We use the Wikipedia definition of EdTech that includes any hardware, software, and educational theory [44]. As such, this definition, when interpreted broadly, encompasses tools that are directly involved in learning activities—such as learning management systems, automated graders, and video conferencing tools—as well as tools that facilitate communications, such as student discussion boards and direct messaging services used in educational settings. That definition also includes algorithms that ship with those tools, such as predictive models (for students' performance or to identify 'at-risk' students [17]), recommending systems for course materials, and automated tutoring agents. For this study, we only consider software tools that HEIs purchase or license from outside vendors and exclude homegrown tools (e.g., research artifacts).

Data collection, breaches, and liabilities. EdTech now manages nearly every aspect of academic activities [12, 16], with tech ecosystems constantly evolving with the availability of plug-and-play third-party 'apps' (such as app marketplaces for Zoom¹ and Canvas²), paving the way for additional data flows. Complexity brings vulnerabilities, as was seen when MOVEit, a popular file transfer tool, was compromised, causing data breaches in almost 900 schools. It exposed students' data including name, IDs, date of birth, contact information, and Social Security number [29]. We lack an understanding of how HEIs conduct security assessments during procurement, how responsibilities for ensuring secure operations are distributed and maintained throughout the life cycle of EdTech, and how liability is negotiated in case of a privacy or security breach.

Privacy in HEI context. Unfortunately, the US does not have any comprehensive data privacy law; the FERPA is the main federal statute that governs the privacy of student data [35]. It requires HEIs to obtain consent before sharing "educational records" unless there is a "legitimate educational interest." However, legal scholar Elana Zeide noted that institutions have almost complete authority to define what is "legitimate educational interest" and decide what data are protected by FERPA and what are not [48]. Russel *et al.* noted that the revision to FERPA in 2008 gives more leverage to private companies by including them as "school officials" and allows disclosure of data to parties such as contractors, consultants, and volunteers; and once data enters the marketplace, it can be freely exchanged since FERPA does not apply to data brokers [35]. Such loopholes allow vendors to use vague statements about what data can be collected, used, and for how long, as Paris *et al.* reported after reviewing publicly available documents at Rutgers University [31]. For example, the contract with Canvas³ says that upon contract termination, the university will lose access to the data collected by Canvas, without mentioning if the data will be deleted or the vendor

¹<https://marketplace.zoom.us/>

²<https://app.learnplatform.com/marketplace>

³<https://www.structure.com/canvas>

may continue to store and use it [31]. This situation also poses legal challenges for HEIs who enroll students from other countries with a stricter privacy policy. As reported in the Findings section, many HEIs we interviewed serve students from the European Union (EU) through online degree programs. Data about these students must be handled following the General Data Protection Regulation (GDPR⁴), which requires a lawful basis and stronger consenting requirements before sharing data with third parties and imposes restrictions on secondary data usage. Thus, there is a critical need for in-depth scrutiny of HEIs' current EdTech acquisition practices—how they leverage federal and state legislature to guide vendor contracts and negotiations to protect data from breaches and misuse, and what challenges they face.

3 Methods

To understand the EdTech acquisition practices within HEIs, we conducted 13 semi-structured interviews [25] with individuals in Educational Technology leadership roles at Universities, including Chief Information Security Officers (CISOs), directors, and other senior individuals.

While our study received exempt status by our Institutional Review Board (IRB), we adhered rigorously to ethical and privacy standards recommended for human subjects studies. This included anonymizing personally identifiable information (PII) such as the participants' and organizations' names, and any other information that could reveal any identity during interview transcription.

Participant recruitment. We recruited EdTech leadership for their pivotal roles in guiding education technology adoption, implementation, and governance at HEIs. Specifically, University leadership provided operational insights, while information security officials offered expertise on security, privacy, and compliance matters. Potential participants were identified through various US institutional websites and the authors' professional contacts. The lead author contacted participants via a recruitment email that outlined study details, compensation, and instructions for interested parties.

We note the significant challenges associated with reaching this target population of leaders at HEIs with specific roles and professional experiences. To address these challenges, we adopted a flexible and iterative recruitment process. We conducted interviews with participants as they became available, continuously inviting new participants. This approach allowed us to engage with a diverse set of participants despite the difficulties in accessing the target population. To expand the participant pool, we also utilized snowball sampling [14], asking initial contacts to recommend additional participants.

In total, we sent 51 invites and interviewed 13 participants from seven US universities between August 2023 and April 2024. All our participants had over three years of experience working with EdTechs, with an average experience of 12 years. Working with hundreds of EdTechs, our participants had first-hand experience dealing with vendors and associated challenges HEIs face in the EdTech landscape. Regarding educational background, four participants had degrees in business administration, and three had graduate degrees related to EdTech. Although we did not ask our

participants directly about their educational background, we noted that information when reaching out to our participants from open sources. Information about the number of EdTechs handled was self-reported during the interview. More participant demographics are detailed in Table 1.

Interview design. Our interview questions explored three aspects of the EdTech life cycle, including privacy policies and practices (regulatory compliance), the acquisition process (including criteria for selection and procurement processes), and post-acquisition challenges. We designed our questions to align with participants' specific roles to gain focused and relevant insights. For example, while all participants discussed privacy policies, individuals from learning enterprises provided detailed perspectives on the EdTech acquisition process, while IT security officials offered comprehensive insights into data management and security post-acquisition. This approach enabled us to capture a comprehensive understanding of various perspectives and responsibilities within the EdTech ecosystem, offering a holistic view of how institutions balance educational goals with security and privacy standards. The interview questions for this study can be found in Appendix A.

Pilot study. The research team collaboratively developed an initial interview questionnaire, which the primary researcher pre-tested through three pilot interviews. All interviewees were EdTech users, including one who regularly conducts workshops and develops tutorials on EdTech use. They offered complementary perspectives, guiding adjustments in our approach and enhancing the quality of the data collected from the study participants [40]. These pilot interviews helped clarify questions, improving the language and interview flow (questions such as 'factors that HEIs look for when acquiring EdTechs' were expanded with more follow-up questions for greater depth), and revealing potential bias (such as maintaining a neutral tone). After multiple refinement rounds, we finalized the questions and conducted the main study.

Data collection. The interviews followed a semi-structured approach, allowing the interviewer to skip or ask follow-up questions as needed, delve into deeper topics, and explore emerging themes. At the outset of each interview, participants provided verbal consent for their participation, with additional consent requested for audio recording purposes. The sequence and emphasis of questions varied based on a) the flow of the conversation and b) the participants' role in EdTech acquisition. During the interview, participants were provided sufficient time to ask clarifying questions and to gather their thoughts or articulate responses, ensuring their perspectives were accurately captured. To minimize interviewer bias, a neutral tone was maintained throughout the process, and the interviewer refrained from expressing personal opinions. Participants were compensated \$50 through an Amazon gift card for a 45-minute online conference call.

Interviews were conducted until data saturation was reached, signifying the point at which no new themes or insights emerge from the data [13], which is a standard in qualitative research. We collected and analyzed data in parallel and observed that saturation was achieved by the 13th interview. This sample size is consistent with qualitative research guidelines, which suggest that saturation

⁴<https://gdpr-info.eu/>

often occurs within 12–20 interviews [6]. More details about the analysis process are provided in the next section.

Data analysis. Anonymized transcripts underwent a thematic analysis procedure [3] using MAXQDA (a qualitative analysis software) [26]. To establish an initial codebook, first, two authors independently analyzed four interviews, compared their initial codes, and collaboratively grouped them into broader categories. The authors used this initial codebook to independently analyze four new transcripts. They updated the codebook with new open codes and created a final codebook with 11 overarching categories by refining already-created categories and resolving any disagreements through discussions [4]. These categories spanned EdTech life cycle at an HEI—from legal guidelines considered during acquisition to factors influencing the acquisition process to security assessments of technologies and post-acquisition challenges. Authors used this final codebook to analyze three transcripts—two new ones and one of the previously coded transcripts—and achieved an inter-coder reliability score for Cohen’s Kappa ($\kappa > 0.8$) [28]. A high inter-rater reliability score, such as this, indicates a consistent coding process, thereby supporting the credibility of the study’s findings [30]. The authors resolved any minor disagreements and continued analyzing the remaining transcripts independently. The final codebook can be found in Appendix A.

Given the qualitative nature of our study, we primarily present our findings qualitatively, occasionally supplementing with counts to highlight prevalent patterns [27]. Additionally, with our semi-structured interviewing approach, it is important to note that any mention of participant counts or IDs associated with specific topics does not necessarily imply exclusivity of those thoughts.

Study limitations. While our study provides valuable insights into EdTech acquisition at HEIs, several limitations warrant consideration. First, due to the semi-structured interview approach, not all follow-up questions were asked to every participant, potentially resulting in variations in data depth and comparability across responses. However, all interviews addressed the key questions necessary to answer our research questions. Second, although we included participants from a range of universities of varying sizes, our recruitment strategy, sample size, and geographic focus may limit the generalizability of our findings. Additionally, generalizability is not typically an “expected” attribute in qualitative research due to the specific focus on particular populations and contexts [24]. To alleviate this concern, we report our findings to highlight both similar and diverse views from our participants wherever applicable. Third, biases such as self-reporting and social desirability could have influenced participants’ responses. We mitigated these risks by framing interview questions neutrally which avoided suggesting any researcher bias and created a comfortable environment for participants to freely share their viewpoints.

Despite these limitations, this study lays a strong foundation for future research on security and privacy in EdTech acquisition within higher education. This study identifies gaps in current practices and aims to promote policy and procedural reforms.

4 Overview of Results

Figure 1 illustrates the EdTech acquisition timeline provided by our participants, key acquisition components and activities, and major findings about each phase in the EdTech acquisition process. Section 5 details HEIs’ privacy posture (RQ1) by examining how federal (§ 5.1), state (§ 5.2), and international regulations (§ 5.3) inform institutional policies of creating governance models (§ 5.4). Section 6 examines EdTech acquisition approaches (RQ2), covering: a) factors considered in acquisitions (§ 6.1), b) the role of contracts for data ownership and liabilities (§ 6.2), c) pre-acquisition security assessments and their challenges (§ 6.3), d) issues with sub-vendors (§ 6.4), and e) the need for collaborative groups among HEIs to share vendor experiences (§ 6.5). Finally, challenges in maintaining, auditing, and discontinuing EdTech (RQ3) are discussed in Section 7. These include issues in determining liability in data breaches (§ 7.1), auditing vendor environments post-acquisition (§ 7.2), refining policies and internal auditing processes (§ 7.3). We note that while we included contextual information (e.g., mentioned pertinent regulations) and condensed data from several participants, all reported findings were solely derived from the interviews.

5 Privacy Posture (RQ1)

HEIs have privacy policies that, at a high level, guide all institutional activities that may require the collection and processing of data from students and staff, as well as the acquisition and use of EdTech from external vendors. While these policies mostly rely on federal privacy regulations, some HEIs also incorporate state laws or form data governance models with additional terms and conditions.

5.1 Inadequacy of Federal Level Policies

Our interviews revealed that most institutional policies are governed largely by domain-specific federal guidelines, such as FERPA [41] and HIPAA [45], for education and health records, respectively. However, most of our participants ($n = 9$), commented that these federal regulations lack adequate protection for sensitive data. For instance, P4 mentioned the need for a comprehensive, rather than sector-specific, privacy law:

“One of the larger challenges right now is that the United States doesn’t have a federal privacy regulation policy. There isn’t a regulatory body that has stepped up into that and there still is debate in Capitol Hill about which regulatory agencies would do that and how they would approach it.”

These existing laws also do not impose sufficiently large penalties on the vendors for violations, as P5 underscored that often vendors may only get a “little reputational hit” in case of data misuse and added that:

“We need better federal laws... it’s a crime that we don’t have better protections for people’s data.”

In response to these challenges, participants emphasized the necessity for advocacy efforts aimed at addressing deficiencies in federal data privacy legislation. Some examples include HEIs coming together to establish standards (P2), provide inputs to the legislation (P3), and eventually, create policies that stand the test of time (P5).

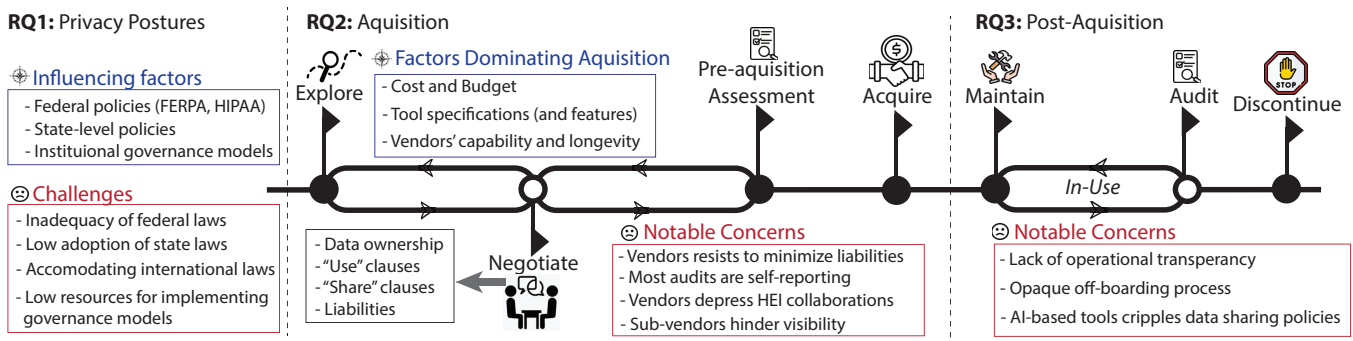


Figure 1: EdTech acquisition and usage timeline learned from our interviews. In light of our final codebook, we also summarize the important findings for each step, both for pre- and post-acquisition phases. Here, ● indicates a mandatory activity, and ○ indicates an optional one.

Takeaway: Inadequate Federal Level Policies

Most HEIs rely on FERPA and HIPAA, meanwhile, vendors may often get away with only a “little reputational hit” in case of data breach or misuse. Participants advocated addressing these deficiencies in federal data privacy regulations.

5.2 Low Adoption of State Legislation

Some states have enacted privacy laws attempting to fill gaps in federal laws [9] (such as the Student Online Personal Information Protection Act [38] in California), but only one institute (out of seven we interviewed) has incorporated state policies into vendor contracts and security assessments. P11 shared how their institution leveraged state-mandated security questions (different from California) for vendor assessment, which is originally designed for state agencies:

“[the state] started looking at all state agencies much more centrally and that created a lot of mandatory questions: what they mean if you answer it no, you’re out. So, you have to answer yes. Or you are deemed non-responsive by the state, and you can’t then sign the contract. So then we took advantage of that to put some of these [for] security questions, both for security and, quite frankly, for accessibility”

P11’s HEI enhanced confidence in data privacy measures by enforcing state-provided five questions—among them, two on security and one on accessibility—during EdTech purchases and noted that “If you [vendor] don’t answer them yes, we just say we can’t do business with you. So the state helped us there, in a way.”

Takeaway: State Level Policies

Most states do not have specific privacy laws, however, the ones that do, allow for their HEIs to leverage laws and achieve higher standards in security.

5.3 Dealing with Intra- and Inter-national Policies

Several participants discussed the challenges of navigating diverse state and international laws, particularly those with campuses spanning multiple states or enrolling international students. These institutions will “follow US law first” (P5), striving to comply with regulations from other countries such as the European Union’s General Data Protection Regulation (GDPR) [42], which imposes huge fines for non-compliance[15]. Although P5’s institution does not have campuses outside of the US, they discussed challenges towards adopting and prioritizing international laws:

“With say GDPR and the Chinese privacy laws (PIPL), you’ve got some tricky things there. I think if there is a conflict, I don’t know how prioritized that would be in terms of meeting the other countries’ regulations”

P4, whose institution offers online courses for international students, added “[HEIs] have to navigate a pretty quickly changing privacy space”. To proactively maintain compliance with the rapidly changing (international) privacy landscape, P4’s institution adopted the NIST privacy framework⁵, invested in automation and tools like ServiceNow⁶, and prioritized privacy protection in their data governance model. We discuss more about governance models in § 5.4.

Takeaway: Intra- and inter-national policies

HEIs struggle to prioritize and comply with multiple state-level and international policies, often giving precedence to US federal law. However, some institutions take GDPR seriously—largely for huge penalties of non-compliance.

5.4 Institutional Data Governance Model

Five participants (P3, P4, P8, P11, and P12) described how their HEIs have developed privacy-focused data governance models, sometimes referred to as General Counsel, to proactively manage EdTech acquisition-related issues. These models review vendor terms for

⁵<https://www.nist.gov/privacy-framework/privacy-framework>

⁶A cloud platform used for efficient collaboration and workflow across campuses, <https://www.servicenow.com/solutions/industry/education.html>

alignment with HEI privacy standards, control access, and curb abusive practices. P4 discussed the workings of their governance model:

“The provost’s office has a governance group that meets monthly [to discuss privacy]. That helps [them] form standards and policy and creates a community of practitioners. [They are] sending folks to Privacy Foundation⁷ specifically to help us build governance models. And then, we spent last year building out a data and privacy governance model for the campus and brought in a really world-renowned expert to help us find how we were going to do data and privacy governance across the campus. A unique model because most people do data governance [only] and they don’t have privacy as a subset and we insisted that privacy was more [peer] to the data governments.”

P3 elaborated on their institution’s internal policies and procedures, which surpass FERPA and HIPAA requirements:

“We have a whole data governance structure in place that governs who can access data, how they can access data, and when they can access data. So, we have a number of policies and procedures in place that really restrict access to only those who have a business need. We audit against that. We review that on a regular basis. So it goes beyond what’s just written within the legislation.”

While it was unclear how “business need” was defined, and if such needs triumph over data subjects’ privacy needs as noted in [5, 31]. In some institutions, internal policies from the governance body often dominate the contracts during EdTech acquisition. One of the participants (P8) shared:

“If a company comes in with a data privacy or data handling agreement that is not aligned with our standards, that will be a stopper for the acquisition”

Takeaway: Institutional data governance policy

Privacy-focused data governance models enable HEIs to set rigorous standards beyond legal regulations, however, not all HEIs have the ability or staff to achieve this. Thus, with varied capabilities, HEIs navigate different challenges in advancing privacy initiatives.

6 Acquisition Phase (RQ2)

This section details EdTech acquisition and how security and privacy issues influence the procurement process. Several participants outlined the intricate processes involved in procuring EdTech solutions, with P6 summarizing them as follows:

“So once we identify that it’s [EdTech is] a solution... we think we want to purchase then they [vendors] provide a proposal that outlines what they are offering and what the costs are. We will negotiate that at a product and cost level to say OK this is just enough of these things, we’re going to need these features, we’re going to need this, whatever. They will put a number against it. Once we have a number that we think is reasonable, we will advance that to procurement. That’s where the procurement process requires an audit and what’s called a purchasing vehicle.”

⁷The Privacy Foundation mentioned here is a research institution that holds seminars for policy-makers, scholars, and other members of the community to educate people on privacy and policy issues faced by HEIs, <https://www.law.du.edu/content/the-privacy-foundation>

This quote reflects broader experiences shared by other participants. For instance, P1, P2, and P3 particularly discussed the necessity of negotiating terms to ensure the inclusion of required features and appropriate pricing.

6.1 Factors Dominating the Acquisition Process

When asked what factors are considered during acquisition, participants emphasized the importance of tool features, robustness, and reliability ($n = 9$), as well as the longevity of a vendor and its business model (P3, P7, P8). For most HEIs, budget is a top priority ($n = 9$), as highlighted by P13:

“The biggest challenge I have is budgets to be honest. You know, they want the top of the line, but they of course didn’t want to spend [the money] on it.”

HEIs also evaluate vendor capabilities, prioritizing those that “enhance the students’ experience” (P1) and “[keep] the learning in mind” (P8), followed by “how stable the product is” (P10). Security considerations typically arise later in the acquisition process. On the other end, vendors often prioritize sales over security, as noted by P5:

“Most of these companies [are] very excited about the product, they’re trying to do good things, but the emphasis is always on sales, and it always feels like security is the afterthought.”

This indicates security is neither a priority by vendors nor by HEIs (at least during the initial stages of acquisition).

Takeaway: Factors impacting EdTech acquisition

HEIs prioritize tool features, vendor capabilities, and longevity, but face budget constraints. Security is a secondary concern. Vendors prioritize sales over security.

6.2 Contractual Safeguards and Negotiations

Contracts serve as crucial tools for HEIs to regulate data collection and usage with vendors. HEIs critically examine vendor contracts, aiming to negotiate *reasonable* data-use and liability terms and create long-term leverage. These negotiations are often time-demanding.

Data ownership. HEIs prioritize maintaining the ownership of the data that EdTechs collect via contracts. According to P2, “it’s always university data” and P5 noted that their “legal department pushes back if vendors claim ownership [of HEI data].” P3 shared a case where a vendor discontinued using student data after “concerns raised in the higher-ed space around privacy” and added that “we’re protected via contracts.” As P9 said, “it all comes down to the agreement.”

All our participants’ institutions ensured that contracts clearly defined data ownership and the scope of data use. Often, contracts outline “confidentiality and privacy” to restrict vendors from “sell[ing] [institutional] data”. P1 stated:

“based on the confidentiality agreement, we would typically not allow them to use our data for any other purposes. It’s not to say that you know somebody like [company name] isn’t mining the data at a high level to enhance their product capabilities and features. But for the most part, we do not sell our student data and records to be shared by other vendors”

Vendors, however, can license institutional data for previously agreed-upon uses and are responsible for its secure handling. To manage data sharing with vendors, P4 mentioned the use of “data transfer agreements” or “data processing agreements.” P6 underscored the importance of contractual clarity in licensing data, cautioning against unfavorable attempts of vendors to claim data ownership through “sloppy agreements.” Thus, the contractual language is a *pillar stone for HEIs and the only safeguard against vendor issues.*

P3 stated that vendors who understand that “data that goes into their systems is the property of [university]” is important. Therefore, according to P3, establishing strong partnerships between HEIs and vendors can be particularly beneficial in the evolving landscape of EdTech. This approach ensures that vendors are not only aware of but also aligned with the needs and expectations of HEIs, facilitating more effective and responsive data management. **Security and privacy terms in contracts.** Participants (P2, P3, P6, and P11) emphasized the importance of vendors agreeing to institutional terms, especially regarding security and privacy. P6 highlighted the process of negotiating terms to ensure compliance with institutional policies:

“Traditionally, we asked the vendor to agree to [the universities] terms and conditions... part of those terms and conditions includes recognition of security and privacy. And to assure that not only are they meeting them in short term, but that they continue to meet them, and that is a contingent part of our agreement going ahead.”

On the other hand, vendors also push for favorable terms and conditions to reduce their liabilities, as P9 discussed:

“If we don’t agree with this term that you [a vendor] have in the contract, we want you to take it out. The vendor says, ‘nope, we need it for our protection, we’re not going to take it out’ then [the university] needs to make the decision: do we keep that term in there or do we pass and not use the piece of software... it’s basically [a] level of risk.”

Lastly, several participants indicated that HEIs struggle to ensure vendors comply with the contracts due to limited insight into vendor operations (§ 7.2). P11 expressed how their office of general counsel was working towards finding out how to “get [checks on vendors] into the contracts, how do we enforce it and how do we check [that they are doing it]?”

Contracts under evolving technological landscape. The constantly evolving nature of EdTech ecosystems presents further challenges, necessitating continual adjustments to policies and contracts to address evolving threats. P11 said:

“The threat surfaces are changing at the speed of light. The bigger the organization, the more you’re playing ‘catch up.’ We are constantly working to adjust our policies and to adjust our contracts to reflect the current scenario”

Rapid advancements in AI technologies also complicate contractual terms for data management and security. AI technologies require a lot of data for training, sometimes accommodated by extending the definition of “legitimate educational interest” [18, 48]. P12 shared their apprehensions:

“So when we enable AI tools which students want to use, faculty want to use, we want to do research on them. How much of that is actually going back and being used? So is it used to train the model? Where is it stored? What do you do with it? Do you share it with anybody? That’s kind of the challenge that I have right now with AI”

Takeaway: Contracts play a critical role!

HEIs rely on contracts to define and regulate data use and access by vendors. Conversely, vendors use contracts to minimize their liabilities. Data mining by EdTech platforms, especially as it relates to emerging AI technologies, is an evolving contractual grey area of concern to HEIs.

6.3 Pre-acquisition Security Assessments

During procurement, vendors undergo security assessments, which can take several forms, such as filling out a questionnaire, most commonly created by a third party like Higher Education Community Vendor Assessment Toolkit (HECVAT) [10]. HECVAT is a 300-question survey that EdTech vendors use to report their security and risk management practices⁸. Some HEIs also prepare their questionnaire. Depending on the EdTech’s complexity and what data it may collect, HEIs might ask vendors to go through *third-party audits* like Service Organization Control Type 2 (SOC 2), but may also be satisfied by vendors’ assessment report. SOC 2 offers HEIs more reliable insights into a vendor’s data and security management practices.

Risk-driven security assessment. Since security assessments are time-intensive and often involve iterative exchanges, several participants highlighted the importance of determining the risk profile of each tool early in the acquisition process. P6 emphasized that their institution’s first step is “to determine the risk profile of a tool” to tailor the intensity of the subsequent security assessment accordingly. P2 explained how different vendors will have different assessments based on the information they are handling:

“It’s different levels of rigor. A relatively small product with no significant data, we’re really focused on assessing security and whether they’re going to protect the integration of the service itself and make sure that it’s not a pivot place for places that hold data. Then there’s a more aggressive one, like we’re going to need to see your plans and get your [assessments] outputs, make sure you send them along with your questionnaire report. If you’re filling out one of the big assessment reports, send those along so we can keep those on record.”

Based on the risk profile assessment, HEIs ask vendors to provide audit reports (from either vendors’ own assessments or third-party audits like SOC2) or fill up questionnaires like HECVAT. P2 asks vendors “for their outputs from their own assessments or an independent party’s assessment of their operation to verify that someone beyond them have said that they’re doing what they’re supposed to.”

Trust in questionnaire-based assessments. While popular, the effectiveness of HECVAT and similar assessments was questioned by several participants: “sometimes vendors answer in very ambiguous ways” (P11), and “honestly, there is probably limited value in the assessment... if they’re going to lie to you, you have no way to detect” (P5).

However, responding to that questionnaire and the response quality can be taken as an indicator of whether a vendor would be a risk to the institution since “not all vendors feel like they would need to respond to a HECVAT” (P12). P12 further explained:

⁸<https://www.ren-isac.net/public-resources/hecvat.html>

“If you [a vendor] answered it in two words for every question, I’m not so sure how serious you were when you filled that out... if a vendor really puts in detailed information, then that gives you a little [bit] of a better feeling that... they have the answers that align to what their product is doing.”

The limitations of questionnaire-based assessments were further underscored by the need for manual review of response reports, a task deemed unsustainable with the growing EdTech ecosystems (*“there’s too many vendors”* P12). As P11 also noted, *“We really don’t have the staff in security to check every single statement that every single vendor makes.”*

Despite these limitations, P12 acknowledged that *“[HECVAT] is useful to gain some trust in that they answered it, but again that’s why I like to look at the SOC 2 and the [penetration] test.”* P2, P3, and P8 similarly take extra measures of penetration tests to make sure that *“they’re [EdTech] secure”* (P3) and *“we aren’t connecting any potential vulnerabilities in our systems”* (P8).

While this multifaceted approach to security assessment reflects institutions’ efforts to mitigate security risks, some EdTech may escape any assessments if they do not go through institutional procurement departments (e.g., being adopted through a less centralized process by the HEI), creating additional vulnerabilities, as P11 disclosed it:

“One of the things that we’ve seen in the last few years is individual departments going out and buying \$40 a month cloud-based applications that never come through the purchasing process. As a result, they’re somewhat invisible to us. ... So we have seen some smaller pieces of cloud software get put in without appropriate oversight.”

Takeaway: Security assessment drivers and trusts

EdTech assessment is limited to a questionnaire (commonly, the HECVAT) and, sometimes, a penetration test. Most participants expressed concerns about the HECVAT’s value, though several suggested that it reflects vendor attitudes and trustworthiness. However, even these limited assessment policies can be bypassed by individual departments that make EdTech acquisitions without oversight.

6.4 Sub-vendors and Associated Concerns

HEIs grapple with additional challenges of visibility and the “unknown” due to the involvement of sub-vendors (vendors’ vendors). This expansion opens the EdTech ecosystem to a larger threat landscape—necessitating HEIs to adapt and adjust their processes accordingly. With more actors playing a part, the risk of unauthorized access to an institution’s data increases drastically, eventually requiring the HEIs to not only consider the risk profile of their vendor but also its sub-vendors. Five of our participants (P2, P3, P5, P11, and P12) found sub-vendor breaches concerning as they might lead to an institution’s data leak.

Risks and challenges of using sub-vendors. P2 expressed that their institution “struggles” with sub-vendor issues, describing the situation as “scary” due to the sheer *“volume of contracts and reviews.”* They compared the complexity to navigating “a spider web.” P2 elaborated:

“There’s a sort of an N plus One problem. You know or like to think about it, if you’re familiar with Minesweeper⁹, you’re going to click on the one file, and it’s going to open up a whole new pathway and it’s going to be different and unknown in each scenario as opposed to being very well understood.”

In some cases where vendors have extensive subcontracting chains—*“vendors who have sub-vendors who have sub-vendors who have sub-vendors”* (P11)—HEIs need to delve deeply to enhance visibility into their main vendor’s sub-vendors. Once vendors (and possibly sub-vendors) send security assessment reports, HEIs’ security team may follow up with further questioning. P3 explains:

“We will sit down with [the vendor’s] security folks [and] go through the HECVAT, get clarity on any kind of questions that they were unable to [answer] on... the HECVAT.”

Additional measures for handling sub-vendors. At the start of an agreement, some HEIs require vendors to disclose their vendors’ information and mandate audits for sub-vendors based on the initial risk assessment. Some institutions, such as P13’s, have dedicated teams working on ensuring that tools from (sub-) vendors that plug into their system are safe by implementing a version of the product on a test bed.

Not all institutions, however, have the resources to thoroughly test (sub-)vendors. P5 expressed that *“smaller universities struggle”* and *“so the process will be as expedient as possible if they’ve got one.”* Notably, all institutions stated that they struggle to keep up with the rising number of audits needed to be performed on vendors and sub-vendors. P11 speculated on potential solutions to the issue:

“I wonder if this kind of chaos is going to lead back to more internally written rather than vendor written software... So you [the university] have control over that path. I don’t know if it will, but I’ve kind of wondered if we’re shooting ourselves in the foot and if we’re going to have a little bit of a swing back to on [premise] management or development to help because it is a big problem.”

Takeaway: Sub-vendor management

HEIs lack visibility into sub-vendor security practices, forcing them to accept HEIs a large threat landscape that they are unable to fully see into.

6.5 Collaboration among HEIs for Better Leverage

The lack of a clear picture of the vendors’ security and privacy posture, exacerbated by reliance on questionnaire-based assessments and limited system access, motivated HEIs to develop collaboration groups, such as Educause¹⁰, Big 10 Academic Alliance¹¹, or the University of California Information Security Council¹². These groups, comprising of HEI security officers (CISOs) and other leadership personnel, facilitate the exchange of notes on security practices and experiences working with different vendors.

HEI collaboration empowers collective influence on vendor practices. While individual institutions may lack leverage, a unified

⁹[https://en.wikipedia.org/wiki/Minesweeper_\(video_game\)](https://en.wikipedia.org/wiki/Minesweeper_(video_game))

¹⁰<https://www.educause.edu/about>

¹¹<https://btaa.org/>

¹²<https://security.ucop.edu/resources/infosec-council.html>

HEI group can effectively advocate for security and privacy priorities, compelling vendors to take notice and prioritize reforms. P2 describes that power by saying:

“The only way that we have found to move [major industry players] is when higher Education as a whole gets together and says you need to be more responsive to our concerns ”

Consultation before finalizing purchase. Despite vendors completing third-party audits, five participants (P1, P2, P3, P4, and P5) expressed their practice of consulting their groups before finalizing a purchase. This collaborative approach facilitates a quick assessment of vendor relationships, their communication practices, and alignment with the needs of HEIs. P5 shared their institution’s practice of “bench-marking” with fellow institutions:

“We had a recent issue with a vendor, one of our schools reached out to us and said... ‘this is the problem we found [with the vendor]’ and we’re like ‘okay, thank you! we’ve got the same problem.’ If I were to acquire something, the first thing we do is [see] what do our other schools [think].”

Information sharing after privacy incidents. P3 described how they use their collaboration with other universities to help inform them about privacy incidents and how to make vendors follow necessary obligations:

“We had an incident in the recent past [with data leakage] that we, let’s just say, leveraged the REN-ISAC that the community is sharing such that others were made aware of some of the issues of this third party. So not only do we leverage the contract and enforce upon the vendor their obligations to notify other institutions that may have been impacted and notify students that may have been impacted, but also leverage that community network that we have to help inform others so that they know the extent of what’s happened.”

Factors impeding HEI collaborations. While most participants were positive about collaborations, P13 stated how collaborations could be hard due to the differences in priority at differing institutions, saying:

“Each university has more of a specific focus on some things, whereas the others don’t. Like a lot of ours is like research and medical... There’s kind of [a] different focus [at] each [university] and so we at least kind of have our own requirements on IT stuff. So yes, it’s different all over the place.”

P5 argues that institutions have limited influence over major vendors and vendors whose primary clientele is outside the higher education sector, and thus, “[big companies are] probably not going to fill one [a security assessment] out for us ” (P5). Moreover, P2 shared that vendors “purposely isolate” HEIs to prevent collaboration and comparison between institutions. They shared that vendors strategically avoid discussing terms negotiated with other institutions and instead focus on showcasing what they can offer “exclusively” to each university. Vendors further ensure that contractual language prevents HEIs from sharing contract details with their peers. Such tactics make it difficult for institutions to negotiate better terms or leverage their collective power.

Takeaway: HEI collaborations and impediments

To help leverage contracts and keep vendor standards, HEIs use collaborations; however, vendors try to isolate HEIs to make it difficult to come together.

7 Post-Acquisition (RQ3):

In this section, we discuss the challenges that institutions perceive and experience as they transition to the post-acquisition phase of integrating EdTech into their university networks.

7.1 Sharing Liabilities during Data Breaches

Privacy incidents, such as data breaches or vendors misusing data, are becoming commonplace in the EdTech space [32, 37]. P9 identified a large number of EdTech tools per institution—ranging from hundreds to upwards of thousands of services—as one of the major causes of privacy issues:

“We deal with so many vendors like probability says someone’s going to get [a leak in data]... it’s a knotted up ecosystem”

Actions and accountability. Regardless of the cause, privacy incidents leave HEIs with little recourse, as P11 stated:

“So, it’s a violation. What’s the remedy? Because how consoled are you if we don’t know till after the fact? Your data is already gone.”

Contracts play a vital role in establishing vendors’ obligations, as P5 noted, “that’s [contract is] the most serious leverage we can have.” Similarly, P1 told us that “if they [vendor] committed to maintaining data privacy, then that’s the contractual obligation, and there’s not much we can do from that standpoint.” In many cases, HEIs make public announcements about the incident, because it is either legally required to give “specific notifications” (P3) based on the data lost or to control reputational damage. P2 stated:

“there’s also just the business part of making sure that our community feels safe and the reputation in the university is known as a place where you can trust to give and provide data and trust that we’re going to take care of it”

Legal implications and ambiguous liability. Privacy incidents can carry significant legal repercussions and hefty fines, particularly if they involve sensitive and protected data (e.g., health or financial data), as discussed by P1:

“With HIPAA data you also have to notify the federal government that there was a HIPAA data loss and you could be fined thousands of dollars per record for that loss.”

Determining liability for privacy-violating incidents can be challenging, particularly when considering whether the university or the affected vendor should be held accountable. This challenge is exacerbated if sub-vendors are involved, given that their contract with the primary vendor and the data flow between them is invisible to the university (§ 6.4). P11 expressed their concern regarding this ambiguity:

“Is it the university [who is liable]? Is it the vendor of the university? Is it the sub-vendor or the sub-sub-vendor? Who is liable? To date, the Department of Education and some very, very sketchy advice says ultimately, the university runs the application, they’re liable.”

While contracts help to some extent in cases of privacy violations, the need for standard guidelines and accountability mechanisms becomes crucial in navigating these situations.

Takeaway: HEI vulnerability to vendor breaches

Vendor breaches pose significant liabilities for HEIs, with contracts often leaving ambiguities and inadequate recourse once data is compromised. Moreover, ambiguities in federal regulations in data breach liability (e.g., as it relates to FERPA) limit HEI's ability to hold vendors accountable.

Takeaway: Post-acquisition trust and obscurity

HEIs struggle to audit vendor environments post-acquisition due to long contract durations, limiting visibility into vendor practices until the next renewal. Trust concerns are raised as HEIs remain unaware of changes in vendor processes or sub-vendors. Thus, HEIs adapt to rely on trust rather than active oversight in an environment of obscurity.

7.2 Visibility and Transparency Concerns

Lack of auditability. The ability to perform or require third-party audits is a large part of the acquisition process, it allows for HEIs to develop a level of trust and transparency with technologies (§ 6.3). However, once the EdTech are acquired, institutions lose their ability to audit those technologies until the next licensing cycle.

HEIs typically maintain contracts ranging from three years (as noted by P1 and P11) to five years (according to P6), aligning with state laws and institutional policies. Consequently, the maximum duration between audits for an EdTech is dictated by the length of the contract. P3 highlighted this gap and described their institution's approach to scrutinizing purchases during the acquisition process.

"We don't have an audit process after the fact. I think that's probably a gap... we do as much scrutiny as we can as a part of that purchase process, working with their technical folks to kind of understand their environment. What processes and controls that they have in place."

Trust issues associated with the obscurity. The inability to audit tools in operation or the vendor's environment invokes potential trust issues, and P5 said that they *"don't particularly have a way to test whether a vendor is abusing our data."*

Participants, including P11, voiced frustrations over vendors' lack of transparency regarding changes, particularly regarding sub-vendors:

"what are you [a university] going to do if they change a sub vendor? Because no vendor on the planet is going to say 'oh we're changing our database vendor. Don't forget to tell [this university]' that ain't gonna happen."

When institutions find out about such changes during contract renewals, products must go through the entire auditing process again. P11 highlighted concerns that HEIs remain unaware of until these contract renewals.

"We continue to refine the processes, but that is one of the product managers responsibility... [Asking] where's that data going? Is the place that it is going, did they get a security evaluation? When was the last time they got a security evaluation? If somebody changes a vendor, until our contract gets renewed and we redo another evaluation, we'd never know that unless the vendor tells us. There's no contractual trigger."

7.3 Refining HEIs' Processes and Practices

Because HEIs cannot audit EdTech in operation, some of them have implemented measures to keep their internal system secure and regularly refine those practices.

Internal audits and penetration tests. To prevent external tools from introducing vulnerabilities into HEIs' internal systems, some institutes conduct annual internal audits and penetration tests, especially with vendors that directly plug into their systems. P3 discussed it:

"I don't know that we [can] do testing from a security and privacy perspective to ensure that the products are not doing something they shouldn't be doing. The only thing that we may do is a penetration test scenario which we do for the most part on internal systems to make sure that they're secure."

HEIs rigorously review their practices as a crucial protective measure. P2 highlighted their institution's periodic audit procedures, noting that top systems are reviewed every two to three years and ancillary systems every three to five years. P2 elaborated on the details of their internal audit process:

"We have an internal policy that does at least two audits on internal systems and services, that does extend [to] the cloud. The audit asks us how we do data security and backup, ask industry standard questions: are we following the right processes, the right controls?"

However, P2 expressed frustrations that such procedures are "unsustainable for checking every vendor that their institution has." To mitigate this, P1 and P3 take proactive steps to strategically reduce the number of technologies employed and opt for solutions that integrate multiple capabilities.

Adapting to evolving technologies. As part of continuously refining audit processes, some HEIs are considering AI tools, while acknowledging associated risks. AI could help HEIs protect their own data and increase surveillance over vendor tools. P11 shared their thoughts on AI-based tools:

"We see an awful lot of new security tools that are AI based...how do you know the algorithm for that AI is good enough to actually work. Well it works 90% of the time, well what about that other ten? Our CIO is concerned that AI will help the bad guys penetrate the network. We also see AI coming on the good guys side within the tools to help monitor millions of rows of log data that our people can't get all the way through."

Overall, HEIs strive to adapt to changing environments and preserve control over their own systems. In the whole *"you have to trust them [vendors] at some point"* (P1), innovative strategies become crucial for HEIs to address post-acquisition challenges and ensure data security.

Takeaway: Refining processes and practices

While not having the ability to confirm that vendors are performing their obligations, HEIs focus on their own practices and look ahead at the growing landscape of emerging tools such as AI that could assist or harm them.

7.4 Off-boarding EdTech

EdTech discontinuation is usually triggered by price increases (P1, P3, P6, and P8), vendor disputes (P5, P6, and P8), or the identification of better tools (P1, P2, and P8). Regardless of the reason, discontinuation raises an important question: What happens to the data collected by the vendor after the discontinuation?

When services are discontinued, vendors typically initiate an off-boarding process outlined in the contract. This process ideally involves erasing all institutional data from their systems and providing a confirmation document. However, vendors may not be as cooperative as one might like; P2 highlighted that:

“Once you tell them [cloud vendors] you’re not buying anything from any of them anymore, they are much less inclined to help you off board ... it might be as little as half of the off boardings, we get the confirmation that — Yes, everything’s shut down and all [the university’s] data is gone”

While vendors may be contractually obliged to provide a “written assurance” (P5) about data erasure, universities need to pursue vendors for that, and that depends on the sensitivity of the records, as P2 explains:

“Unless it’s in the contract, unless we want to pursue long squabbles with a disinterested party, we essentially have to come to a conversation. If we had 200,000 of our alumni with all their social security numbers (SSNs) on it, you can bet we’re going to pursue it. If we have 200,000 people with their addresses and a phone number, we’re likely to pursue it. If we have 1000 people and it’s just their first name, we’re probably going to write it off.”

P12 emphasized that while a signed assurance provides “some legal protection”, it may not fully address concerns if data remains with the vendor despite claims of erasure. P5 shared that due to vendors’ non-cooperation, the off-boarding for complex systems could take multiple years. P3 highlighted a trend towards centralized systems managing various functions from “student information systems to HR to finance.” They expressed concerns if institutions wish to adopt such systems:

“Once you get there, it’s going to be harder to get off and go anywhere else because you’ve got one integrated system. There are advantages to that, but then you’re locking yourself in and it gets harder to break yourself away.”

Takeaway: Off-boarding process

When discontinuing an EdTech, HEIs have no power or ability to investigate if their data has been fully erased from a vendor’s system. HEIs are made to trust a written verification made by the vendor itself, and even that is only sometimes provided.

8 Discussions and Conclusions

This study aims to surface specific security and privacy challenges faced by HEIs that come from the specific tools they use, the resources and processes they have in place, and HEI-specific policies. Thus, while some findings may apply to many HEIs, we do not aim for generalizability. Below, we discuss our results by highlighting key challenges and providing recommendations based on our study that can help strengthen security and privacy within HEIs regarding EdTech acquisition, retention, and discontinuation.

Limitations of FERPA in safeguarding HEI. Overall, our study suggests that HEIs’ privacy posture is largely underlined by the protections offered by FERPA (§ 5.1). This result is consistent with Brown and Klein [5], who found that out of 151 institutional policy documents, 81 were FERPA-related. Given the limited data types that are protected by FERPA [41], the coverage of entities (only federally funded organizations) [41], and the existence of loopholes that can be leveraged to evade compliance altogether [33], the privacy of institutional data is in jeopardy.

Security assessments dilemmas. Although compliance with FERPA and HEIs’ security practices is mandated [41], currently there is no way to guarantee it from the vendors’ end (§ 6.3). HEIs mostly rely on a questionnaire-based assessment of EdTech before procurement (e.g., HECVAT [10]), which does not provide any actual visibility into EdTech and its auditing capabilities. Some HEIs even lack resources to examine those assessment reports in the face of a growing number of vendors (§ 6.4). We note that assessments like HECVAT and SOC2 focus on securing the data from unauthorized access, but they do not cover how this data can be used by “authorized” entities. Alarming, many vendors are unwilling to complete even these basic assessment procedures, highlighting a critical gap in data protection that has significant implications for student privacy and institutional accountability.

Contractual challenges and power asymmetry. Contracts dictate data ownership and use. However, we uncovered that HEIs face challenges to create contracts with concrete and privacy-focused terms and conditions regarding data collection and use (§ 6.2). First, the web of sub-vendors (§ 6.4) exacerbates the issues of ensuring robust data protection practices at the contract level. Second, even when HEIs try to address ambiguity in security assessments and contract documents or revise terms to reflect proper data use and accountability, vendors *often* push back. This reveals a significant power asymmetry between vendors and HEIs, including those among the largest public universities in the US. While some universities form coalitions to improve their negotiating position, these efforts face collaboration challenges and often exclude smaller universities and community colleges. Third, HEIs face inherent difficulties in making concrete statements about data collection since they cannot independently examine the tools, making it incredibly difficult to verify from outside what these tools are doing, what data they are collecting, and where they are sending this data. Consequently, this power imbalance and lack of transparency forces HEIs to trust vendors’ assurances, a practice fraught with risk, as highlighted by our participants.

Data ownership challenges. Having addressed the broader contractual challenges, our study also uncovered grim realities regarding data ownership and use in the context of EdTech acquisitions (§ 6.2). Firstly, HEIs strive to retain ownership of their data, but achieving exclusive control often depends on the institution’s negotiating power. Secondly, contracts may permit vendors to use HEI data without owning it, raising concerns about control and usage rights. Thirdly, vendors can grant similar data usage privileges to sub-vendors, who can change at any time and may have arbitrary contracts that allow them to share or sell data with other parties (including the primary vendor) or data brokers without any restriction. Lastly, vendors might gain sole ownership of data if services are discontinued, as HEIs lack mechanisms to ensure data deletion. The complex relationships of HEIs with vendors and sub-vendors underscore the necessity for further research to unravel the implications of these dynamics on data control and privacy in EdTech acquisitions.

Obscurity of (improper) data utilization. Data also creates new data. For example, behaviors often correlate with demographics and other personal factors, which can be “inferred” from interaction data collected by various EdTech tools [17], allowing vendors to learn more about the students. In addition to deciding the ownership for this newly created data, their use has profound ethical implications [23] and the potential to disproportionately harm vulnerable groups (e.g., outing a non-binary student by “inferring” their gender [17]).

Indeed, one of the primary uses of student data is creating such machine learning-based predictive models [17]. Companies like Instructure, which developed Canvas, offer services based on these models [19]. When such companies are acquired by for-profit equity firms, as seen with Canvas [47], it further intensifies the ethical issues surrounding the digitization of education. The ethical standards of the previous company may not necessarily be upheld under new ownership, and the prospect of misuse or sale of extensive student data raises serious concerns [33, 36]. Educational technologists are increasingly raising alarms about such company takeovers and other issues with the EdTech ecosystem [21, 46].

Risks and liabilities of AI/ML in EdTech landscape. We also learned that while (sub-)vendors profit from data by building and selling new (AI/ML-based) services, they don’t share the associated risks and liabilities (§ 7.3). Moreover, vendors often include terms that waive their liability, as also noted in prior research [31]. Even when liability terms are included in contracts, the complex chain of (sub-)vendors—who can update their privacy policies without HEIs’ knowledge—makes it nearly impossible to assign liability in the event of a privacy violation. Although HEIs may seek protection for the data collected by vendors, there are no safeguards for the ML models trained from that data. Notably, recent research shows that ML models are susceptible to attacks that can reveal the very data used to train them through model inversion or querying [34]. These emerging concerns necessitate a reevaluation and enhancement of regulatory frameworks.

Recommendations:

For regulators. We strongly advocate for comprehensive federal privacy laws to provide HEIs with baseline protection regardless of

their negotiating power with vendors. We also recommend federal or state mandates for standard contracts with default data protection clauses that will apply to all (sub-) vendors to prevent privacy issues caused by complex vendor chains and non-compliant EdTech deployments.

For researchers. Significant research efforts are required to enable transparency, auditability, and accountability in the EdTech ecosystem. Researchers could develop non-intrusive tools to record and analyze data flows for oversight and auditing purposes, especially post-acquisition. Privacy researchers could study techniques to establish private-by-design paradigms, such as differential privacy for aggregated data, minimal purpose-specific data collection, and verifiable data deletion protocols. Legal scholars could scrutinize existing laws to identify policy gaps that vendors exploit and create more comprehensive frameworks benefiting HEIs and student privacy. Researchers in business and management could investigate the impact of company acquisitions on data policies and handling practices, particularly for large vendors like Canvas. Finally, policy researchers could develop methods to ensure vendors comply with data protection standards before and after procurement, enabling HEIs to be confident that vendors are not misusing or mishandling their data throughout the procurement process.

For HEIs. We recommend HEIs increase inter-HEI collaborations to exchange knowledge about best practices and experiences with vendors, which will help strengthen their data security and privacy posture. Forming coalitions will enable HEIs to be more confident in their vendor agreements and enforce better data protection practices. For example, by being upfront and asking for clarifications on vague clauses, supply chains (including third-party vendors, sub-vendors, shadow IT, etc.), and the off-boarding process. Additionally, involving cybersecurity and policy can help HEIs design rigorous vendor assessment processes, negotiate clear data use agreements, and conduct regular audits and compliance checks. HEIs could also adopt data minimization practices, advocate for stronger regulations, and leverage more third-party security assessments.

We end with optimism noting that some HEIs are forming data governance boards, we hope that security and privacy researchers will collaborate to identify appropriate privacy-enhancing tools, as well as research and develop new tools that meet the specific needs of education processes.

Acknowledgments

We thank the anonymous reviewers and shepherd for their valuable feedback for this paper. We are also thankful to the participants for taking time out of their busy schedules to share their insights into the EdTech acquisition process for HEIs. Sazzadur Rahaman was supported by the IT4IR TRIF program from the University of Arizona’s Office of Research, Innovation & Impact (RII). Rakibul Hasan was partially supported by the Air Force Office of Scientific Research (AFOSR) (Award #FA9550-24-1-0227). Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Air Force Office of Scientific Research (AFOSR). This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001124C0362. Any opinions, findings, and conclusions or recommendations expressed

in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA). This material is based upon work supported by the National Science Foundation under Award No. 2247954. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] David G. Balash, Elena Korkes, Miles Grant, Adam J. Aviv, Rahel A. Fainchtein, and Micah Sherr. 2023. Educators' Perspectives of Using (or Not Using) Online Exam Proctoring. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 5091–5108. <https://www.usenix.org/conference/usenixsecurity23/presentation/balash>
- [2] Kyle Barr. 2022. Remote Learning Software Tracked Kids' Data to Sell to Brokers. <https://gizmodo.com/remote-learning-data-brokers-privacy-1848975202>
- [3] Virginia Braun and Victoria Clarke. 2012. *Thematic analysis*. American Psychological Association.
- [4] Virginia Braun and Victoria Clarke. 2022. Conceptual and design thinking for thematic analysis. *Qualitative psychology* 9, 1 (2022), 3.
- [5] Michael Brown and Carrie Klein. 2020. Whose data? Which rights? Whose power? A policy discourse analysis of student privacy policy documents. *The Journal of Higher Education* 91, 7 (2020), 1149–1178. <https://doi.org/10.1080/00221546.2020.1770045> Publisher: Routledge.
- [6] Kelly Caine. 2016. Local standards for sample size at CHI. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 981–992.
- [7] Jake Chanenson, Brandon Sloane, Navaneeth Rajan, Amy Morril, Jason Chee, Danny Yuxing Huang, and Marshini Chetty. 2023. Uncovering Privacy and Security Challenges In K-12 Schools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–28. <https://doi.org/10.1145/3544548.3580777>
- [8] Shaanan Cohney, Ross Teixeira, Anne Kohlbrenner, Arvind Narayanan, Mihir Kshirsagar, Yan Shvartzshnaider, and Madelyn Sanfilippo. 2021. Virtual Classrooms and Real Harms: Remote Learning at U.S. Universities. In *Seventeenth Symposium on Usable Privacy and Security (SOUUPS 2021)*. USENIX Association, 653–674. <https://www.usenix.org/conference/soups2021/presentation/cohney>
- [9] Student Privacy Compass. 2024. State Student Privacy Laws. <https://studentprivacycompass.org/state-laws/>.
- [10] EDUCause. 2024. Higher Education Community Vendor Assessment Toolkit. <https://library.educause.edu/resources/2020/4/higher-education-community-vendor-assessment-toolkit>.
- [11] Todd Feathers. 2022. This private equity firm is amassing companies that collect data on america's children. <https://themarkup.org/machine-learning/2022/01/11/this-private-equity-firm-is-amassing-companies-that-collect-data-on-americas-children>
- [12] Barbara Fedders. 2018. The Constant and Expanding Classroom: Surveillance in K-12 Public Schools. *North Carolina Law Review* 97, 6 (2018), 1673–1726. <https://heinonline.org/HOL/P?h=hein.journals/nclr97&i=1722>
- [13] Patricia I Fusch Ph D and Lawrence R Ness. 2015. Are we there yet? Data saturation in qualitative research. (2015).
- [14] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [15] Leila Halawi and Alpesh Makwana. 2023. The GDPR and UK GDPR and its impact on US academic institutions. *Issues in Information Systems* 24 (2023).
- [16] Rakibul Hasan. 2023. Understanding EdTech's Privacy and Security Issues: Understanding the Perception and Awareness of Education Technologies' Privacy and Security Issues. *Proceedings on Privacy Enhancing Technologies* 2023, 4 (Oct. 2023), 269–286. <https://doi.org/10.56553/popets-2023-0110>
- [17] Rakibul Hasan and Mario Fritz. 2022. Understanding Utility and Privacy of Demographic Data in Education Technology by Causal Analysis and Adversarial-Censoring. *Proceedings on Privacy Enhancing Technologies* 2022, 2 (April 2022), 245–262. <https://doi.org/10.2478/popets-2022-0044>
- [18] Louise Hooper, Sonia Livingstone, and Kruakae Pothong. 2022. Problems with data governance in UK schools: The cases of Google Classroom and ClassDojo. (2022). Publisher: Digital Futures Commission, 5Rights Foundation.
- [19] Instructure. 2024. Data insights for student success. <https://www.instructure.com/highereducation/solutions/data-analytics>.
- [20] Kyle ML Jones. 2019. Learning analytics and higher education: a proposed model for establishing informed consent mechanisms to promote student privacy and autonomy. *International Journal of Educational Technology in Higher Education* 16, 1 (2019), 1–22.
- [21] Kyle M. L. Jones, Alan Rubel, and Ellen LeClerc. 2020. A matter of trust: Higher education institutions as information fiduciaries in an age of educational data mining and learning analytics. *Journal of the Association for Information Science and Technology* 71, 10 (2020), 1227–1241. <https://doi.org/10.1002/asi.24327>
- [22] Mark Klose, Vasvi Desai, Yang Song, and Edward Gehringer. 2020. EDM and privacy: Ethics and legalities of data collection, usage, and storage. *International Educational Data Mining Society* (2020). Publisher: ERIC.
- [23] Kyriaki H Kyritsi, Vasilios Zorkadis, Elias C Stavropoulos, and Vassilios S Vervykios. 2019. The Pursuit of Patterns in Educational Data Mining as a Threat to Student Privacy. *Journal of Interactive Media in Education* 2019, 1 (2019).
- [24] Lawrence Leung. 2015. Validity, reliability, and generalizability in qualitative research. *Journal of family medicine and primary care* 4, 3 (2015), 324–327.
- [25] Natasha Mack, Cynthia Woodsong, Kathleen M MacQueen, and Greg Guest. 2005. *Qualitative research methods*. Family Health International.
- [26] MAXQDA. [n. d.]. MAXQDA. <https://www.maxqda.com/>.
- [27] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
- [28] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica* 22, 3 (2012), 276–282.
- [29] Phil Muncaster. 2023. Almost 900 US Schools Breached Via MOVEit. <https://www.infosecurity-magazine.com/news/us-900-schools-breached-moveit/>.
- [30] Clíodhna O'Connor and Helene Joffe. 2020. Inter-coder reliability in qualitative research: debates and practical guidelines. *International journal of qualitative methods* 19 (2020), 1609406919899220.
- [31] Britt Paris, Rebecca Reynolds, and Catherine McGowan. 2022. Sins of omission: Critical informatics perspectives on privacy in E-Learning systems in higher education. 73, 5 (April 2022), 708–725. <https://doi.org/10.1002/asi.24575> Number of pages: 18 Place: USA Publisher: John Wiley & Sons, Inc. tex.issue_date: May 2022.
- [32] Alex Psilakis. [n. d.]. EdTech Data Breach Emphasizes Need for Bolstered Data Privacy and Security Protections. <https://digitalwatchdog.org/edtech-data-breach-emphasizes-need-for-bolstered-data-privacy-and-security-protections/>.
- [33] Amy Rhoades. 2020. Big Tech Makes Big Data Out of Your Child: The FERPA Loophole EdTech Exploits to Monetize Student DataEdTech Exploits to Monetize Student Data. *American University Business Law Review* 9 (2020).
- [34] Maria Rigaki and Sebastian Garcia. 2023. A survey of privacy attacks in machine learning. *Comput. Surveys* 56, 4 (2023), 1–34.
- [35] N Cameron Russell, Joel R Reidenberg, Elizabeth Martin, and Thomas B Norton. 2018. Transparency and the marketplace for student data. *Va. JL & Tech.* 22 (2018), 107.
- [36] D Samuelson. 2017. Student data being sold to the highest bidder as privacy rights have gone down the drain. <https://www.naturalnews.com/2017-05-12-student-data-being-sold-to-the-highest-bidder-as-privacy-rights-have-gone-down-the-drain.html>
- [37] Natalie Schwartz. 2023. Data breaches cost higher education and training organizations \$3.7M on average in 2023. <https://www.highereddiver.com/news/data-breaches-cost-higher-education-colleges/689499/>.
- [38] Common Sense. 2024. Student Online Personal Information Protection Act (SOPIPA). <https://www.common SenseMedia.org/kids-action/about-us/our-issues/digital-life/sopipa>.
- [39] Kaiwen Sun, Christopher Brooks, Abraham H Mhaidli, Florian Schaub, and Sonakshi Watel. 2018. Taking student data for granted? A multi-stakeholder privacy analysis of a learning analytics system. In *EDM 2018 workshop on policy and educational data mining*.
- [40] Lehana Thabane, Jinhui Ma, Rong Chu, Ji Cheng, Afisi Ismaila, Lorena P Rios, Reid Robson, Marroon Thabane, Lora Giangregorio, and Charles H Goldsmith. 2010. A tutorial on pilot studies: the what, why and how. *BMC medical research methodology* 10 (2010), 1–10.
- [41] US Department of Education. [n. d.]. Family Educational Rights and Privacy Act (FERPA). <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>.
- [42] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10, 1352676 (2017), 10–5555.
- [43] Kristen L. Walker, Kiya Bodendorf, Tina Kiesler, Georgie de Mattos, Mark Rostom, and Amr Elkordy. 2023. Compulsory technology adoption and adaptation in education: A looming student privacy problem. *Journal of Consumer Affairs* 57, 1 (2023), 445–478. <https://doi.org/10.1111/joca.12506> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/joca.12506>
- [44] Wikipedia. [n. d.]. Educational technology. https://en.wikipedia.org/wiki/Educational_technology/.
- [45] Wikipedia. 2024. Health Insurance Portability and Accountability Act. https://en.wikipedia.org/wiki/Health_Insurance_Portability_and_Accountability_Act.
- [46] Ben Williamson, Sian Bayne, and Suellen Shay. 2020. The datafication of teaching in Higher Education: critical issues and perspectives. *Teaching in Higher Education* 25, 4 (May 2020), 351–365. <https://doi.org/10.1080/13562517.2020.1748811> Publisher: Routledge_eprint: <https://doi.org/10.1080/13562517.2020.1748811>
- [47] Jeffrey R Young. 2020. As Instructure Changes Ownership, Academics Worry Whether Student Data Will Be Protected. <https://www.edsurge.com/news/2020-01-17-as-instructure-changes-ownership-academics-worry-whether-student>

data-will-be-protected/.

[48] Elana Zeide. 2015. Student privacy principles for the age of big data: Moving beyond FERPA and FIPPS. *Drexel L. Rev.* 8 (2015), 339.

Appendix

A Extended Version and Supplementary Study Materials

The extended version of this paper can be accessed from <https://arxiv.org/abs/2405.11712>. This version includes our study's

supplementary materials such as interview questions, recruitment email, and study codebook in its appendix. The codebook provides detailed information, including theme names, category names, category descriptions, and open codes.

Table 1: Participant Demographics

ID	Current Job Title	Role in EdTech Acquisition (Years in current Role)	Years of Exp. in EdTech	EdTech Handled	Region
P1	Chief Technology Officer (CTO)	Oversees security review (4)	23	~100	West Coast
P2	Chief Information Officer and VP for Information Technology	Oversees security audits (4)	15	200+	East Coast
P3	Deputy Chief Information Officer	Oversees security audits (2)	23	-	Mid-west
P4	Chief Information Security Officer	Oversees security reviews and audits (3)	3	200+	South
P5	Chief Information Security Officer	Oversees procurement of EdTech handling sensitive data (5)	10	-	Mid-west
P6	Deputy Chief Information Officer	Implements privacy controls in contracts (4)	4	200+	South
P7	Senior Director of Learning	Oversees procurement (7)	7	200+	South
P8	Senior Director of Learning Experience and IT Services	Oversees procurement (4)	4	200+	South
P9	Executive Director of Data Analysis	Handles on premise technology data storage (5)	22	200+	South
P10	Enterprise Partner	Negotiations with vendors (13)	13	200+	South
P11	Deputy Chief Information Officer	Oversees security reviews and purchasing committees (9)	25	~100	Mid-west
P12	Chief Information Security Officer	Oversees security reviews and data handling (2)	12	200+	North-east
P13	IT Support	Handles Edtech post-acquisition (5)	5	~100	South