NASFAA

NATIONAL ASSOCIATION OF STUDENT FINANCIAL AID ADMINISTRATORS

Use of Artificial Intelligence in Financial Aid Offices:

Findings from NASFAA Member Listening Sessions

NASFAA Research Department

Use of AI in the Financial Aid Office Task Force

Table of Contents

Acknowledgements

For this work, NASFAA wishes to thank the members of the Use of AI in the Financial Aid Office Task Force, including:

- Task force chair, Samantha Hicks, Coastal Carolina University
- NASFAA Board of Directors Observer, Celena Tulloss, Auburn University

Task force members:

- Robin Adamson, Coffeyville Community College
- Carly Blagg, Southeast Missouri State University
- Gena Boling, Cornell University
- Crystal Crank, Butler Community College
- Sarah Everitt, Gonzaga University
- Noemi Garcia Tagorda, University of Southern California
- Delia Gomez, Texas A&M University-San Antonio
- Andrew Harvey, Garrett College
- Wayne Kruger, St. Petersburg College
- Jarod Paulson, Capella University
- JoEllen Price, San Jacinto College
- Shara Solero, University of San Diego

We also wish to thank the NASFAA members who participated in these listening sessions. Without your time and feedback, we would not have been able to provide the valuable analysis included in this report to the community.

Definition of Artificial Intelligence

For the purposes of this report, and all work from NASFAA's Use of AI in the Financial Aid Office Task Force, the following definition of AI is used: Artificial Intelligence (AI) is software that simulates human intelligence by interpreting language, recognizing patterns, and generating or analyzing data.

Executive Summary

In late 2025, the National Association of Student Financial Aid Administrators (NASFAA) convened the Task Force on the Use of Artificial Intelligence in Financial Aid Offices to examine how the profession is engaging with Artificial Intelligence (AI) and to develop recommendations for how NASFAA can best support members as this technology evolves. The task force's work is organized in three phases: a review of relevant literature and reports from peer associations; a national member survey that included an institutional policy scan and listening sessions with financial aid professionals; and a final report to the NASFAA Board of Directors with recommendations. This report presents findings from six listening sessions conducted in early March 2026. It is the third report in this task force's series, following Findings from a Survey of Financial Aid Professionals and Findings from a Review of Institutional AI Policies. The task force will use the information gathered across these three reports to inform its final report, which will include recommendations for potential future resources, training, or policy resources to support members and will be submitted to the NASFAA Board. The final report will be published after it has been reviewed and approved by the NASFAA Board in the summer of 2026.

Key Findings

- Most financial aid offices are in early or exploratory stages of AI adoption, with use concentrated in drafting and editing written communications.
- Decision-making about AI use is fragmented; most offices operate without formal processes, and individual staff discretion is the most common model.
- Data privacy is the dominant concern and an active constraint on adoption, with Federal Tax Information (FTI) and other sensitive information in appeals identified as specific gaps not addressed by existing institutional policies. Other personal and ethical concerns around the use of AI were also identified. Most financial aid offices are in early or exploratory stages of AI adoption, with use concentrated in drafting and editing written communications.
- Formal AI training specific to financial aid is nearly absent; where training exists, it is largely self-directed, vendor-provided, or institution-wide and not tailored to the needs of financial aid professionals.
- Participants drew a clear line around professional judgment, compliance functions, and student counseling, identifying these as areas where human judgment must remain central.
- Student-facing AI raises distinct bias concerns, particularly for first-generation, low-income students, while AI-generated student submissions, including SAP appeals and professional judgment requests, are already creating operational challenges for institutions.
- Participants want NASFAA to provide neutral, financial aid-specific resources, including model policy language, training resources, and decision frameworks, but not tool endorsements or advocacy for AI adoption.
- Significant variation in institutional resources, governance structures, and student populations shapes what AI adoption looks like and what support is needed, and NASFAA resources should account for that range.

Detailed Findings

The NASFAA Research Department conducted six listening sessions in March 2026. Sessions were organized around five topic areas based on the task force subgroups: AI Foundations and Types, Data Governance and Privacy, Ethical Principles and Student Impact, Policy-First AI Use, and User Experience and Service Design. Two sessions were convened on AI Foundations. Each session began with a set of common questions on current use, decision-making, and concerns, then shifted to topic-specific discussion. Findings from all six sessions are presented together below, organized first by the themes that emerged across all or most sessions, and then by the additional themes that surfaced within individual topic areas.

Most financial aid offices are in early or exploratory stages of AI adoption, with use concentrated in communication drafting and editing.

Across all six sessions, participants described AI adoption in their financial aid offices as nascent and uneven. The most common use, reported in every session, was the use of generative AI tools to draft, edit, or reformat written communications: student-facing emails, aid offers, policy documents, and internal templates. Participants also described using AI for policy review, data analysis, Excel automation, and knowledge base development, though these uses were less consistent across institutions.

Roughly half of the participants across sessions characterized their office as being at or near the ground floor of AI exploration, still determining what they want the technology to do for them, or having only recently initiated a request through their institution to begin investigating AI options. Several participants reported feeling that their financial aid office lagged behind other campus units, including admissions, institutional research, and marketing. This pattern is consistent with the task force's national survey, which found that 54% of financial aid administrators reported using AI in their work over the prior six months, compared to 94% in a comparable survey of higher education professionals[1].

> ***"Other departments are strong and heavy into AI. As far as the financial aid piece, we're in the beginning stages."***

A smaller group of participants, concentrated in the second AI Foundations session, represented institutions with significantly more advanced adoption profiles: AI-first institutional strategies, enterprise licenses deployed to hundreds of users, AI-related performance goals built into annual evaluations, and active use of AI for audit processes, recruitment screening, and organizational planning. These participants offered useful contrast and described concrete applications that generated interest from others in the room. Their experience reflects what is possible with sustained institutional investment, and also surfaces the downstream questions in areas such as governance, bias, and managing staff resistance that other institutions are likely to encounter as adoption grows.

Student-facing AI tools, primarily chatbots, represent a distinct and more widely deployed category. Several participants described their institution's chatbot as a tool students use unevenly and expressed reservations about its current quality. This distinction between back-end staff use and front-end student-facing deployment is important: many offices that are not yet using generative AI for staff workflows are nonetheless already delivering AI-mediated service to students.

1. The Impact of AI on Work in Higher Education

Decision-making about AI use is fragmented, with most offices operating without formal processes or dedicated guidance.

When asked how their offices decide whether and how to use AI, participants described three distinct models. In some institutions, AI governance flows from the top: an institutional committee reviews and approves tools, enterprise licenses are issued centrally, and use parameters are set at the institutional level. In others, the financial aid director exercises discretion, choosing which tools to use and setting informal expectations for staff. In the most common scenario, however, no formal process exists: staff who are comfortable with AI use it, and those who are not do not, with little coordination or oversight.

> ***"There is no structure for using AI. It's open to whoever wants to use it and how they want to use it, as long as they fit within the large umbrella acceptable use."***

Participants described this state with a mix of pragmatism and unease. Some viewed the absence of formal structure as appropriate given how new the technology is and how rapidly it is changing. Others expressed concern about staff using tools inconsistently or without understanding which data should and should not be entered into those tools. A recurring scenario across multiple sessions was the possibility of a well-meaning staff member inadvertently entering student records or other protected information into a public AI tool because no one had told them not to.

Shadow adoption, the use of AI tools not officially approved or provisioned by an institution, was present but described as less prevalent in financial aid than in other higher education contexts. The task force's national survey found that 37% of respondents reported using institutionally unapproved AI tools, compared to 56% in a survey of higher education professionals in other offices. Listening session participants generally found this gap plausible and attributed it to the compliance orientation of financial aid culture: administrators accustomed to working within strict regulatory frameworks tend to be more cautious about adopting unvetted tools.

Data privacy is the primary concern and a constraint on adoption.

In every session, data privacy emerged as the most prominent concern participants associated with AI use in financial aid. The concern was not abstract: participants described specific fears about accidentally including student names, identifiers, or other protected information in AI prompts; uncertainty about where that information is stored and who can access it; and doubts about whether vendor data protection practices are adequate or transparent.

> ***"I just don't want to accidentally release data... I don't want to accidentally pop a student's information out there."***

Many participants described a practical workaround that has emerged organically across institutions: using de-identified or synthetic data when using AI for tasks that would otherwise require student records. Multiple participants across different sessions independently described building data dashboards or testing AI prompts with fictional student data. This approach reduces risk but also significantly limits what AI can do in financial aid contexts.

Enterprise AI licenses, which are platforms that contain institutional data within a closed environment and do not support external model training, were consistently described as a critical safeguard. Participants whose institutions had deployed Microsoft Copilot or comparable enterprise tools described a meaningful distinction between those tools and public-facing AI platforms, and characterized the enterprise environment as the line between appropriate and inappropriate data use. Several noted that, even within enterprise tools, staff need resources on which categories of information are appropriate to include in AI interactions.

2. The Impact of AI on Work in Higher Education

FTI was identified in the Data Governance and Policy-First sessions as a specific concern that remains underaddressed. Unlike general student records protected under FERPA, FTI is protected by distinct federal statutory protections under the Internal Revenue Code. Participants noted that no institutional AI policy they were aware of addressed how to handle FTI in AI interactions, and that staff may not understand the distinction. One participant raised the specific scenario of students bringing AI transcription tools into virtual counseling sessions, in which FTI is discussed, a gap not addressed by any current institutional policy reviewed in the task force's policy scan.

Participants in the Data Governance session also raised the issue of health information appearing in SAP appeals and professional judgment documentation as a category requiring heightened attention. Financial aid offices routinely receive documentation that includes medical diagnoses, mental health histories, and other sensitive personal health information. This layer of protected information exists within the broader category of student records but may carry additional sensitivity, and its presence in AI-assisted workflows has not been addressed in any institutional policy the task force reviewed.

Survey Connection

The task force's national survey found that 79% of respondents cited data privacy as a top concern about AI use in financial aid, and 67% cited FERPA compliance. Participants in the listening session confirmed both concerns and identified FTI and health information in appeals as additional dimensions not captured in the survey's closed-ended options.

Formal AI training is nearly absent, and the training that does exist is not specific to financial aid.

Across all six sessions, participants reported that formal AI training was rare or nonexistent. The most common experience was that AI tools arrived with little or no accompanying resources. Where training did occur, it typically took one of three forms. Some staff engaged in self-directed learning through online platforms, vendor-provided onboarding, or institution-wide awareness sessions, but none of these addressed the specific questions financial aid administrators face.

We weren't given any training ... we had to figure it out ourselves. It was just, 'Oh, here it is, by the way, and it's recording your meetings now.

A few participants described more intentional approaches. One described a voluntary monthly AI coffee chat that generated ongoing engagement and a shared team channel for resources and discussion. Another described an institutional unit offering structured literacy courses available to both staff and students. These models were seen by other participants as worth emulating, but were acknowledged as requiring staff capacity and institutional infrastructure that many offices do not have.

Participants were specific about what they wanted from the training. Across sessions, the most frequently requested topics were: responsible and ethical use; prompt writing and iterative prompting; data privacy guardrails specific to financial aid; use cases and workflows relevant to financial aid administration; and how to evaluate and verify AI output. Participants at different career stages and role levels described different training needs: directors navigating institutional governance and vendor evaluation need different preparation than frontline staff learning to use tools appropriately in daily work.

Several participants emphasized the importance of training that is impartial and not driven primarily by institutional liability concerns, but grounded in the realities of financial aid practice and developed by an organization with professional credibility in the field.

There is a lack of impartial training. We need some kind of third option here.

Financial aid administrators draw a clear line around professional judgment, compliance functions, and student counseling.

When asked where human judgment must remain central, participants across all six sessions converged on a consistent set of boundaries. Professional judgment determinations, such as dependency overrides, special circumstance adjustments, and satisfactory academic progress (SAP) appeals, were uniformly identified as areas where AI should not make or substantially drive decisions. Compliance functions, including regulatory interpretation, policy application, and verification, were similarly protected. Participants expressed concern that staff without sufficient subject-matter expertise might use AI to answer questions they should answer themselves and that, over time, reliance on AI for compliance guidance could erode the professional knowledge base the field depends on.

Do not ask it questions that are replacing what you are hired to do. You can use it to finesse language or develop training materials, but not to replace subject-matter expertise.

Student counseling was a particular area of concern across multiple sessions. Participants described financial aid counseling as inherently relational and contextual, requiring hearing the question behind the question, recognizing when a student's situation calls for slowing down rather than providing a quick answer, and being willing to deliver difficult news rather than offering false reassurance. Several participants noted that AI tools are designed to be helpful and agreeable in ways that may be antithetical to these functions: a tool optimized to give encouraging responses is not well-suited to a conversation where the honest answer is discouraging.

The art and science of financial aid

A framing introduced organically by a participant in the first AI Foundations session resonated widely across sessions: the distinction between the art and the science of financial aid work. The science - calculations, recalculations, proration, return of Title IV funds, verification workflows, document processing, and form completion was seen as more amenable to AI assistance. The art - student advocacy, professional judgment, interpreting regulations in ambiguous circumstances, and recognizing what a student actually needs in a given moment was deemed to require professional judgment that AI cannot replicate.

A lot of the science pieces of financial aid could be, if not replaced, at least supported by AI, whereas the art piece — how we talk to students, hearing the unasked questions — that can't be replaced.

This framing was not offered as a rigid rule but as a working intuition that participants found useful. Several noted that even science-side tasks require human review; for example, AI assistance with an R2T4 calculation does not eliminate the need for a staff member to verify the result. And the line between art and science is not always clear: drafting a communication to a student in financial difficulty has both technical and relational dimensions. But the framing offered a practical starting point for thinking about where AI adds value and where it introduces risk, and it may be useful as an organizing concept in guidance for financial aid offices.

Student-facing AI raises distinct bias concerns, and AI-generated student submissions are already an operational challenge.

The task force designed this line of inquiry with algorithmic bias in mind; specifically, the risk that AI systems trained on historical data can encode and reproduce existing patterns of disadvantage, producing outcomes that systematically harm particular student populations. In financial aid contexts, where AI tools may influence communications, eligibility determinations, fraud screening, or appeals processing, this risk is not theoretical: decisions shaped by biased models can affect students' access to aid, their ability to resolve holds or errors, and their overall experience navigating institutions.

When asked about AI bias, however, participants often interpreted the question as referring to students using AI in ways that felt inappropriate, such as: generating appeals, researching eligibility, or submitting AI-generated documentation, rather than to bias embedded in the AI tools institutions deploy. This pattern likely reflects the reality that most financial aid offices are currently at an early stage of AI adoption, focused primarily on staff-facing generative AI tools rather than algorithmic decision-support systems where bias risks are most acute. The one concrete example of institutional algorithmic bias that surfaced — a fraud detection tool whose outputs disproportionately affected undocumented students, generating automatic holds that the affected students could not easily resolve — came from a participant describing a past experience rather than a current practice under active review. That this risk did not surface organically in practitioner discussion, despite being a central design question for the task force, suggests the profession is at an early stage of engagement with algorithmic bias as a distinct category of concern.

Participants across multiple sessions expressed concern about the impact of AI on students, particularly first-generation, low-income, and other underserved students. The core concern was that students who are less familiar with higher education and financial aid processes may not know how to ask the right questions of an AI tool, may receive inaccurate information that they lack the background knowledge to identify and question, and may become discouraged and disengage from the financial aid process before ever contacting their institution. Participants noted that these dynamics exist independent of AI, but felt that AI has the potential to scale and accelerate them.

First-generation students... unfamiliar with higher ed and maybe bureaucracy in general, getting discouraged early and dropping out of the funnel. That's my biggest area of concern.

This concern is not hypothetical. Participants described real cases in which students received incorrect information from AI tools about dependency status, financial aid eligibility, and the cost of attendance, and in some cases, could not be persuaded that the AI had provided incorrect information. One participant described students who had researched their options using AI and arrived at counseling appointments prepared to argue based on what an AI had told them, regardless of the regulatory reality.

AI-generated student submissions

A related concern, distinct from AI-generated submissions, involved students using AI to research and navigate the financial aid process. Participants acknowledged this is a natural form of student self-advocacy, not unlike consulting a knowledgeable friend or doing independent research. The concern is less about intent than about accuracy: AI advice in this domain is often wrong, and students who act on incorrect guidance about dependency status or professional judgment criteria may find themselves in difficult positions. Participants also noted, more cautiously, that AI could in some cases lead students to present their circumstances in ways that don't fully reflect their actual situation — a dynamic with potential implications for program integrity, though one that is difficult to distinguish from ordinary advocacy.

Students using AI to complete their SAP appeals... it's very clear they are using AI for their appeals.

Chatbots and the limits of current student-facing AI

Participants described widespread student frustration with chatbot tools. The core complaint was consistent: chatbots answer the question that is asked, but cannot identify the question that should have been asked. A student checking their aid amount may not know to ask whether they are approaching a SAP warning. A student asking about study abroad funding may not know to ask how dropping below full-time enrollment will affect their eligibility. Financial aid professionals understand these nuances, and current AI tools do not replicate that skill.

If you've worked in financial aid, it's the question that's not asked. They don't know what they need to ask. We do, just because we've been in it so long.

Several participants described their institutions as actively transitioning away from first-generation chatbot tools toward more capable vendors, citing student frustration and the reputational risk posed by chatbots that give incomplete or incorrect answers. The bar students now set for these tools is high; they expect authenticated data access, accurate follow-up responses, and an experience that compares favorably with other digital services they use. Current tools, in most participants' assessments, do not yet meet that bar.

Participants want NASFAA to provide neutral, FA-specific resources, with clear boundaries on that role.

Participants were asked what resources they could use to navigate AI use in the financial aid office. Participants identified the following as the most needed resources, roughly in order of frequency across sessions:

- Model policy language for financial aid office AI use, with acknowledgment of institutional variation and a framework for adapting general institutional policies to financial aid-specific contexts
- Training modules specific to financial aid administration, at multiple levels of experience and role, developed by a source that is impartial with respect to specific tools and vendors
- Case studies and examples of how peer institutions are using AI in financial aid, organized by function: student-facing, staff operations, and institutional research, so that offices can identify relevant models
- Risk checklists for evaluating AI use in financial aid contexts
- A vendor evaluation framework, structured questions that financial aid offices should ask when evaluating AI tools, without endorsement of specific products
- A decision framework for AI use, analogous to existing NASFAA decision tools, to guide offices through questions of appropriate use for different functions and data types
- Ethical standards and resources specific to AI use in financial aid

Participants were equally clear about what they do not want from NASFAA:

- Endorsement or ranking of specific AI tools or vendors
- Use of AskRegs or other proprietary NASFAA resources as training data for AI models
- Promotional framing that could be read by institutional leaders as suggesting AI can reduce or replace the need for financial aid staff
- Resources priced or structured in ways that are inaccessible to community colleges and smaller institutions

The word that recurred most often in the Policy-First session and echoed across others was "neutral." Participants want NASFAA to acknowledge AI as a reality of the current environment and provide practical, grounded resources, without cheerleading for adoption or amplifying fears. Several participants noted that NASFAA's credibility as an impartial professional association makes it uniquely positioned to fill a gap that institutions and general higher education organizations cannot fill on their own.

I think there is a real lack of impartial training. The trainings at my institution have been provided by Microsoft. We need some kind of third option here.

One point of tension worth noting: a small number of participants, primarily those at institutions with more advanced AI adoption, expressed the view that NASFAA should be more proactive in encouraging engagement with AI, and should consider incorporating AI into member-facing resources where appropriate. These participants were in the minority, but their perspective reflects a constituency within the membership that does not want NASFAA's caution to translate into a failure to help members navigate a technology that is already here. It is worth noting that how NASFAA develops and delivers tools and resources for members is a distinct operational question from the member-facing resources role described in this section; the findings here speak to the latter.

Institutional context shapes AI adoption in ways that resources need to account for.

One consistent thread across sessions was the significant variation in institutional context, in resources, governance structures, institutional culture, and student population, which shapes what AI adoption looks like and what support is needed. This variation is not simply a matter of more or less adoption. It reflects genuinely different resource environments, governance infrastructures, and institutional priorities that shape what is feasible. Student population also matters. Participants from institutions serving large shares of first-generation, low-income, or undocumented students described bias concerns about AI. Participants also noted that populations most at risk from AI misguidance, students who lack the background knowledge to recognize when AI has given them wrong information, or who face structural barriers when AI systems produce incorrect outcomes, are concentrated at the institutions that already face the greatest resource constraints.

Personal and ethical concerns among staff

Across sessions, participants described a contingent of staff at their institutions who are personally opposed to AI use or who have ethical objections to the technology. These concerns included: environmental impact, particularly the energy and water consumption associated with large AI data centers; distrust of the technology itself and its commercial underpinnings; and discomfort with being expected to use a tool that a staff member believes poses professional or ethical risks. In some institutions, this has created visible internal tension between staff who are enthusiastic about AI and staff who do not want to engage with it at all.

There is a contingent of people who don't want to use AI, don't want anything to do with it. Getting that buy-in has been challenging in some areas.

These concerns appeared in the survey data as well, particularly in qualitative responses from frontline staff, where environmental concerns and categorical opposition to AI were recurring themes. Participants in the listening sessions generally framed this as a management challenge, wondering how to create a culture in which those who want to use AI can do so effectively without requiring or pressuring staff who have genuine objections. Some described requiring a baseline of AI literacy from all staff, even those who do not personally use AI tools, as a reasonable institutional expectation. Others expressed uncertainty about where to draw that line.

The task force's national survey found that 62% of respondents described their own attitude toward AI in financial aid as cautious, while 51% characterized their leadership's attitude as enthusiastic. This gap between practitioner caution and leadership enthusiasm was a palpable undercurrent in several listening sessions, and it points to a broader management and culture challenge.

Topic Area Findings

In addition to the common questions asked across all sessions, each session devoted a portion to a discussion specific to its topic area. The following summaries capture the key themes from those topic-specific discussions. Cross-cutting themes addressed in the findings above are not repeated here

AI Foundations and Types

The two AI Foundations sessions covered the broadest conceptual ground of any topic area and drew participants with notably varied levels of AI familiarity. A significant portion of both sessions focused on correcting two persistent misconceptions: that AI functions as a reliable knowledge base that retrieves accurate information, and that current AI capabilities are either far beyond or far below what they actually are. Participants described staff presenting AI-generated policy research as authoritative without recognizing that AI generates plausible-sounding language rather than verified facts, and described other staff who remain reluctant to engage with AI at all because they associate it with an imagined future of autonomous systems. Both misconceptions, participants noted, can lead to poor decisions: overconfidence in AI output in one direction, and blanket avoidance of useful tools in the other.

The AI Foundations sessions also surfaced the sharpest discussion of financial aid's cross-functional context. Participants noted that AI decisions made in admissions, institutional research, or student affairs can have direct downstream implications for financial aid, and that financial aid offices are often not included in institutional AI governance conversations where those decisions are made. One participant described learning only through informal collaboration that the admissions office had deployed a fraud-screening AI whose outputs affected students who then appeared in the financial aid pipeline. Participants expressed strong interest in cross-association research sharing between NASFAA and other higher education organizations to better understand how AI use in adjacent offices affects financial aid practice.

The question of how AI might affect staffing levels was more prominent in these sessions than in others. Participants described institutional leaders asking whether AI could reduce the number of financial aid staff needed — and expressed concern that this framing misunderstands how financial aid actually operates. Several noted a less-discussed but real consequence of AI handling frontline interactions: newer staff learn financial aid by doing it, and if AI takes over the routine transactions that have historically been the training ground for early-career professionals, the field may face a knowledge pipeline problem over time.

Data Governance, Privacy, and Risk Management

The Data Governance session produced the most detailed discussion of data-handling practices, and participants described a field that is managing real risk largely through informal workarounds rather than formal policy. The practice of using synthetic or de-identified data to test AI tools without exposing actual student records emerged as a common approach developed independently by participants, without institutional guidance. While effective as a workaround, participants noted it depends on staff awareness and discipline, and that a single staff member who does not know the informal rule can create a data exposure incident.

Participants also discussed the significant variation in how institutions are structuring data governance for AI. One institution described a formal data classification system using a four-level scale that determines which data categories can enter which tools, with AI use restricted above a certain classification threshold. Most participants, however, described no comparable structure and expressed interest in NASFAA providing a model framework rather than having each institution develop its own from scratch. Vendor transparency was a recurring concern: participants want to know how vendors handle data, whether data is used to train models, and what contractual protections exist, but described difficulty getting clear answers from vendors on these questions.

Ethical Principles and Student Impact

The Ethical Principles session was the most explicitly values-focused of the six and included direct discussion of personal ethical concerns alongside institutional and professional ones. The group discussed an important distinction early in the session: the difference between ethical concerns about how AI is used in financial aid practice and personal ethical objections to AI itself, including concerns about the technology's environmental footprint, its commercial origins, and the discomfort of being expected to use a tool one does not trust. Both types of concern were present in the room, and participants appreciated having both acknowledged.

The session also produced the most direct discussion of bias. One participant described a fraud detection AI that had generated automatic holds disproportionately affecting undocumented students, a real case with concrete harm that grounded what might otherwise have been an abstract conversation. Participants discussed the difficulty of auditing AI tools for bias before deployment, the limits of vendor assurances, and the particular vulnerability of populations who cannot easily navigate institutional error-correction processes. The session closed with a discussion of what ethical AI use in financial aid would require: not just compliance with data privacy rules, but also active attention to who is helped and who is harmed by AI-mediated decisions and communications.

Policy-First AI Use in Financial Aid

The Policy-First session drew participants who had thought most systematically about institutional policy, and it produced the most detailed discussion of specific policy gaps. Participants identified meaningful gaps in financial aid-specific coverage, particularly around FTI, transcription tools in counseling sessions, and the absence of any named financial aid administrator in institutional AI governance structures. This finding reinforced what the task force's policy scan had found: that even when formal institutional AI policies exist, they address financial aid staff primarily as incidental users of general institutional tools rather than as professionals with distinct regulatory obligations.

Participants in this session were the most explicit about what they want NASFAA to produce and how they want it framed. The word "neutral" recurred throughout the discussion: participants want NASFAA to acknowledge AI as a present reality and provide practical guidance, without promoting adoption or amplifying concern. Several described the challenge of translating high-level institutional policy into office-level practice, and expressed strong interest in a decision-tree framework, similar to NASFAA's existing data-sharing decision tree, that would give financial aid offices a structured way to think through AI use decisions for different functions and data types.

User Experience and Service Design

The User Experience session focused most directly on the student-facing dimension of AI in financial aid and produced the most concrete discussion of current chatbot limitations. Participants described a consistent gap between what students expect from AI-powered service tools and what current chatbot implementations can deliver: students want authenticated data access, accurate follow-up responses, and interactions that feel current and capable, while most deployed chatbots are still operating on architectures that fail on follow-up questions and cannot access real-time student account data. Several participants described their institutions as actively transitioning to more capable vendors, primarily due to student frustration.

The session also addressed how financial aid offices are, and are not, measuring the impact of AI on their work. Most participants described their primary metric as time saved, assessed informally rather than through any structured evaluation framework. None described systematic measurement of AI's impact on student outcomes, service quality, or bias. Participants acknowledged this as a gap and noted that without better measurement, it is difficult to make the case internally for continued investment or to identify cases where AI is producing unintended consequences. The session closed with a discussion of what useful AI impact measurement in financial aid might look like, with particular attention to whether and how offices could track changes in the student populations reaching human staff, and whether those changes represent improvement or displacement.

Methodology

Data collection

NASFAA's Research Department held six Zoom listening sessions in March 2026. Session invitations were extended to NASFAA institutional members who volunteered to be on the Use of AI in the Financial Aid Office Task Force but were not selected and to NASFAA's Rapid Response Network. Participation was voluntary and self-selected. Sessions were organized around five topic areas, with two sessions convened on AI Foundations and Types to accommodate interest:

- AI Foundations and Types (two sessions)
- Data Governance, Privacy, and Risk Management
- Ethical Principles and Student Impact
- Policy-First AI Use in Financial Aid
- User Experience and Service Design for Financial Aid

Each session was facilitated by NASFAA's Director of Research and by the Task Force Chair for the Use of AI in the Financial Aid Office, and followed a semi-structured protocol. The first portion of each session used a common set of questions addressing current AI use, decision-making processes, concerns, guardrails, and training. The second portion used topic-specific questions tailored to the session's focus area. Sessions were recorded with participant consent and transcribed using Zoom's automatic transcription feature. Zoom chat logs were retained as supplemental data for sessions where substantive discussion occurred in the chat.

Analysis

Transcripts were analyzed thematically using an abductive coding approach, consistent with the methodology used in the task force's national survey and institutional policy scan. An initial deductive framework was developed from the session protocol questions and themes identified in the prior phases of the task force research. Codes were then added and revised inductively based on content emerging from the transcripts. A codebook was developed prior to coding and reviewed by NASFAA's Director of Research.

Generative AI tools supported this analysis, including thematic coding, pattern identification across the transcript corpus, and drafting of the prose summaries that appear in this report. This work was conducted using a locally hosted AI model, consistent with NASFAA's internal AI Best Practices policy, which requires that member data and proprietary organizational information not be transmitted to external AI systems. Prior to analysis, all transcripts were reviewed and redacted to remove participant names, institution names, and other identifying details. All AI-assisted analysis was reviewed and edited by the NASFAA Research Department and the task force prior to publication; all interpretive judgments in this report reflect staff analysis.

Limitations

As with all qualitative research, findings are subject to standard limitations, including self-selection bias. The voluntary nature of participation means respondents may hold stronger or more developed opinions about AI than non-participants, and the session format favors practitioners willing to discuss their experiences in a group setting.

Findings describe participants' perspectives and experiences and should not be interpreted as representative of all financial aid offices. Where the task force's national survey provides a more representative quantitative picture, it is referenced throughout the findings section. Frequency language in this report refers to how consistently themes appeared across sessions rather than to counts of individual participants.

Finally, the AI landscape is evolving rapidly. These sessions were conducted in January-March 2026, and the profession's understanding of, access to, and use of AI tools continues to change. The task force encourages readers to consider these findings as a snapshot of a moment in time rather than a static picture of the profession.

Confidentiality

Participants were informed at the start of each session that their comments would be used in aggregate findings and that no names or institution names would be attributed in any published materials. All quotations in this report have been de-identified. Quotations have been lightly edited for readability without alteration of meaning.

Further Information

For further information on the listening session protocols or the analysis code book, please email NASFAA's Research Department at Research@nasfaa.org.