Civic Summit 2026: Analysis of Findings

Executive Summary

The UMass CEP Civic Summit 2026 convened a panel of experts and brought together a mix of people from the community who care about education, to think together about the role of AI in public schools. The conversation was organized around three practical scenarios: students using AI for tutoring, classrooms using AI for brainstorming, and teachers using AI to create content or give feedback. Before the summit, 73 unique participants completed a registration survey, and after the summit, 21 participants completed a follow-up survey. These 21 respondents were matched across both surveys.

The two surveys are not treated as a simple "before and after" measure of the same outcome because they were designed to capture different things. The pre-summit survey helped us to understand where participants were starting from, i.e., their pre-existing exposure to AI and the messages about AI they had heard or personally agreed with. The content of the pre-survey was predicated on the ways the entry of AI into K12 classrooms in the United States has been framed thus far (Nolan-Greene, 2025). The post-summit survey helped us to see where participants landed after the panel and audience discussion about the scenarios, including a proposed set of AI-use "guardrails" and how they described the possible benefits and risks of AI in education.

Headline Findings

- Most participants identified Practical Framing (i.e., "AI is here, let's figure out how to use it well") as both what they hear most in public conversation and what they personally believe. Second most commonly identified was Economic Framing (i.e., "We need AI in schools so students can compete in the workforce"), though less than one-third of respondents personally agreed with it.
- 29% of pre-summit respondents personally agreed with the same message they reported hearing most. The other 71% personally believed something different from what they hear in public discourse.
- After the summit, 81% of respondents agreed or strongly agreed that AI can help teaching and learning when there are proper safeguards in place (average score of 4.05 out of 5). This matches the expert panel's overall view that AI in schools needs guardrails.
- Participants saw about twice as many risks as benefits when AI is used by students (mean of 4.71 risks vs. 2.33 benefits), but maintained a much more balanced view when it comes to teachers using AI (3.14 risks vs. 2.62 benefits). Interestingly, the experts did not show this same pattern—they were more concerned about whether schools provide guidance and structure, regardless of who is using AI.
- One important limitation is that not a single teacher or staff member completed the post-summit survey. The post-summit results mostly reflect the views of students (52%) and community members. Future summits will need a better plan to engage teachers in the post-event survey.

Context and Framing

Instead of having participants sit and listen to experts, the summit invited people to talk through three real-world scenarios with each other and with a panel of experts. The scenarios captured the kinds of situations that dominate public conversation right now: a student turning to AI to help with her math homework (Scenario 1), a teacher telling her class to use AI to brainstorm essay topics (Scenario 2), and a teacher independently adopting an AI platform for grading and feedback

without any guidance from her school (Scenario 3). The expert panel rated Scenarios 1 and 2 as needing guardrails, and Scenario 3 as sitting between "discouraged" and "needs guardrails." What's worth noticing is that no scenario received an outright "red light" from the experts.

The two surveys were intentionally designed to ask different kinds of questions. The pre-summit survey gathered information about where participants stood when they walked in: what messages about AI in education they hear most often in public conversation, and which of those messages they personally agreed with. The post-summit survey looked at where they ended up: how strongly they agreed that AI-use, with appropriate guardrails, can help teaching and learning, and what specific benefits and risks they saw for students and teachers.

The analysis takes the following approach. First, each survey is described on its own. Second, for the participants who completed both surveys, the report looks at how their starting positions connect to where they landed. Third, the post-summit responses are read against the expert panel's ratings to see how the deliberation played out in relation to expert reasoning.

Methods

The pre-summit survey collected respondent name, email, role, AI use frequency on a 1-5 scale, a multi-select item on how participants had learned about AI in education to date, and two five-option, single-selection items asking respondents which of five framings of AI in education they hear most in public discourse and which best reflects their own position. The five framings were Learning-Recovery, Economic, Hope, Practical, and Resist.

The post-summit survey collected name, role, a single 1-5 Likert item on the proposition "With the appropriate guardrails, AI can enhance teaching and learning," and four multi-select items on the perceived benefits and risks of AI in the classroom, asked separately for students and teachers. Duplicate entries in both files were cleaned up by name, keeping each person's most recent submission. After this step, the pre-summit dataset had 73 unique respondents and the post-summit dataset had 21.

The analysis is organized into three parts. The first describes each survey on its own, looking separately at what respondents said pre- and post-summit. The second focuses on the 21 people who completed both surveys, asking whether the framing they endorsed before the summit relates to where they landed on the guardrails question afterward, and whether their prior AI use relates to how they identified benefits and risks. Because the matched group is small, this analysis is exploratory. The third part reads the overall post-summit results against the expert panel's scenario ratings to see how participant reasoning lined up with expert reasoning.

Pre-Summit Findings

Of the 73 pre-summit respondents, 24 (33%) were undergraduate or graduate students, 14 (19%) were teachers, 10 (14%) were school administrators, 9 (12%) were community members, and the remainder were faculty, staff, consultants, or identified themselves with the category "Other."

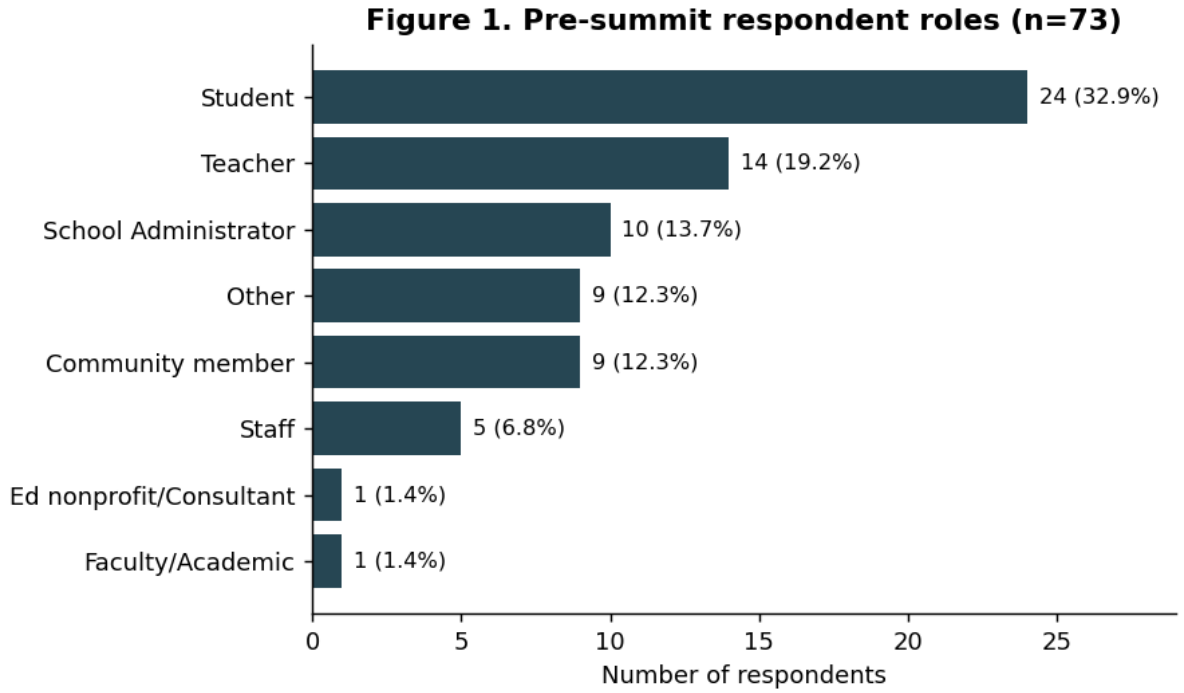

Figure 1. Pre-summit respondent roles.

Self-reported AI use was distributed across the full 1-5 scale, with a mode at 2 (Rarely; n = 24) and substantial representation at 3, 4, and 5. This suggests the registered population was not pre-selected on either high enthusiasm for or strong skepticism toward AI tools.

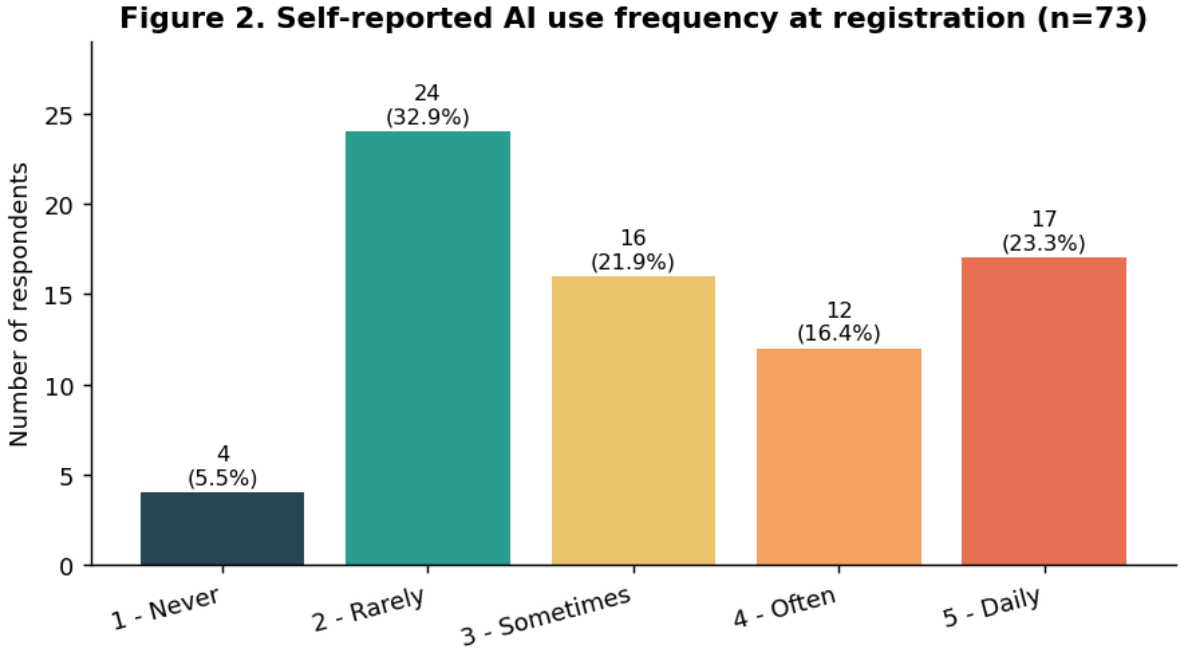

Figure 2. Pre-summit self-reported AI use frequency.

Participants drew on a mix of informal and formal channels to learn about AI in education. Personal experimentation, news media, and conversations with colleagues were the three most common sources.

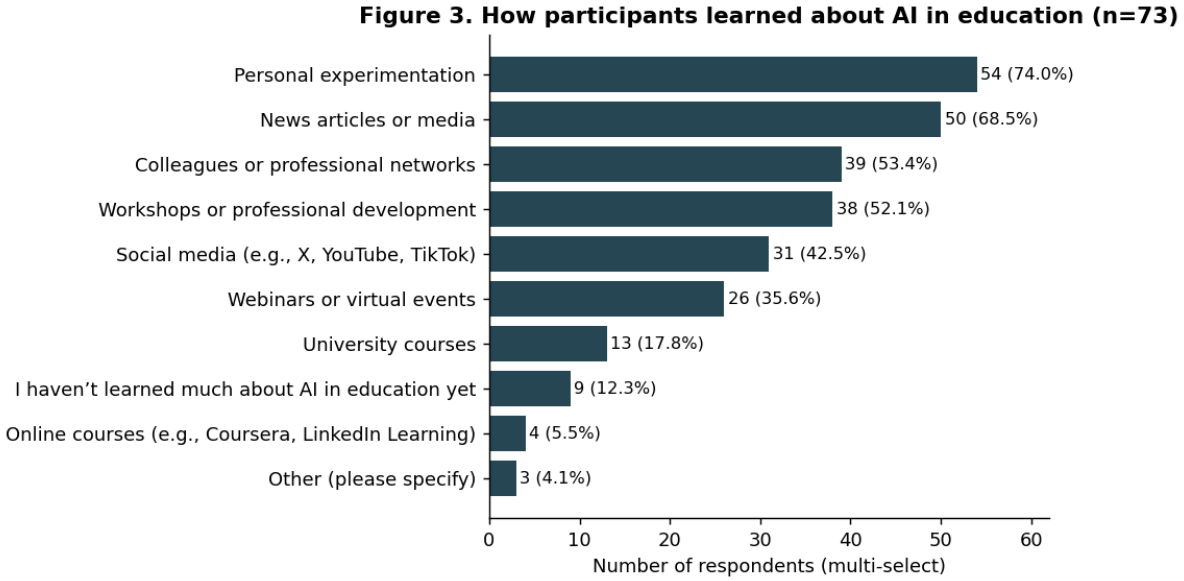

Figure 3. Self-reported sources of learning about AI in education (multi-select).

Framings: What People Hear vs. What They Endorse

The pre-summit survey asked participants to identify which of five predominant framings of AI (Nolan-Greene, 2025) in education they hear most in public discourse, and separately, which best reflects their own position. The five framings were Learning-Recovery (AI as response to learning loss and personalization), Economic (AI adoption as workforce-preparation imperative), Hope (reassurance that AI will not replace teachers or worsen inequality), Practical (AI is here, focus on use), and Resist (AI threatens the human core of teaching and learning).

Economic Framing was identified by 15 respondents as the message they hear most, but only 4 endorsed it personally. Practical Framing rose from 33 heard to 38 endorsed, and Resist Framing rose from 16 heard to 24 endorsed.

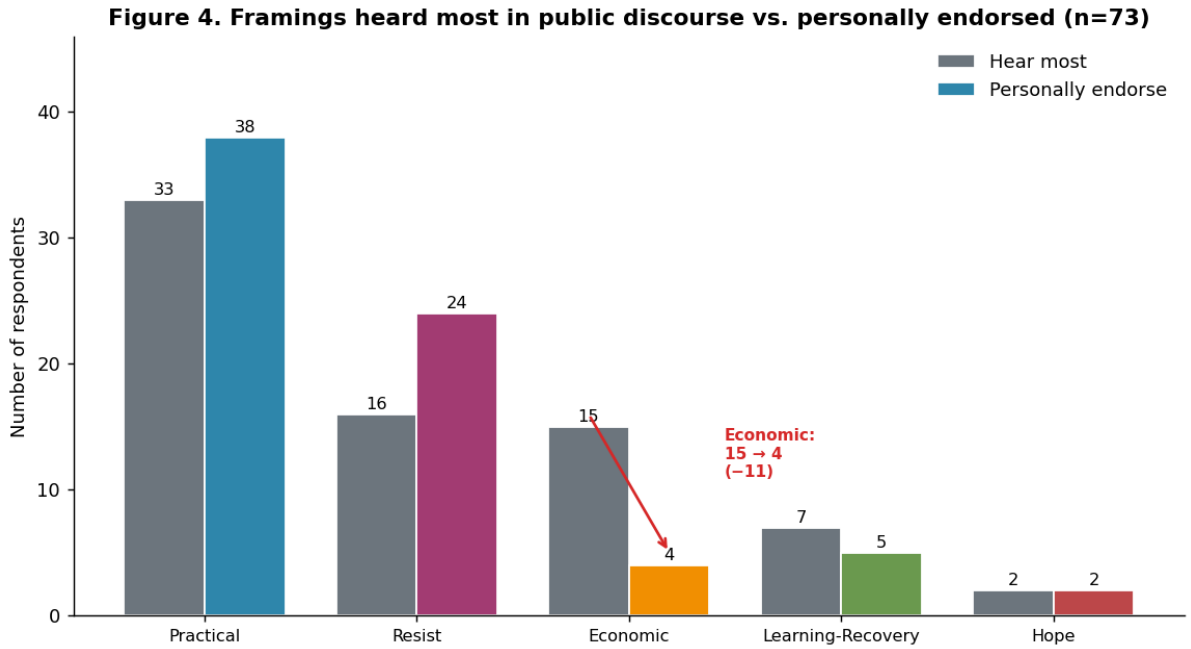

Figure 4. Framings heard most in public discourse versus personally endorsed. The Economic Framing gap (− 11) is the most distinctive feature.

Examined as a crosstab, the same pattern is even sharper. Only 29% of respondents heard and endorsed the same framing (the diagonal of Figure 5); 71% diverged. Among the 15 respondents who heard Economic Framing most, 8 endorsed Practical and 6 endorsed Resist; one endorsed Economic. Among the 16 who heard Resist most, 9 endorsed Practical. Among the 7 who heard Learning-Recovery most, 5 endorsed Practical.

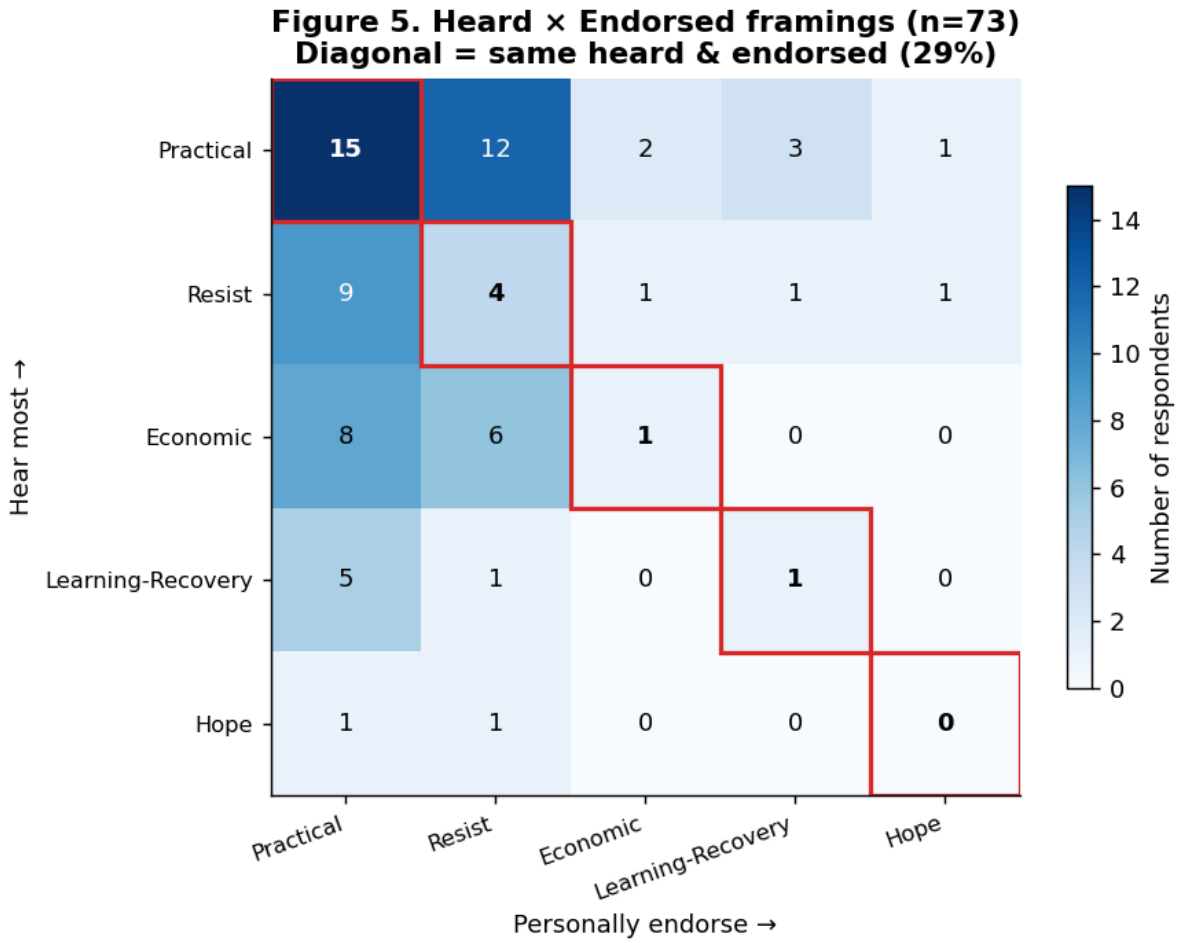

Figure 5. Heard & Endorsed framing crosstab. Diagonal cells (outlined) represent respondents who heard and endorsed the same framing.

Framing Endorsement by Role and AI Use

Framing endorsement varied by role. Teachers were the only role group to endorse Resist Framing more often than Practical (8 vs. 5). Students endorsed Practical at twice the rate of Resist (13 vs. 7). School administrators leaned strongly Practical (6 of 10), and were also the group most likely to endorse Economic Framing (3 of 4 Economic endorsers across all roles came from administration).

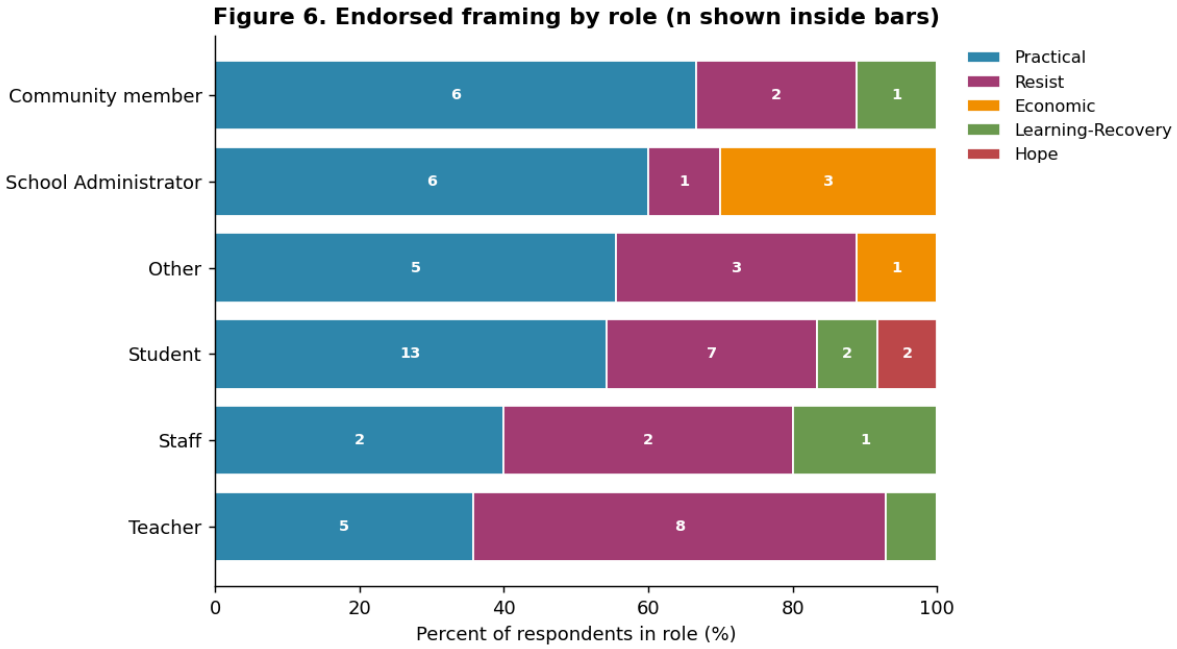

Figure 6. Endorsed framing by respondent role. Roles with fewer than two respondents are omitted.

Higher self-reported AI use was associated with stronger endorsement of Practical Framing and lower endorsement of Resist Framing. Daily AI users (level 5) were the most concentrated Practical endorsers and also the most likely Economic Framing endorsers; Resist endorsement was concentrated among those who used AI rarely (level 2).

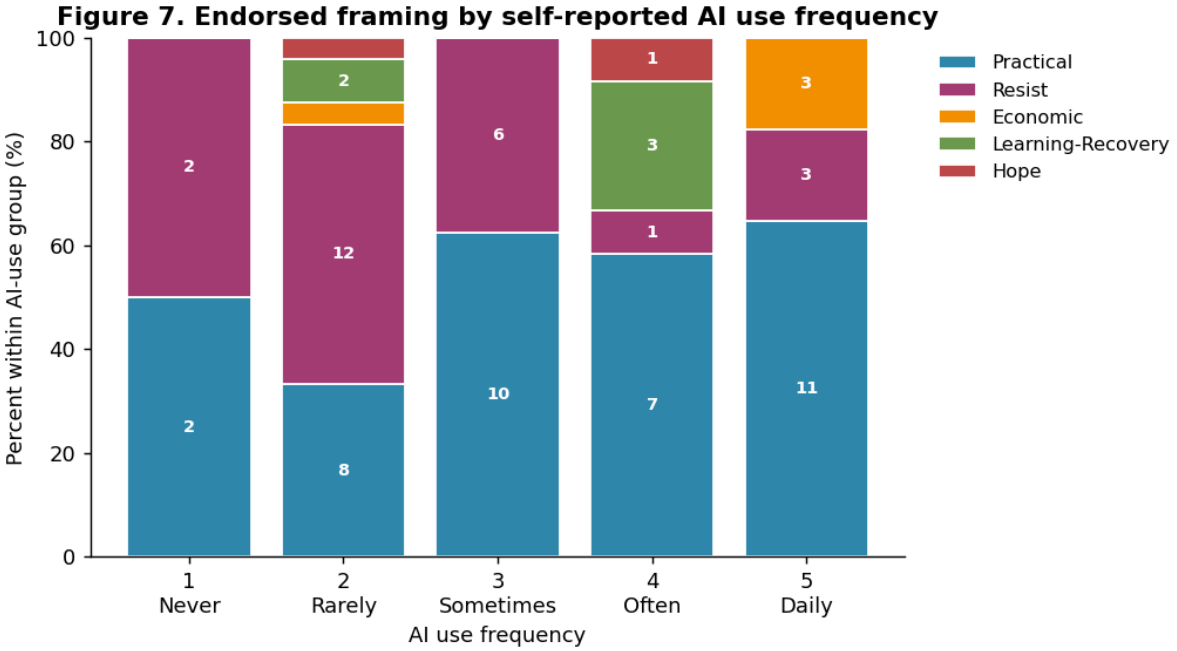

Figure 7. Endorsed framing by self-reported AI use frequency.

Who Returned for the Post-Summit Survey

Of the 73 pre-summit respondents, 21 (29 %) completed the post-summit survey. Return rates differed sharply by role. Students returned at 54% (13 of 24). Community members returned at 33% (3 of 9), and school administrators at 30% (3 of 10). Critically, zero of 14 teachers and zero of 5 staff respondents completed the post-summit survey. As a result, the post-summit sample is heavily student-dominated.

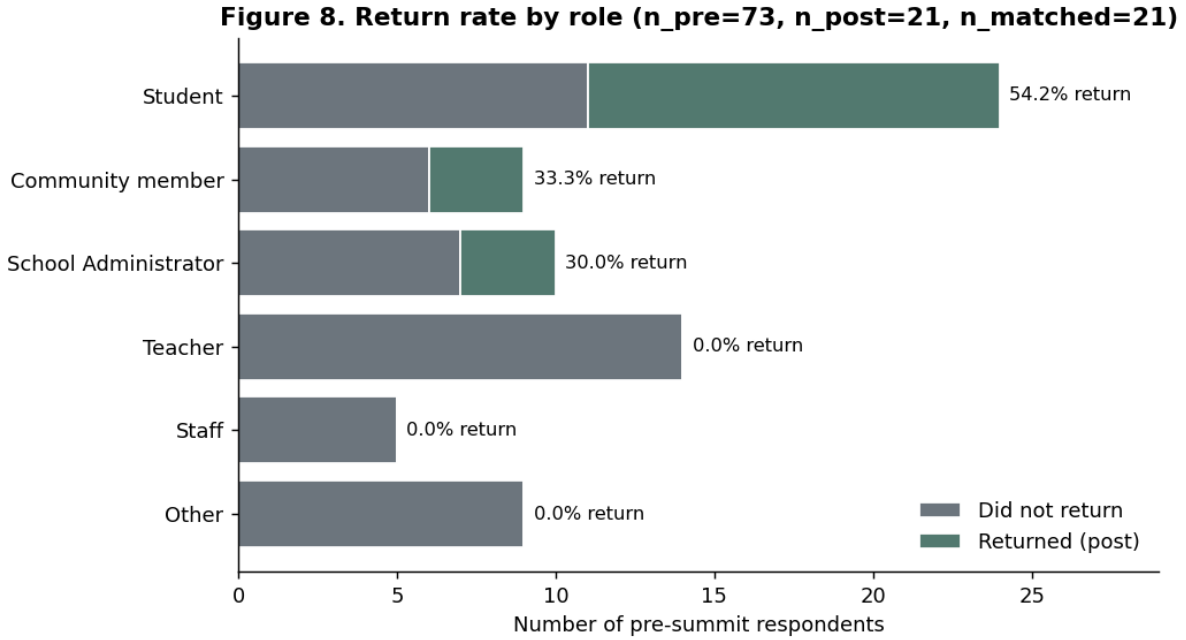

Figure 8. Return rates by role. Teachers and staff are entirely absent from the post-summit data.

Return rates also varied modestly by pre-summit framing endorsement: 1 of 2 Hope endorsers returned, 8 of 24 (33%) Resist endorsers, 10 of 38 (26%) Practical, 1 of 4 (25%) Economic, and 1 of 5 (20%) Learning-Recovery. The teacher and staff absence is the most consequential limitation of the post-summit analysis.

Post-Summit Findings

Agreement with the Guardrails Proposition

Asked to rate their agreement with the proposition "With the appropriate guardrails, AI can enhance teaching and learning" on a 1-5 scale, 81% of post-summit respondents agreed or strongly agreed (M = 4.05, SD = 0.67). No respondent strongly disagreed or disagreed; the lowest response was neutral, given by 4 respondents (19%). Twelve respondents (57%) selected 4 and five (24%) selected 5.

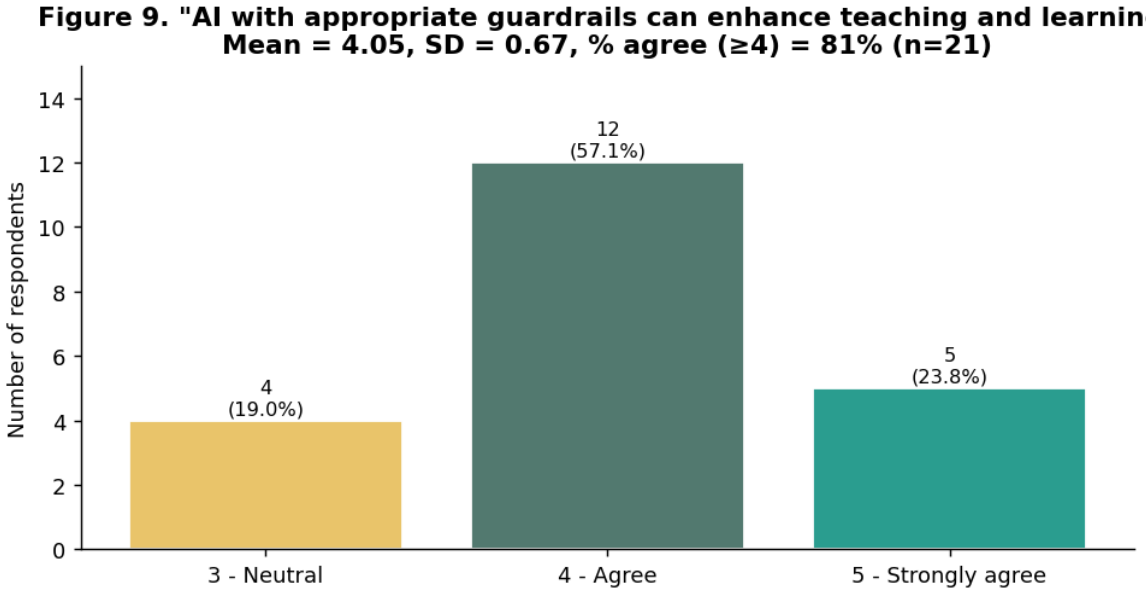

Figure 9. Post-summit agreement with the guardrails proposition (n = 21).

This pattern lines up closely with where the expert panel landed. Across all three scenarios, the experts settled on guardrails or careful, context-dependent restraint. They never gave the "green light" to AI without conditions, and they never told people to reject it outright. The post-summit responses follow the same shape. Agreement with the guardrails idea is strong but not universal, and the steady group of neutral responses suggests that some participants left the summit neither sold on AI's promise nor ready to dismiss it.

Articulated Benefits and Risks

When participants were asked separately about the benefits and risks of AI in the classroom for students and for teachers, clear patterns emerged in each group.

For students, the most commonly chosen benefits were help with organization and studying (76%), explaining concepts and personalized tutoring (52%), and support with brainstorming and research (33% each). For teachers, the top benefits were generating instructional materials (62%), brainstorming instructional content (57%), personalizing lessons (48%), and handling administrative tasks (43%).

Risks were chosen at consistently high rates across the board. For students, the top concerns were less interaction between students and teachers (90%), inaccurate or incomplete information

(81%), less interaction among students themselves (81%), and reduced critical and creative thinking (81%). For teachers, the top concerns were inaccurate or incomplete lessons and content (90%), less interaction with students (76%), generic or impersonal lessons (76%), and job displacement (67%).

Figure 11. Most-selected post-summit benefits and risks

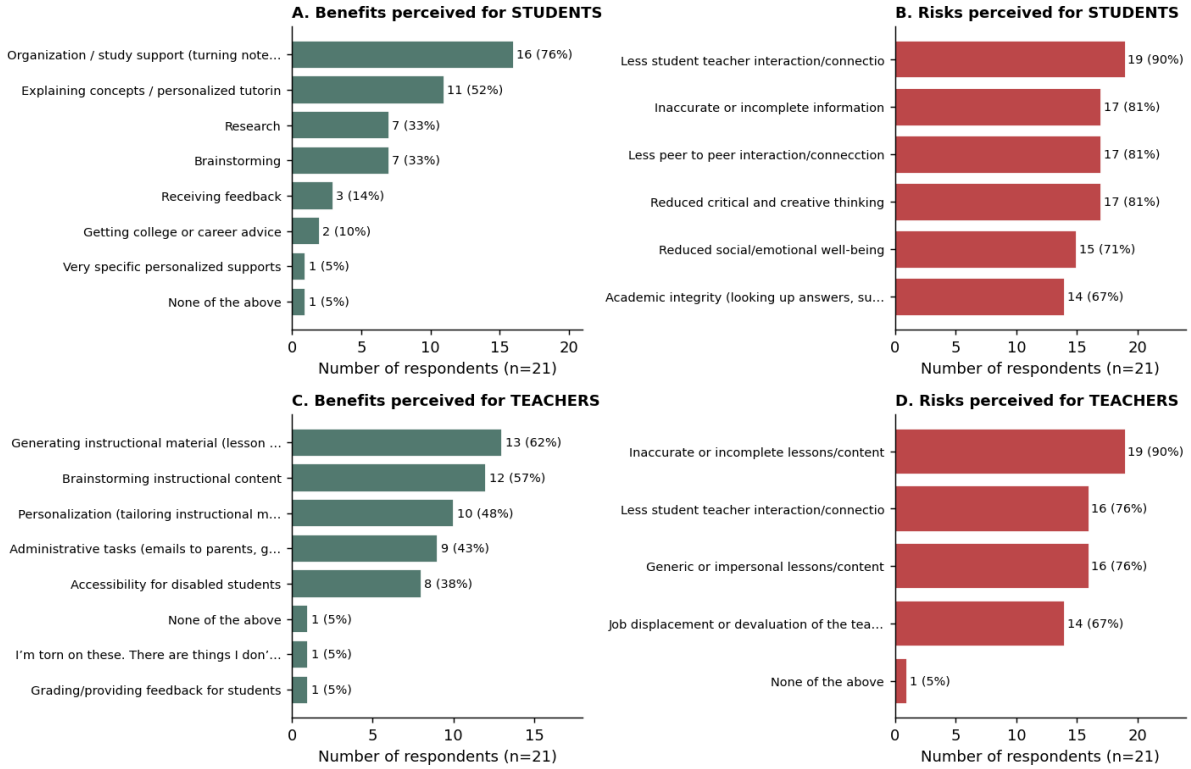

Figure 11. Most-selected post-summit benefits and risks, by target group. Panels A and C are benefits; panels B and D are risks.

Students vs. Teachers

A clear pattern stands out in how participants weighed benefits against risks for students versus teachers. On average, they identified 2.33 benefits and 4.71 risks for students, meaning they saw twice as many risks as benefits (a ratio of about 2-to-1). For teachers, the numbers were closer: 2.62 benefits and 3.14 risks, or roughly 1.2 risks for every benefit. Participants viewed AI for students as twice as risky as it is helpful, but saw the picture for teachers as much more balanced.

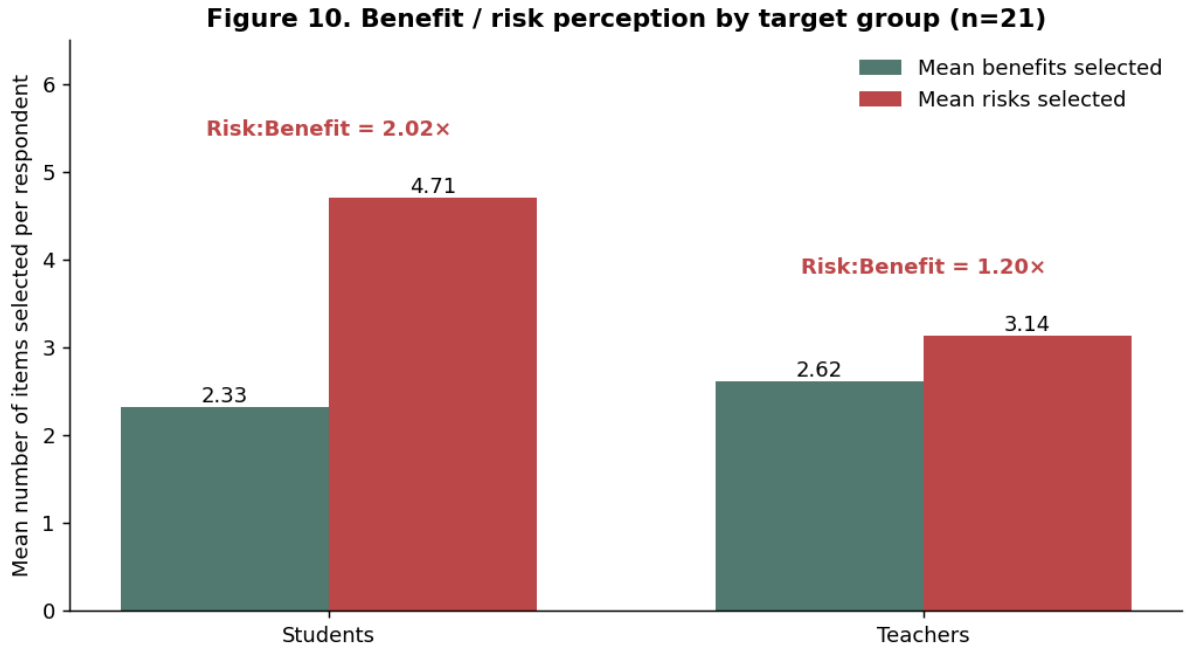

Figure 10. Mean number of benefits and risks selected per respondent, by target group.

Looking at this gap through the three scenarios, part of it makes sense and part of it does not quite match what the experts thought. Scenarios 1 and 2—a student using AI on her own for tutoring, and a teacher asking her class to use AI for brainstorming—were both rated by the experts as needing guardrails. Scenario 3, in which a teacher started using an AI platform on her own without any guidance from her school, was rated as somewhere between discouraged and needing guardrails. But the experts' concern, in this case, had less to do with the teacher using AI and more to do with the fact that she was doing it without any support or guidance from her school.

The experts did not treat students and teachers differently. What mattered most to them was whether the school provided structure and guidance for AI-use. Participants, on the other hand, were noticeably more worried about students than about teachers. They seem to have brought a stronger "protective" instinct into the conversation, one that runs parallel to, but separate from, the experts' focus on institutional support. Both ways of thinking show up in the participant responses — they raised concerns about lost connection, lower-quality content, and weakened professional judgment for both groups — but their worry leans more heavily toward students in a way the expert ratings do not.

Linked Analysis: Pre-Summit Framing and Post-Summit Positions

Among the 21 people who took both surveys, the framing they endorsed beforehand lined up with where they landed afterward. People who endorsed Practical Framing before the summit gave an average score of 4.20 on the post-summit guardrails question (n = 10). People who endorsed Resist Framing averaged 3.75 (n = 8). The three respondents who had endorsed Economic, Learning-Recovery, and Hope Framings each came in at 4.00, 4.00, and 5.00 respectively. The Resist group did not change their minds. They stayed more skeptical than the Practical group, but they did not slip into disagreement either.

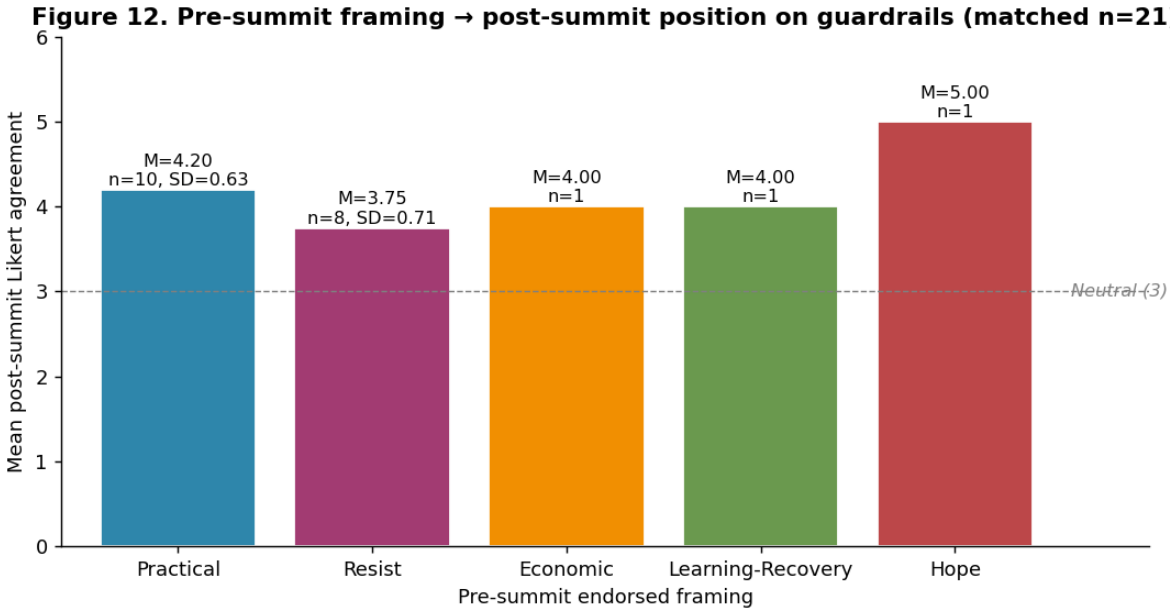

Figure 12. Mean post-summit Likert agreement by pre-summit endorsed framing, matched subset.

Expert panel ratings

The three scenarios were rated by the expert panel as needing guardrails (Scenario 1—independent ChatGPT tutoring; Scenario 2—class-required AI brainstorming) and as between discouraged and needing guardrails (Scenario 3—independent teacher adoption of an AI platform without school guidance). All three ratings fall within a guardrails-or-conditional-restraint space; no scenario received unqualified discouragement, and the more cautious rating of Scenario 3 was located primarily in the institutional vacuum surrounding the teacher's adoption rather than in the use of AI itself.

Limitations

The pre-summit and post-summit samples differ in composition and size (n = 73 vs. n = 21), and the post-summit sample is heavily student-weighted. The matched subset of 21 respondents is too small to support inferential testing. All linked findings are descriptive.

Conclusion

The 2026 Civic Summit asked whether AI in education is a friend or a foe. Ultimately, surveys suggest the summit did something more valuable than producing a single answer. Stakeholders came in with strong and often conflicting views about what AI in schools should mean. They left sharing a common view that AI can help teaching and learning when the right guardrails are in place, while still maintaining different views about where the risks lie and who is most affected.

The pre-summit results showed a community whose personal views often pull away from the loudest messages in public conversation, especially the economic argument for AI in schools. The post-summit results showed a careful stance that matches how the expert panel approached the scenarios. The one notable difference: participants worried more about students than the experts did, while the experts focused more on whether schools provide structure and guidance.