Article

Perceived Adoption, Calibrated Mistrust, and Career Anxiety: Generative AI Through the Eyes of AI-Engaged Undergraduates at an Ivy League University

Francesco Andrea Causio [1,*] and Anushka Sheoran [1]

[1] University of Pennsylvania, Philadelphia, PA 19104, USA;
* Correspondence: fcausio@sas.upenn.edu; asheoran@seas.upenn.edu

Abstract

The rapid proliferation of generative artificial intelligence (GenAI) tools has reshaped higher education, yet fine-grained survey data on student adoption patterns, trust calibration, and career anxiety at technology-forward institutions remain sparse. This study presents exploratory findings from a convenience survey of 119 University of Pennsylvania students, recruited at a campus AI builder community event in February 2026. The sample is self-selected and demographically skewed: 72% of respondents identify as male, 69% as Asian, and 67% report prior AI/ML coursework. Within this cohort, perceived peer adoption of GenAI chatbots is near-universal: 96% of respondents report peer use of ChatGPT, 87% Claude, and 83% Gemini, with 81% estimating daily or near-daily peer usage. A notable 85% report that peers maintain paid subscriptions (72% at $15–50 per month; 13% exceeding $50). Respondents perceive widespread AI-mediated shifts in academic behaviour: 74% agree that peers consult AI in lieu of traditional search engines, 68% believe AI summaries displace assigned readings, and 66% perceive reduced office-hour attendance attributable to AI consultation. Despite 80% of respondents reporting clear understanding of institutional AI policies, only 46% trust AI-generated information to be accurate, a descriptive gap we characterise as a perceived adoption–trust discrepancy. Exploratory gender differences emerge: male respondents rate perceived peer academic displacement behaviours higher, while female respondents report stronger policy understanding and greater career concern ($N_{\text{female}} = 30$; all comparisons uncorrected for multiplicity). AI-related career anxiety is substantial: 74% report that AI has altered their career thinking, 60% worry their specific career path will be negatively affected, and 57% expect AI to surpass human capabilities by 2056, with this anxiety correlating positively with demand for expanded institutional AI curricula. We situate these findings within the broader literature on student AI adoption [1–5] and discuss implications for AI literacy education, institutional policy design, and equity in premium tool access. Given the self-selected, technology-forward composition of the sample, we caution that these estimates almost certainly represent perceived upper-bound adoption rates among the most AI-engaged undergraduates and cannot be generalised to the broader Penn student body, let alone to wider higher-education populations.

Keywords: generative AI; ChatGPT; higher education; student attitudes; AI adoption; academic integrity; career anxiety; gender differences; AI literacy; peer perception; convenience sample

1. Introduction

The release of ChatGPT in November 2022 constituted an inflection point in the relationship between artificial intelligence and higher education. Within three years, the landscape transitioned from tentative experimentation to pervasive integration: the United Kingdom's Higher Education Policy Institute (HEPI) reported that student use of generative AI for assessments rose from 53% in 2024 to 88%

in 2025, a shift the authors characterised as "almost unheard of" in magnitude [3]. At the University of California, Los Angeles, the majority of graduating seniors in the Class of 2025 (73%) reported using GenAI tools to support their coursework [4]. Baek, Tate, and Warschauer [2024], in a survey of 1,001 U.S. college students, documented that ChatGPT use was widespread but stratified by academic discipline, linguistic background, and institutional context. Cabrera et al. [2] subsequently characterised student GenAI engagement as "widely used but barely trusted," capturing the ambivalent posture that has come to define the student–AI relationship. These data underscore a reality that universities are still endeavouring to integrate into their pedagogical frameworks and governance structures.

Notwithstanding this growing evidentiary base, the quantitative portrait of student AI adoption remains unevenly developed along several critical axes. Most extant surveys, including the HEPI study ($N = 1{,}041$), the UCLA Senior Survey ($N = 6{,}639$), the ACT survey of high school students ($N = 3{,}816$; Schiel 5), the Baek et al. [1] multi-campus investigation ($N = 1{,}001$), and the Cabrera et al. [2] trust-focused study, capture broad populations but at relatively low granularity with respect to the interplay between technology adoption, trust calibration, and career anxiety. Moreover, these surveys predominantly rely on self-reported behaviour, which may be susceptible to social desirability and recall biases. Surveys conducted at individual elite institutions offer a complementary lens: they illuminate how AI is perceived in environments characterised by high technological literacy, selective admission, and competitive academic cultures, wherein the stakes of AI adoption (both for learning processes and for career outcomes) are particularly salient.

The present study addresses this gap through an exploratory survey of 119 students at the University of Pennsylvania, an Ivy League institution, administered in February 2026 at a campus event organised by the Penn Claude Builder Club. Critically, our instrument captures peer-observed, rather than self-reported, AI tool adoption, usage frequency, premium subscription behaviour, perceived academic displacement effects, trust in AI-generated information, privacy concerns, career anxiety, and attitudes toward institutional AI policy. We complement quantitative Likert-scale data with open-ended qualitative responses, treated here as illustrative rather than systematically coded, regarding course-level AI policies, perceived AI limitations, creative AI use cases observed among peers, and hopes and concerns for the technology's trajectory.

We extend prior findings in three respects. First, we document a descriptive gap between perceived peer AI adoption (near-universal) and self-reported trust in AI accuracy (held by a minority), a pattern consistent with Cabrera et al.'s [2025] characterisation of student GenAI as "widely used but barely trusted" and which may reflect emerging pragmatic scepticism rather than uncritical acceptance (although we note that our peer-observation and self-report measures are not commensurate and thus not directly comparable). Second, we identify exploratory gender differences in AI-related perceptions that extend beyond simple usage gaps: female respondents in our sample report greater career concern and stronger policy understanding, while male respondents rate perceived peer academic displacement behaviours more highly. Third, we find that perceived AI-related career anxiety is pervasive (74% of respondents report that AI has altered their career thinking) and that this anxiety is associated with desire for expanded institutional AI education, suggesting interest in AI literacy curricula among this self-selected cohort.

We are forthright about the limitations of our sample. Respondents were recruited from a campus AI builder community event, yielding a population that is overwhelmingly male (72%), Asian (69%), and computer-science-literate (67% having completed AI/ML coursework). These demographic skews are not representative of Penn's student body, which is approximately 55% female and 27% Asian, let alone the broader U.S. higher-education landscape. Our estimates should be interpreted as perceived upper-bound adoption rates within a self-selected, technologically sophisticated cohort. All subgroup analyses (including gender comparisons based on 30 female respondents) are exploratory and require replication in larger, more representative samples before substantive conclusions can be drawn.

2. Methods

2.1. Study Design and Recruitment

This study employed a cross-sectional, single-institution, exploratory survey design. The instrument was administered via Google Forms and distributed electronically to attendees of a Claude Builder Club event held at the University of Pennsylvania in February 2026. The survey remained open for approximately three hours during the event. Participation was voluntary and anonymous; no financial or academic incentives were offered. All participants provided informed consent at the commencement of the survey. The full survey instrument is available as a supplementary file.

2.2. Response Rate and Missing Data

We received 119 completed responses. Because attendance at the event was not formally enumerated and the survey link was distributed without a closed recipient list, we are unable to compute a response rate or estimate the denominator of eligible participants. Among the 119 respondents, item-level nonresponse was minimal: no Likert-scale item exceeded 2% missing values. Missing responses were excluded pairwise from the relevant analyses. No imputation was performed.

2.3. Participants

Demographic characteristics are summarised in Figure 1 and described in Section 3.1. In brief, respondents were predominantly male (72.3%), Asian (68.9%), and drawn from the Class of 2027 (46.2%). A majority (67.2%) had completed at least one computer science course at Penn pertaining to AI or machine learning. Graduate students (Master's and PhD) comprising approximately 3% of the sample. Academic concentrations encompassed engineering, computer science, data science, the natural sciences, the social sciences, medicine, and business. A minority of respondents (19.3%) reported receiving need-based financial aid.

2.4. Measures

The survey instrument comprised 37 items distributed across five domains: (1) demographics and academic background, (2) peer-observed AI tool adoption and usage patterns, (3) perceived academic impact of AI on peer behaviour, (4) personal attitudes toward AI (trust, privacy, career impact, institutional policy, societal concerns), and (5) open-ended qualitative reflections. Table 1 summarises the Likert-scale items employed in the primary analyses. All Likert items utilised a five-point agreement scale: 1 = Strongly Disagree, 2 = Moderately Disagree, 3 = Neither Agree nor Disagree, 4 = Moderately Agree, 5 = Strongly Agree. The full questionnaire with item wording is provided in the supplementary materials.

A critical design choice warrants explicit discussion: the adoption, usage frequency, and academic displacement items were framed in terms of peer-observed behaviour ("My peers often consult AI. . . ," "My classmates are less likely to. . . ") rather than self-reported behaviour. This design was deliberate, as it captures perceived social norms, which independently influence behaviour [6], and may attenuate social desirability bias in the reporting of potentially norm-violating behaviours. It follows, however, that these measures reflect perceived peer behaviour and must not be interpreted as direct behavioural prevalence estimates.

Table 1. Likert-scale survey items employed in primary analyses. All items used a five-point agreement scale (1 = *Strongly Disagree*; 5 = *Strongly Agree*). The full instrument is available as a supplementary file.

Domain	Item Stem (Abbreviated)
Perceived Peer AI Behaviour	Peers consult AI instead of Wikipedia/Google
	Classmates less likely to complete readings (AI summaries)
	Classmates less likely to attend office hours (AI consultation)
	Classmates less likely to attend lectures (AI availability)
Policy & Trust	I understand Penn class rules regarding GenAI use
	Worried peers gain unfair advantage through AI
	I trust AI-generated information to be accurate
	Worried about privacy of peer conversations with AI
Societal Concern	Worried governments/corporations will censor AI
	AI products created with people like me in mind
Personal Impact	Surprised by rapid AI advancement 2024–2026
	AI has influenced which classes I choose at Penn
	AI has changed how I think about my future career
Institutional Demand	I wish Penn offered more AI-use classes
	I wish Penn offered more AI-society/future classes
Career Anxiety	Worried my career path will be negatively affected by AI
	Worried AI increases economic inequality
AI Futures	Expect AI $>$ humans in nearly all regards by 2056
	Mitigating AI extinction risk should be a global priority

Multi-select items captured peer-observed AI products, use cases, and premium subscription behaviour. Open-ended questions solicited qualitative responses regarding course-level AI policies, perceived AI limitations, creative AI use cases observed among peers, hopes and concerns about AI's future, and desired institutional resources.

2.5. Data Analysis

Descriptive statistics (means, standard deviations, percentages) were computed for all Likert-scale and categorical items. Independent-samples Welch's t-tests were employed for exploratory subgroup comparisons across gender (male, $n = 86$; female, $n = 30$) and computer science background (prior AI/ML coursework: $n = 80$; no prior coursework: $n = 39$). Given the large number of comparisons conducted (20 items $\times$ 2 grouping variables), these tests must be interpreted as exploratory; no correction for multiple comparisons was applied, and all p-values are reported for transparency with the explicit caveat that none would survive conservative Bonferroni correction. Effect sizes (Cohen's d) are reported alongside p-values for nominally significant comparisons. Pearson correlations were computed among key Likert items to examine associative patterns between adoption perceptions, trust, and career anxiety; these are descriptive associations and do not warrant causal interpretation. Qualitative responses were analysed through close reading and thematic categorisation by the research team; no formal intercoder reliability protocol was employed, and qualitative themes must be treated as illustrative rather than systematically derived. All analyses were conducted in Python 3.11 using NumPy for computation and matplotlib 3.8 for visualisation. The anonymised dataset and analysis code are available from the corresponding author upon reasonable request.

3. Results

3.1. Demographics

The sample exhibited demographic skew relative to the broader Penn undergraduate population (Figure 1). Male respondents outnumbered female respondents by a ratio of approximately 2.9:1 (72.3% vs. 25.2%), and Asian students comprised 68.9% of the sample, followed by White (19.3%), Hispanic/Latino (5.9%), and Black/African-American (2.5%). The Class of 2027 constituted the modal cohort (46.2%), followed by the Classes of 2026, 2028, and 2029 (each approximately 16.8%). A minority of respondents (19.3%) reported receiving need-based financial aid. These demographic characteristics are consistent with self-selection from a campus technology builder community and must not be interpreted as representative of the Penn student body.

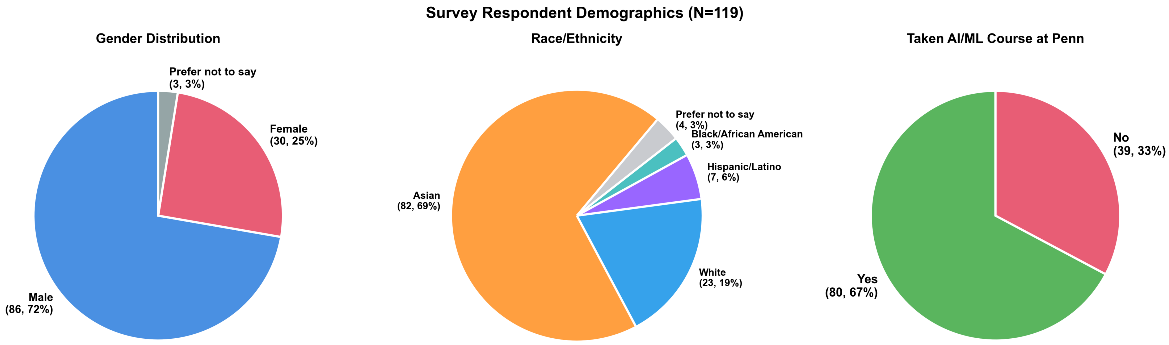

Figure 1. Demographic profile of survey respondents ($N = 119$). Panels display gender, ethnicity, class year, and self-reported CS/AI course background. Note the pronounced skew toward male (72.3%), Asian (68.9%), and CS-experienced (67.2%) respondents, consistent with self-selection from an AI builder community event.

3.2. *Perceived Peer AI Adoption and Usage Patterns*

Perceived peer AI chatbot adoption was near-universal within this self-selected sample. Every respondent reported knowing at least one friend or classmate who uses generative AI products. The most commonly observed products were ChatGPT (96% reported peer use), Claude (87%), and Gemini (83%), followed by Perplexity AI (39%), Grok (32%), and DeepSeek (31%). Midjourney (8%) was the sole image-generation tool mentioned with notable frequency (Figure 2). The predominance of text-based chatbots over multimodal generation tools is consistent with the primarily academic and programming-oriented use cases reported by this population.

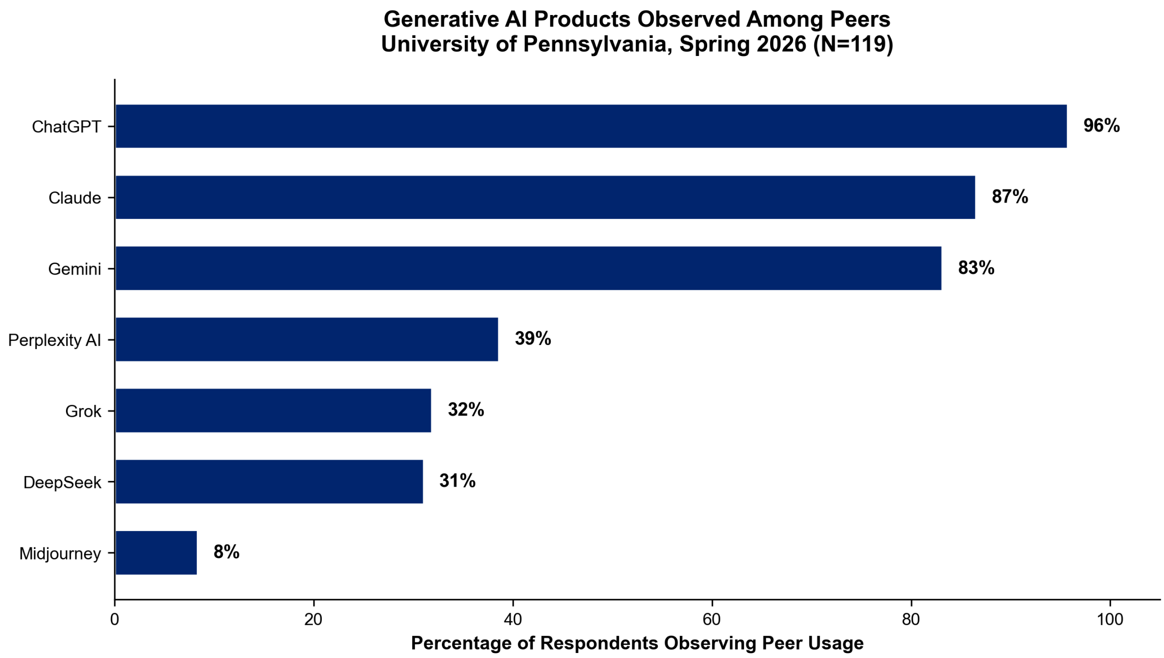

Figure 2. Percentage of respondents reporting perceived peer use of each generative AI product. Multiple selections were permitted. $N = 119$.

Perceived peer usage frequency was exceptionally high: 81% of respondents estimated that peers use GenAI chatbots daily or almost daily, with an additional 9% reporting every-other-day use, 6% reporting weekly usage, and merely 4% estimating biweekly usage. No respondent estimated less frequent use than biweekly. These estimates substantially exceed the 88% of UK undergraduates who self-reported ever using AI for assessments in the HEPI survey [3] and the 73% of UCLA seniors who self-reported using AI to support coursework [4]. We emphasise, however, that these metrics are not directly comparable: our measure captures perceived peer daily usage frequency, whereas the HEPI and UCLA instruments capture self-reported ever-use for specific academic purposes. The discrepancy likely reflects both genuine adoption differences in this AI-engaged cohort and measurement divergence.

The most commonly observed use cases were programming assignments (83%), writing assignments encompassing ideation, drafting, and proofreading (82%), and general knowledge or search queries (82%). Writing emails (78%), data processing (66%), and translation or language learning (36%) were also frequently cited. Graphics, art, and creative work (35%) and entertainment or companionship (30%) completed the observed use cases (Figure 3). The heavy concentration in programming and writing reflects both the sample's STEM orientation and the documented strengths of current-generation AI in code generation and text composition.

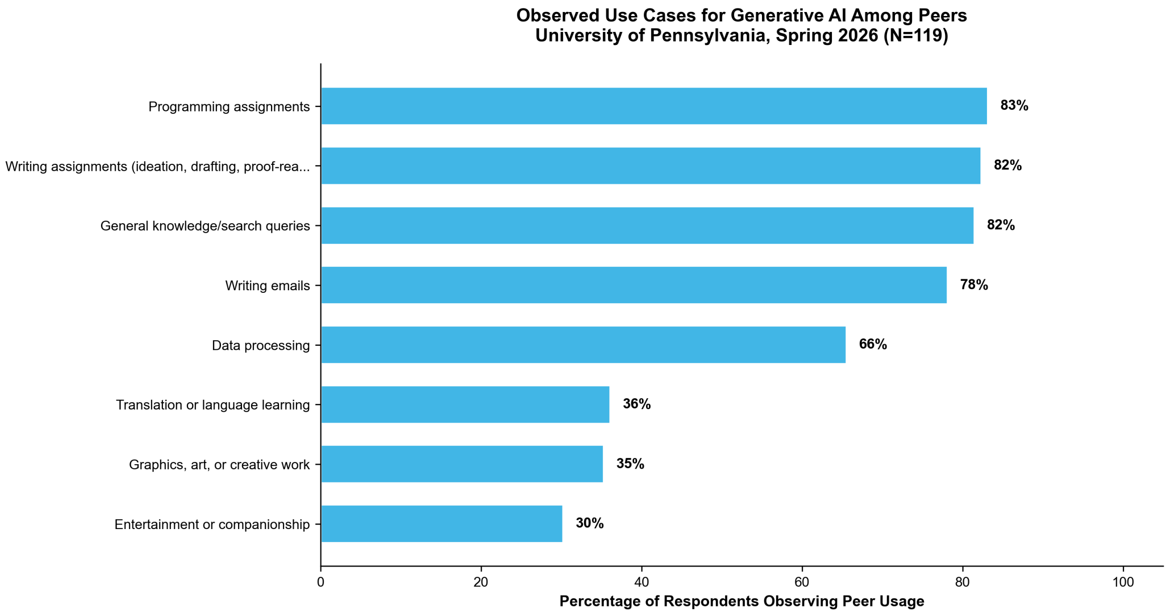

Figure 3. Percentage of respondents perceiving each peer use case for generative AI. Multiple selections permitted. $N = 119$. Data reflect perceived peer behaviour in a self-selected, AI-engaged sample.

3.3. Premium Subscriptions

A notable 85% of respondents perceived that at least some peers maintain paid premium subscriptions to GenAI tools. Specifically, 72% estimated that peers expend $15–50 per month, while 13% estimated peer expenditure exceeding $50 per month. Only 10% believed peers exclusively utilised free versions, and 5% were uncertain. Perceived premium adoption did not differ significantly between students with and without CS backgrounds, suggesting that willingness to pay for AI access in this sample is not restricted to those with formal technical training (Figure 4). This finding, while derived from perceived peer behaviour, raises equity concerns: if premium AI access confers academic advantages—faster code generation, more capable models, larger context windows—then differential access predicated on ability to pay constitutes a potential new dimension of the digital divide in higher education. We caution, however, that perceived peer expenditure may be inflated by projection or salient exemplars, and direct spending data are required to corroborate these estimates.

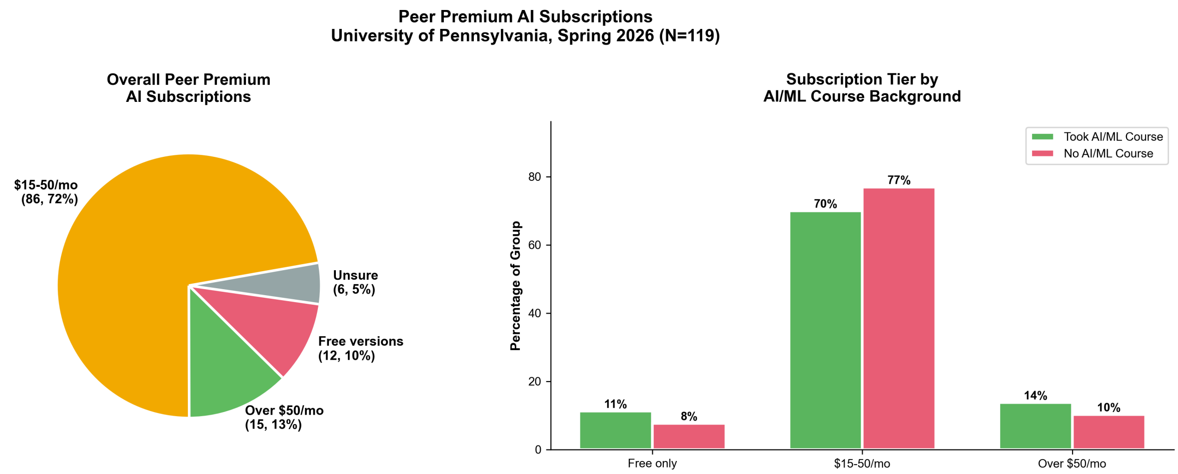

Figure 4. Perceived peer premium subscription expenditure: overall distribution (left) and comparison by self-reported CS/AI coursework experience (right). $N = 119$. Data reflect respondents' estimates of peer spending, not verified expenditure.

3.4. Perceived Academic Impact

Respondents perceived substantial AI-mediated shifts in peer academic behaviours (Figures 5 and 6). A strong majority (74%) agreed that peers consult AI instead of using Wikipedia or Google ($M = 3.92$, $SD = 1.23$). Similarly, 68% agreed that classmates are less likely to complete assigned readings because they employ AI to summarise them ($M = 3.73$, $SD = 1.11$), and 66% agreed that AI consultation is perceived to reduce office-hour attendance ($M = 3.74$, $SD = 1.16$). Perceived impact on lecture attendance was weaker: 35% agreed that AI reduces lecture attendance ($M = 3.10$, $SD = 1.16$), with 29% expressing disagreement. This gradient, with strongest perceived displacement for search and reading, moderate displacement for office hours, and weakest displacement for lectures, suggests that, in respondents' perception, AI currently substitutes for asynchronous, information-access activities to a greater extent than for synchronous, community-based academic experiences. Whether these perceptions correspond to actual behavioural changes or reflect projected norms cannot be determined from the present data.

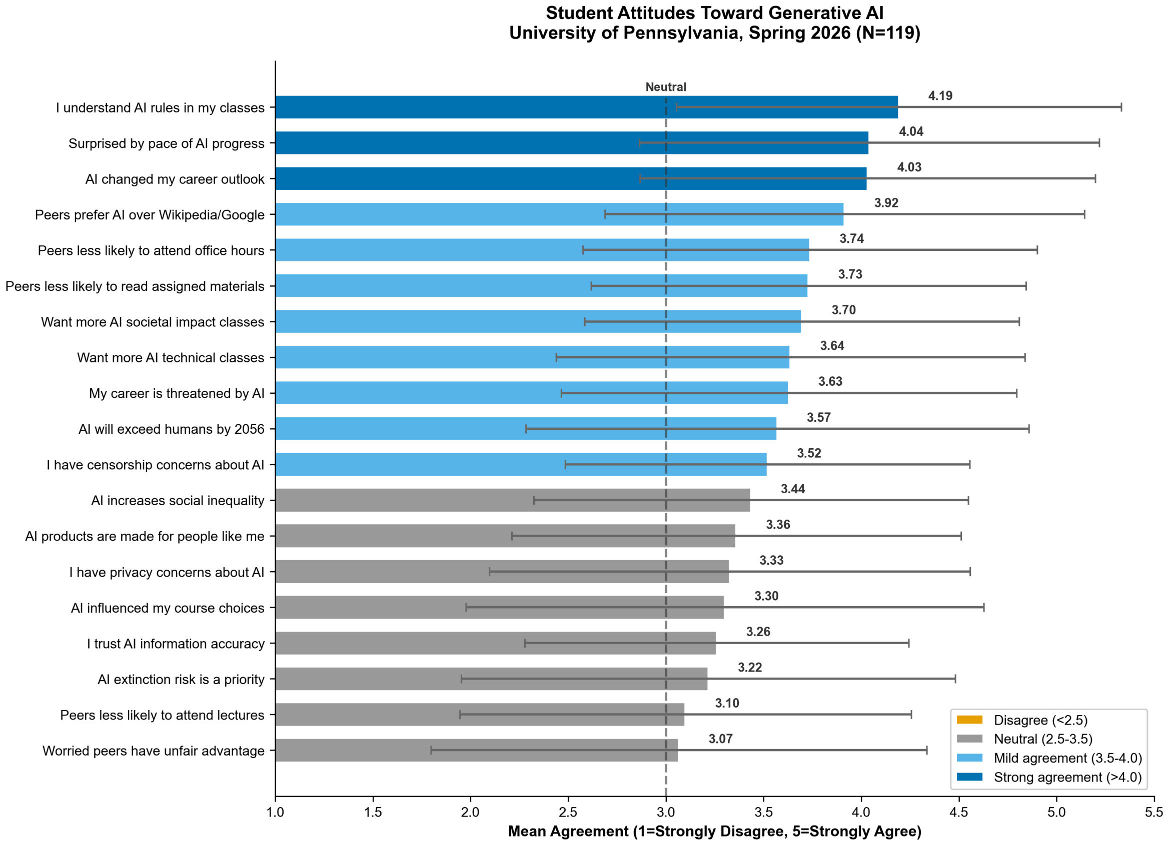

Figure 5. Mean agreement scores for all Likert-scale items ($N = 119$). Error bars represent ± 1 standard error. Dashed vertical line at 3.0 marks the neutral midpoint. Items are sorted by descending mean agreement. All measures reflect respondent perceptions (peer-observed for academic behaviour items; self-report for trust, policy, career, and societal items).

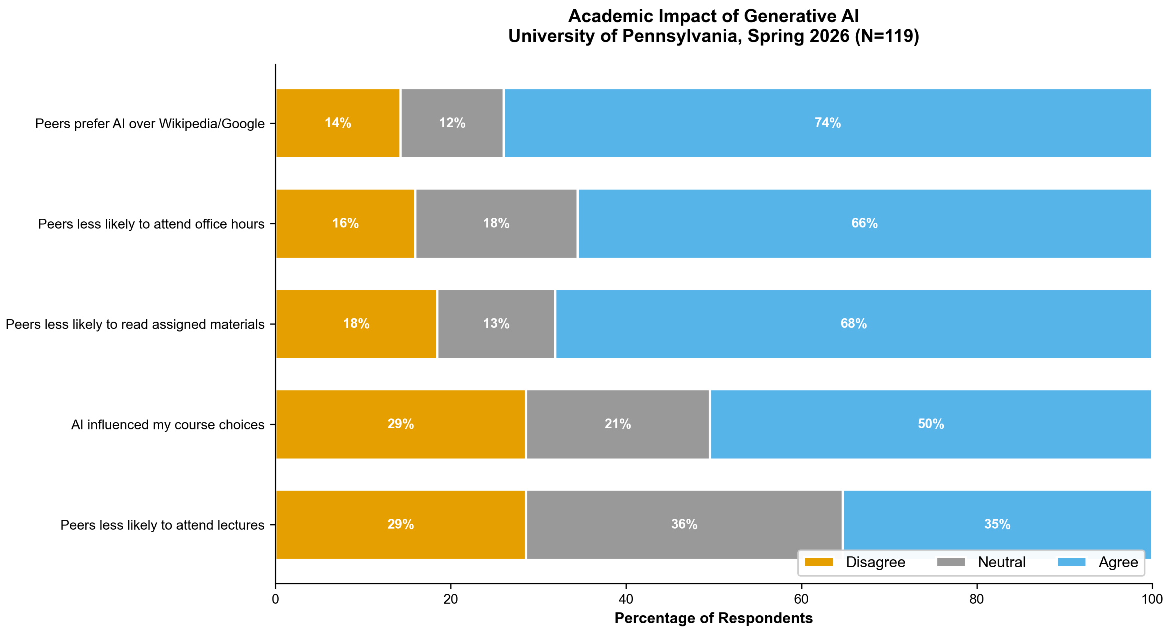

Figure 6. Perceived academic impact of AI: percentage of respondents agreeing, neutral, or disagreeing with each statement ($N = 119$). Items reflect perceived peer behaviour, not verified behavioural observation.

Perceived AI influence on academic decision-making was moderate: 50% agreed that AI has influenced which classes they choose at Penn ($M = 3.30$, $SD = 1.33$), with 29% disagreeing. This suggests that for a meaningful subset of this self-selected, AI-interested cohort, the presence of AI is not merely a tool utilised within courses but a factor that shapes curricular trajectory.

3.5. Policy Awareness and Trust

A large majority of respondents (80%) reported understanding the rules governing AI use in their current Penn courses ($M = 4.19$, $SD = 1.14$), one of the highest-agreement items in the survey. This finding aligns with HEPI's [2025] report that 80% of UK students found their institution's AI policy "clear." However, policy understanding did not translate into uniformly reduced concern about fairness: 37% were concerned about peers gaining an unfair advantage through AI ($M = 3.07$, $SD = 1.27$), with a comparable 34% expressing no such concern, a near-even split that suggests ambivalence regarding AI's role in academic integrity.

Self-reported trust in AI-generated information was notably low relative to perceived peer usage rates (Figure 7). Only 46% agreed that they trust AI-provided information to be accurate ($M = 3.26$, $SD = 0.98$), while 24% expressed disagreement. This descriptive gap between perceived near-universal peer adoption (measured via peer-observation items) and minority self-reported trust is consistent with Cabrera et al.'s [2025] characterisation of student GenAI as "widely used but barely trusted." We emphasise, however, that these constructs are assessed at different levels of analysis (peer-observed behaviour versus self-reported attitudes) and the apparent discrepancy may partly reflect this measurement asymmetry. The observed positive correlation between trust in AI accuracy and AI-altered career thinking ($r = 0.51$) indicates that respondents who trust AI more are also those most engaged with its career implications, though the directionality of this association cannot be established from cross-sectional data. Several interpretations are compatible with the data: genuine pragmatic scepticism, habitual use despite acknowledged unreliability, task-specific trust calibration, or social desirability in the reporting of low trust.

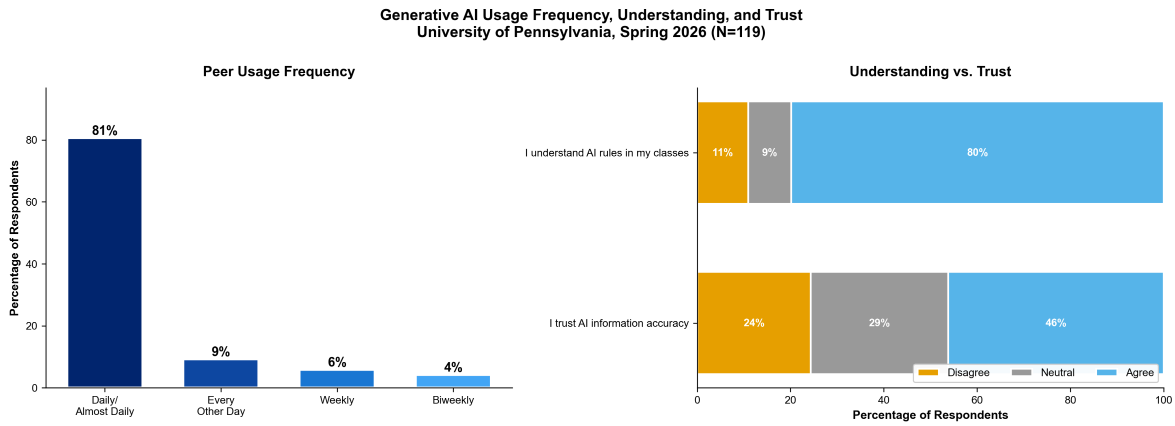

Figure 7. Left: perceived peer AI chatbot usage frequency ($N = 119$). Right: self-reported trust in AI accuracy versus self-reported policy understanding. Note that frequency data measure respondent estimates of peer behaviour, while trust and policy items are self-report measures.

3.6. Career and Societal Concerns

AI-related career concern emerged as one of the survey's most prominent findings (Figure 8). A large majority of respondents (74%) agreed that AI has altered the way they think about their future career ($M = 4.03$, $SD = 1.17$), and 60% expressed worry that their specific career path will be negatively affected by AI automation ($M = 3.63$, $SD = 1.17$). Nearly half (48%) worried that AI will increase economic inequality among Penn graduates and the wider world ($M = 3.44$, $SD = 1.11$). The 60% career-worry rate is broadly consistent with Pew Research Center [2024] data indicating that 62% of Americans believe AI will exert a major impact on jobholders in the next two decades, and with ACT's [2025] finding that high school students expressed significant concern about AI's impact on employment.

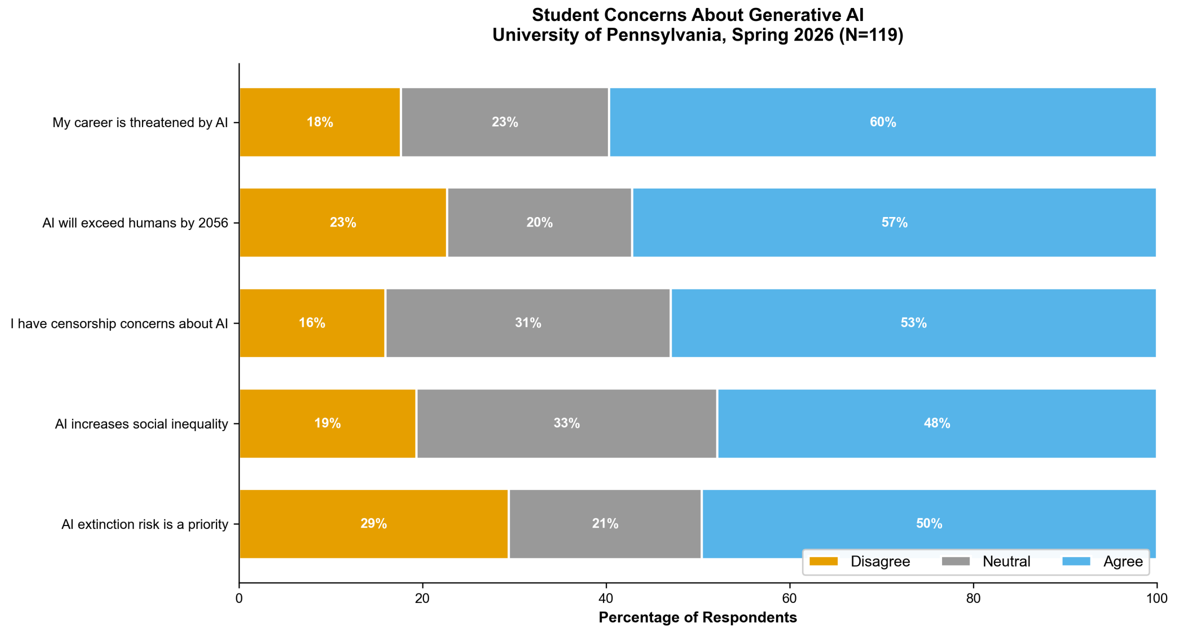

Figure 8. Career and societal concerns about AI: percentage of respondents agreeing, disagreeing, or neutral on each item ($N = 119$). Career anxiety items exhibit high agreement rates; note that these findings from a self-selected AI-builder cohort may overestimate campus-wide concern.

With respect to AI's longer-term trajectory, 75% reported being surprised by the pace of AI advancement between 2024 and 2026 ($M = 4.04$, $SD = 1.18$). A majority (57%) expected AI to surpass human capabilities in almost all regards by 2056 ($M = 3.57$, $SD = 1.29$). Precisely half (50%) agreed that mitigating AI extinction risk should be treated as a global priority alongside pandemics and nuclear war ($M = 3.22$, $SD = 1.26$), while 29% disagreed, a distribution broadly consistent with expert opinion surveys on AI existential risk [8]. Privacy concerns were moderate: 50% worried about the privacy of peer conversations with AI ($M = 3.33$, $SD = 1.23$), and 53% worried about governmental or corporate censorship of AI ($M = 3.52$, $SD = 1.04$). Only 49% felt that AI products were created "with people like me in mind" ($M = 3.36$, $SD = 1.15$), suggesting that a meaningful portion of respondents, particularly those from underrepresented backgrounds, perceive AI development as insufficiently inclusive.

3.7. Exploratory Gender Differences

Exploratory Welch's t-tests revealed several nominally significant gender differences ($p < 0.05$, uncorrected; Figure 9). Given the small female subsample ($n = 30$), the imbalanced design, and the large number of comparisons conducted, these findings must be treated as hypothesis-generating rather than confirmatory; none would survive conservative Bonferroni correction, and confidence intervals are wide. With these caveats in place, we note the following descriptive patterns: female respondents were descriptively less likely to agree that classmates forgo office hours due to AI ($M_{\text{female}} = 3.27$ vs. $M_{\text{male}} = 3.94$, $p = 0.010$, $d = 0.52$) and that classmates forgo lectures ($M_{\text{female}} = 2.70$ vs. $M_{\text{male}} = 3.28$, $p = 0.023$, $d = 0.48$). They reported substantially stronger understanding of AI rules in their courses ($M_{\text{female}} = 4.70$ vs. $M_{\text{male}} = 4.06$, $p < 0.001$, $d = 0.70$), greater surprise at AI advancement ($M_{\text{female}} = 4.53$ vs. $M_{\text{male}} = 3.91$, $p < 0.001$, $d = 0.66$), and higher career concern ($M_{\text{female}} = 3.97$ vs. $M_{\text{male}} = 3.55$, $p = 0.045$, $d = 0.37$). These exploratory patterns are directionally consistent with ACT's [2025] finding that male students are more optimistic about AI across multiple dimensions; however, the small and imbalanced subsamples preclude strong conclusions. No significant differences emerged between students with and without prior CS/AI coursework on any Likert-scale item, suggesting that formal AI education (at least as captured by a binary coursework indicator) does not substantially differentiate attitudes in this already technologically immersed convenience sample.

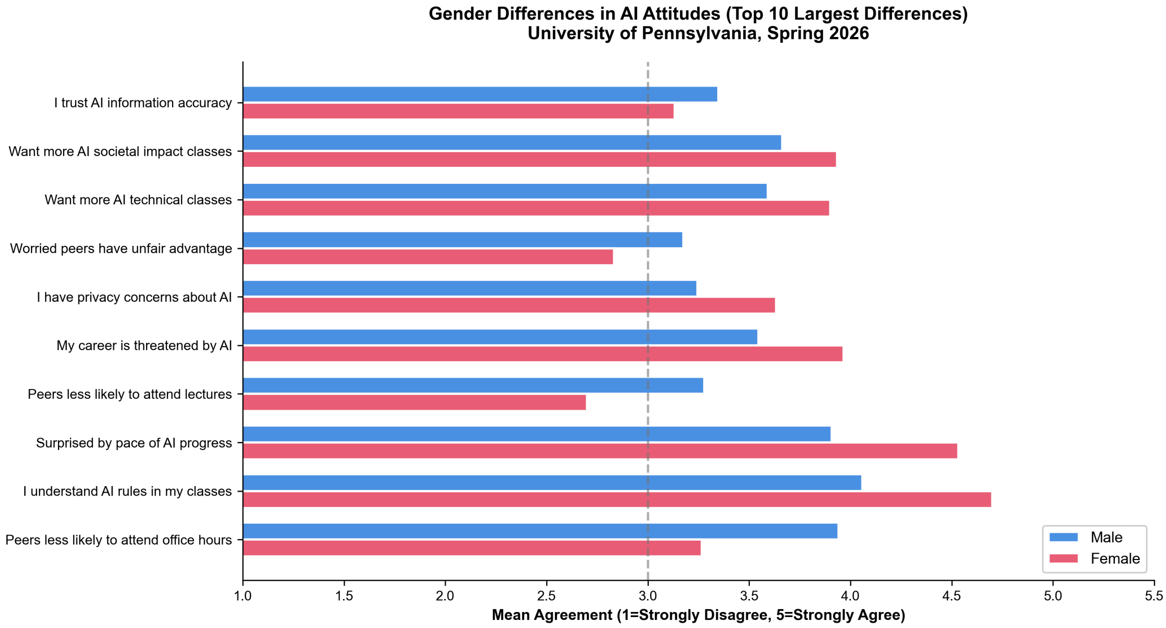

Figure 9. Exploratory gender differences in Likert-scale responses ($N_{\text{male}} = 86$, $N_{\text{female}} = 30$). Only items with nominally significant Welch's t-tests ($p < 0.05$, uncorrected) are displayed. Error bars represent ± 1 SE. All p-values are exploratory; no correction for multiple comparisons was applied. The small and imbalanced female subsample precludes strong inferential claims.

3.8. Correlational Patterns

Pearson correlations among key survey items revealed several notable associative clusters (Figure 10). Perceived academic displacement items were strongly intercorrelated: the perception that peers use AI instead of search engines correlated positively with perceptions of reduced reading ($r = 0.69$), reduced office-hour attendance ($r = 0.58$), and reduced lecture attendance ($r = 0.40$). Career-relevant items exhibited substantial associations: agreement that AI has altered career thinking was positively correlated with desire for more AI classes ($r = 0.66$), expectations of AI surpassing human capabilities by 2056 ($r = 0.64$), and trust in AI accuracy ($r = 0.51$). These correlations are cross-sectional and do not licence causal interpretation; multiple directional accounts are possible (e.g., career concern driving demand for AI education, AI education increasing career concern, or a common latent variable such as general AI engagement driving both). Expectations of AI surpassing human capabilities by 2056 correlated moderately with trust in AI accuracy ($r = 0.58$), desire for more AI classes ($r = 0.53$), and career concern ($r = 0.34$). The AI extinction risk item evidenced modest but consistent associations with most other items, correlating most strongly with career concern ($r = 0.46$) and unfair advantage worry ($r = 0.37$). This pattern is consistent with the interpretation that existential AI concern in this population may be more tightly coupled to personal and academic impacts than to abstract technological forecasting.

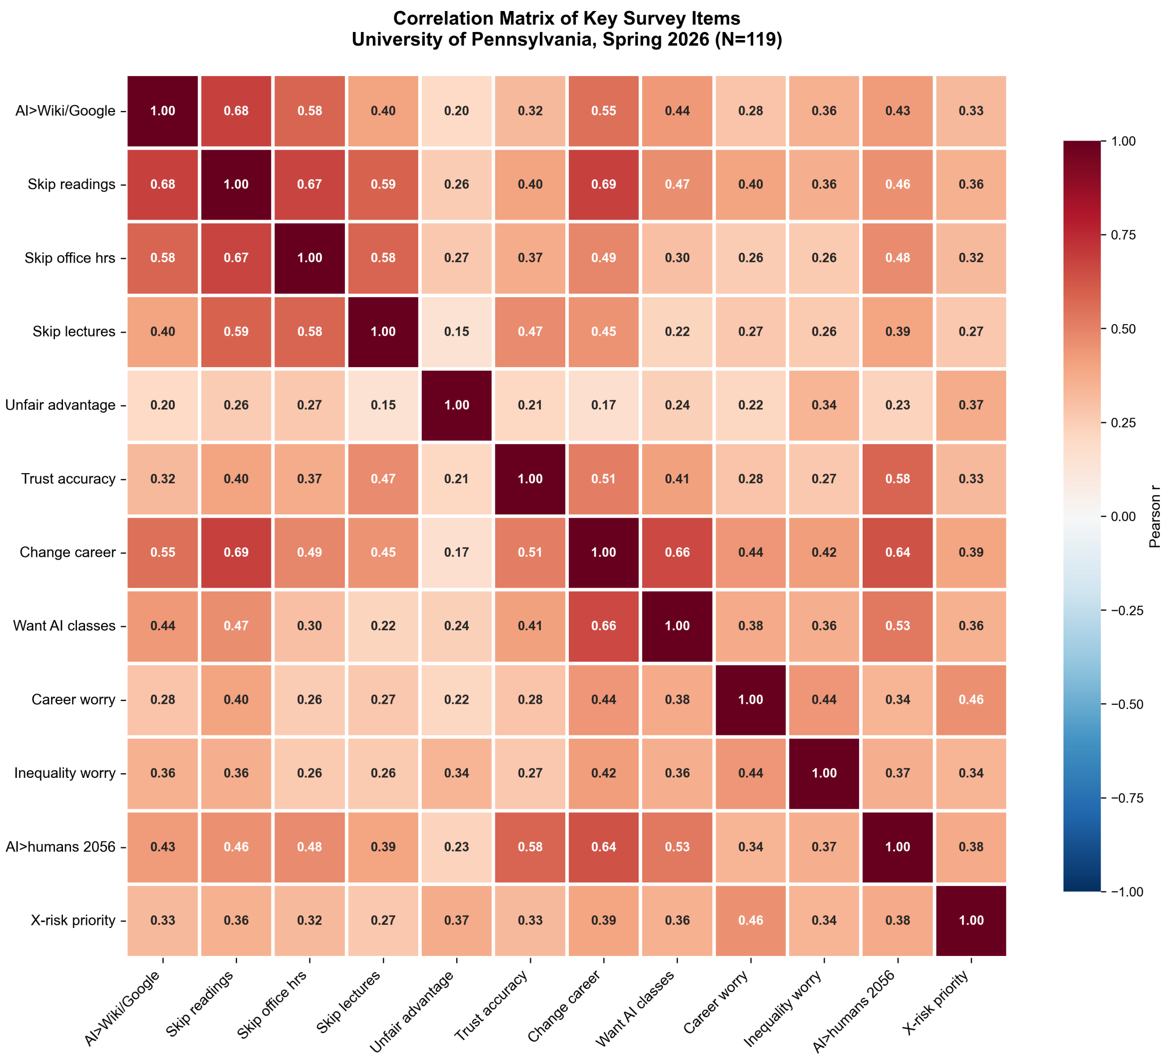

Figure 10. Pearson correlation matrix of key survey items ($N = 119$). Colour intensity reflects correlation strength (red = positive, blue = negative). All correlations are associative and cross-sectional; no causal inference is implied. Only $|r| > 0.3$ is displayed for visual clarity.

3.9. Illustrative Qualitative Observations

The observations that follow are drawn from open-ended survey responses and must be treated as illustrative anecdotes rather than systematically derived themes. No formal intercoder reliability protocol was employed; categorisation was performed through close reading and discussion by the research team.

Open-ended responses revealed a heterogeneous policy landscape across Penn courses. Students described a spectrum spanning from "largely total prohibition" to "encouraged to use, but not turn in as a final product" to "free to use as long as you quote it." Several respondents noted temporal evolution in institutional posture: "Over my 3 years here, [I've] seen a bit of moderation develop over time, with some feeling conflicted about using it for coursework and concerns that it is putting holes in their foundational knowledge in the field." One respondent observed that "a good 2/3 of the classes I've been to are more against using AI," while another characterised Penn's approach as one that "hasn't caught up to just the extent that students use it." These observations suggest a perceived gap between institutional AI policy and student AI practice, although we cannot assess the prevalence or representativeness of these views from the present data.

When asked what generative AI struggles with most, responses clustered around hallucination and inaccuracy ("being 100% accurate," "citing its sources and not hallucinating"), the handling of context and extended interactions ("it struggles with remembering details over long conversations"), and domain-specific limitations in mathematics, "low level programming," and "complex projects requiring large context window." Students also noted stylistic limitations: one described AI output as "still sound[ing] super annoying and obnoxious, always writes in the same manner."

Hopes and concerns regarding AI's future reflected a dual narrative of optimism and anxiety among these illustrative responses. Hopes centred on scientific and medical advancement ("my biggest hopes are in the medical data field"), productivity gains ("hope it grants me back more time"), and human flourishing ("it will allow humans to practise their passions and live dignified, comfortable, and fulfilling lives"). Concerns clustered around job displacement ("I hope it doesn't delete jobs for entire industries"), economic inequality ("concern the power of AI doesn't get concentrated behind paywall"), and authenticity ("I hope the word choice of AI will become less yes-man like and more genuine/honest"). One respondent articulated a vision of careful integration: "Hopes that it is integrated into society with human advancement at the forefront of its goals." When asked about desired institutional resources, students requested AI literacy education, researcher talks, and builder events, reflecting the sample's orientation toward hands-on AI engagement at a builder-focused gathering.

4. Discussion

This exploratory survey of 119 students attending a Penn AI builder community event reveals a campus microcosm in which generative AI is perceived to have achieved near-total integration into daily academic life: universal peer awareness, daily usage estimated by the overwhelming majority, and widespread willingness to expend personal funds on premium access. These perceived adoption rates are substantially higher than those reported in larger, more representative surveys employing self-report measures. The HEPI study documented that 88% of UK undergraduates had self-reported using AI for assessments [3]; the UCLA Senior Survey found that 73% had self-reported using AI to support coursework [4]; and Baek et al. [1] reported that ChatGPT use among 1,001 U.S. college students was widespread but varied substantially by discipline and institutional context. Our finding of 81% perceived daily peer usage and 96% ChatGPT awareness almost certainly reflects the combined effects of a genuinely high-adoption subpopulation, self-selection into an AI builder event, and measurement via peer-observation rather than self-report. We emphasise that these estimates cannot be generalised to the broader Penn student body and represent perceived upper bounds among the most AI-engaged undergraduates.

The descriptive gap between perceived near-universal peer adoption and minority self-reported trust in AI accuracy (46%) is consistent with the emerging student GenAI literature. Cabrera et al. [2] explicitly characterised student GenAI engagement as "widely used but barely trusted" and interpreted this as evidence that students maintain critical distance from AI outputs even as they depend on them extensively. Our data are compatible with this interpretation but admit multiple accounts. The adoption and trust measures operate at different levels of analysis (perceived peer behaviour vs. self-reported attitudes) and address distinct constructs (frequency of tool utilisation vs. perceived information accuracy). This measurement asymmetry entails that the observed gap could partly reflect genuine pragmatic scepticism, partly habitual use despite acknowledged unreliability, and partly the difficulty of directly comparing peer-perception and self-report instruments. Longitudinal designs incorporating matched self-report and peer-perception measures, ideally supplemented by behavioural data, would be required to adjudicate among these interpretations.

The exploratory gender differences we observe are directionally consistent with prior work but must be interpreted with considerable caution, given the small female subsample ($n = 30$), the imbalanced design, and the absence of correction for multiple comparisons. The ACT [2025] survey of high school students found that males were significantly more optimistic about AI across all dimensions, while the HEPI [2025] survey reported a 14-percentage-point gap in AI usage favouring male students. Our data hint at a more granular pattern: female respondents in our sample are descriptively less likely to perceive AI-mediated disengagement from academic support structures while being more concerned about AI's career implications and more surprised by the pace of AI advancement. One possible interpretation—offered here as speculation requiring replication—is that female students in this AI-engaged cohort maintain stronger commitments to traditional academic support structures and are more alert to AI's societal-level implications, even as they are equally or more engaged with the technology. This pattern may reflect broader gender differences in technology risk perception [9] rather than AI-specific dynamics. Larger, balanced samples employing validated instruments are required before substantive conclusions can be drawn.

The perceived career anxiety findings join a growing literature documenting AI-related labour-market concern among students. The 74% agreement rate for "AI has changed the way I think about my future career" is among the highest in our survey, and the 60% who express concern about negative career effects is notable even within this elite institutional context. These figures are broadly consistent with broader societal polling: Pew Research Center [2024] reported that 62% of Americans believe AI will exert a major impact on jobholders in the next two decades, and ACT's [2025] high school survey documented significant job-displacement concern. What may distinguish our sample is the strong association between career concern and demand for AI-related institutional education ($r = 0.38$ for career worry with desire for AI classes; $r = 0.66$ for changed career thinking with desire for AI classes). This association is consistent with the interpretation that, among this AI-engaged cohort, AI literacy education is perceived as a mechanism for managing career uncertainty. We cannot, however, exclude alternative explanations: students who are generally more anxious may both worry about AI more and desire more coursework across domains; the builder-event context may select for individuals who view AI education as intrinsically valuable; and the cross-sectional design precludes establishing that career concern drives educational demand rather than the reverse, or that both are driven by a common cause.

The perceived displacement of traditional academic behaviours (search, reading, office hours) raises important questions about AI's effect on learning depth that our data can illuminate but not resolve. If respondents accurately perceive that 68% of peers read less because AI summarises texts, and 66% forgo office hours in favour of AI consultation, then the aggregate effect on learning quality may be substantial. This pattern echoes concerns regarding "cognitive offloading" in the context of large language models [10]. Three important caveats apply, however. First, our data capture perceived displacement, not verified behavioural change; social projection and availability heuristics may inflate these estimates. Second, AI may displace shallow engagement (skimming summaries of texts that would not have been read carefully in any event) rather than deep engagement, and our data cannot distinguish these scenarios. Third, the correlational structure of these items (all four displacement items loading together) may partly reflect a general "AI displacement schema" rather than independent behavioural assessments. Direct behavioural measurement (e.g., learning management system engagement data, library usage metrics, office-hour attendance records) would be required to validate these perceptions.

The perceived premium subscription data, while striking (85% report peer expenditure, with 13% above $50 per month), warrant cautious interpretation. These are estimates of peer behaviour, not verified expenditure, and may be inflated by projection or salient exemplars. Nonetheless, if even directionally accurate, they raise equity concerns: differential access to premium AI tools predicated on ability to pay represents a potential new dimension of the digital divide in higher education. This concern was echoed by a respondent who worried about "the power of AI [getting] concentrated behind paywall." Prior scholarship on the digital divide in higher education has focused on hardware, broadband, and software access [11]; AI premium subscriptions may constitute an emergent tier of inequality that institutions will need to consider as part of their technology equity strategies. Verified expenditure data from representative samples would be needed to substantiate this concern.

5. Limitations

This study is subject to several important limitations that constrain the generalisability, internal validity, and robustness of its findings. We enumerate these here, distinguishing—where germane—between limitations that can be addressed through textual qualification and those that require additional data collection.

Sample selection and size. The sample is self-selected from attendees of a campus AI builder community event, yielding a population that is overwhelmingly male (72%), Asian (69%), and CS-literate (67%). These demographics are not representative of Penn's broader student body (approximately 55% female, 27% Asian), let alone the U.S. higher-education landscape. Our estimates of perceived AI adoption, premium subscription rates, and daily usage frequency almost certainly represent upper bounds relative to what would be observed in a representative sample of the same institution. The modest sample size ($N = 119$) further restricts statistical power and precludes multivariate analyses that could disentangle the effects of gender, CS background, and class year. The female subsample ($n = 30$) is especially small, rendering all gender comparisons exploratory. We are also unable to compute a response rate because event attendance was not formally tracked and the survey was not distributed via a closed recipient list. [Requires new data collection to address: representative sampling, larger and balanced subgroups, and tracked denominator.]

Construct validity of peer-observation measures. Our adoption, usage frequency, and academic displacement items capture perceived peer behaviour, not verified behavioural data. Peer-observation measures are susceptible to social projection (assuming others behave as one does), availability heuristics (overweighting salient exemplars), and pluralistic ignorance (misperceiving group norms). Although perceived norms independently influence behaviour [6] and may attenuate social desirability bias relative to self-report, they cannot substitute for direct behavioural measurement. All claims about "peer adoption" or "peer usage" must be interpreted as claims about perceived peer behaviour. [Unfixable without behavioural measures such as learning management system logs, direct observation, or diary studies.]

Measurement-level mismatch. Our central interpretive construct, the perceived adoption–trust gap, compares perceived peer adoption (measured via peer-observation items) with self-reported trust (measured via a single self-report Likert item). These operate at different levels of analysis and address distinct constructs; the observed gap may partially reflect this measurement asymmetry rather than a genuine psychological paradox. We have recast this as a descriptive observation and enumerated alternative interpretations in the Discussion. [Fully resolving the question would require matched self-report and behavioural measures administered in the same instrument.]

Single-item measurement. Each construct was assessed with a single Likert item, which precludes assessment of internal consistency, measurement error, or latent variable structure. Future iterations of this work should incorporate validated multi-item scales such as the General Attitudes toward Artificial Intelligence Scale [12,13] or domain-specific scales developed for educational AI contexts. [Unfixable without instrument redesign and re-administration.]

Cross-sectional design. The survey captures a single temporal snapshot (February 2026). The velocity of AI development and adoption is such that findings may date rapidly. Longitudinal tracking would be necessary to distinguish durable attitudinal patterns from transient reactions to specific product releases or media cycles, and to establish temporal precedence for the correlational associations reported herein. [Unfixable without longitudinal follow-up surveys.]

Exploratory statistical approach. The multiple comparisons conducted across 20 Likert items and two grouping variables increase the risk of Type I error. No correction for multiple comparisons was applied. The reported p-values must be interpreted as exploratory and hypothesis-generating; none would survive conservative Bonferroni correction. Effect sizes (Cohen's d) are reported to facilitate interpretation, but confidence intervals are wide owing to small and imbalanced subgroup sizes. [Fully addressing this limitation would require pre-registered hypothesis testing with larger, balanced samples.]

Qualitative analysis limitations. Open-ended responses were categorised through close reading and research team discussion without a formal coding framework, intercoder reliability assessment, or systematic thematic analysis protocol. Qualitative observations must be treated as illustrative anecdotes that enrich the quantitative findings rather than as systematically derived themes. [Rigorous qualitative analysis would require a structured, multi-coder protocol applied to existing or expanded textual data.]

Item ambiguity. Several Likert items may admit heterogeneous interpretations across respondents. "Trust AI-generated information to be accurate" does not specify domain, model, or task type. "AI has changed the way I think about my future career" captures a broad construct that could encompass optimism, pessimism, or strategic recalibration without distinguishing among these. [Unfixable without item revision and re-administration.]

Missing variables. We lack data on several constructs that would substantially enrich interpretation: actual (rather than perceived) AI usage frequency, academic performance metrics, socioeconomic background beyond a binary financial aid indicator, and prior AI attitudes before college matriculation. [Unfixable without a redesigned survey instrument.]

IRB status. Institutional Review Board approval was pending at the time of writing. The final manuscript should include the IRB determination and protocol number once available. [Fixable: add determination when received.]

6. Conclusion

This exploratory survey of 119 students attending a Penn AI builder community event provides a granular, perception-focused portrait of generative AI's perceived integration into campus life among AI-engaged undergraduates at an Ivy League university in early 2026. Within this self-selected, technology-forward cohort, perceived near-universal peer adoption, daily usage, and substantial premium expenditure coexist with widespread self-reported distrust of AI accuracy, a descriptive pattern consistent with the emerging literature on student "pragmatic scepticism" toward AI tools [2]. Perceived career anxiety is pervasive and is positively associated with interest in expanded institutional AI education within this sample. Exploratory gender differences in AI perceptions hint at divergent calibrations of AI's academic and career impacts, although these findings require replication in larger, balanced samples before substantive conclusions can be drawn. As universities navigate the rapid evolution of AI capabilities, the student perspectives captured here, however circumscribed in generalisability, underscore interest in AI literacy curricula that address not only technical proficiency but also critical evaluation, career navigation, and the ethical dimensions of an AI-augmented professional future. Future research should prioritise representative sampling strategies, matched self-report and behavioural measures, longitudinal designs capable of establishing temporal precedence, and validated multi-item instruments to build upon the suggestive patterns identified in this exploratory investigation.

Author Contributions: Conceptualization, F.A.C. and A.S.; methodology, F.A.C. and A.S.; software, F.A.C.; formal analysis, F.A.C.; investigation, F.A.C. and A.S.; data curation, F.A.C.; writing—original draft preparation, F.A.C. and A.S.; writing—review and editing, F.A.C. and A.S.; visualization, F.A.C.; project administration, F.A.C. and A.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The study was reviewed by the University of Pennsylvania Institutional Review Board [determination pending]. All participants provided informed consent at the commencement of the survey.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The anonymised dataset and analysis code are available from the corresponding author upon reasonable request.

Acknowledgments: We thank the attendees of the Penn Claude Builder Club event for their participation and candid responses. This survey was designed and administered by F.A.C. and A.S. The analysis was conducted collaboratively by the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Baek, C.; Tate, T.; Warschauer, M. "ChatGPT seems too good to be true": College students' use and perceptions of generative AI. *Computers and Education: Artificial Intelligence* **2024**, *7*, 100294. https://doi.org/10.1016/j.caeai.2024.100294.
2. Cabrera, C.; Cooper, A.; Sapsed, J. Widely used but barely trusted: Understanding student perceptions on the use of generative AI in higher education. *International Journal of Inclusive Education* **2025**. https://doi.org/10.1080/13603108.2025.2595453.
3. Freeman, J. Student Generative AI Survey 2025. Technical report, Higher Education Policy Institute (HEPI), 2025.
4. UCLA Teaching & Learning Center. Generative AI use & perspectives from the Class of 2025. https://teaching.ucla.edu/news/ucla-student-ai-use-perspectives/, 2025.
5. Schiel, J. How students' perspectives on AI are related to AI training in high school, planned college major, and gender. Technical Report R2520, ACT, 2025.
6. Cialdini, R.B.; Goldstein, N.J. Social influence: Compliance and conformity. *Annual Review of Psychology* **2004**, *55*, 591–621. https://doi.org/10.1146/annurev.psych.55.090902.142015.
7. Pew Research Center. How Americans view artificial intelligence. https://www.pewresearch.org/short-reads/2024/11/21/how-americans-view-artificial-intelligence/, 2024.
8. Stein-Perlman, Z.; Weinstein-Raun, B.; Grace, K. 2023 Expert survey on progress in AI. https://aiimpacts.org/2023-expert-survey-on-progress-in-ai/, 2024.
9. Henwood, K.L.; Parkhill, K.A.; Pidgeon, N.F. Science, technology and risk perception: From gender differences to the effects made by gender. *Equal Opportunities International* **2008**, *27*, 662–676. https://doi.org/10.1108/02610150810916730.
10. Risko, E.F.; Gilbert, S.J. Cognitive offloading. *Trends in Cognitive Sciences* **2016**, *20*, 676–688. https://doi.org/10.1016/j.tics.2016.07.002.
11. Van Dijk, J. *The Digital Divide*; Polity Press, 2020.
12. Schepman, A.; Rodway, P. The General Attitudes towards Artificial Intelligence Scale (GAAIS): Confirmatory validation and associations with personality, corporate distrust, and general trust. *International Journal of Human–Computer Interaction* **2023**, *39*, 2724–2741. https://doi.org/10.1080/10447318.2022.2085400.
13. Sindermann, C.; Sha, P.; Zhou, M.; Wernicke, J.; Schmitt, H.S.; Li, M.; Sariyska, R.; Stavrou, M.; Becker, B.; Montag, C. Assessing the attitude towards artificial intelligence: Introduction of a short measure in German, Chinese, and English language. *KI – Künstliche Intelligenz* **2021**, *35*, 109–118. https://doi.org/10.1007/s13218-020-00689-0.