Beyond AI Adoption: Trust as a Task-Level Appraisal in Academic LLM Use

Yvonne Chang
UCL
66-72 Gower Street
London, United Kingdom
WC1E 6EA
yvonne.chang.23@ucl.ac.uk

Jonathan Mella
UCL
66-72 Gower Street
London, United Kingdom
WC1E 6EA
jon.mella@ucl.ac.uk

Anna Dowthwaite
UCL
66-72 Gower Street
London, United Kingdom
WC1E 6EA
anna.rudnicka.15@ucl.ac.uk

Yoana Ahmetoglu
UCL
66-72 Gower Street
London, United Kingdom
WC1E 6EA
yoana.petrova.15@ucl.ac.uk

Anna L. Cox
UCL
66-72 Gower Street
London, United Kingdom
WC1E 6EA
anna.cox@ucl.ac.uk

Large language models (LLMs) have been positioned as tools that could reduce workload and support productivity in academic work, yet adoption remains selective. Prior research attributes this to differences in trust but offers limited insight into how trust is enacted in practice. We report a qualitative study with 18 UK academics, examining how they decide when and how to trust LLMs in their work. Trust is not a stable attitude but a situated, task-level appraisal, governed primarily by whether outputs can be verified and at what cost. We identify verification burden to describe how, in academic knowledge work, the effort required to validate LLM outputs can exceed the benefit of using them, making restraint the reasonable choice. LLMs also serve as emotional infrastructure for managing the psychological demands of academic work, while creating tensions with academic identity. At the institutional level, we identify institutional paralysis: unclear policies, absent norms, and a lack of endorsed tools constrain adoption even among willing users. Together, these reframe cautious LLM use as a nuanced interplay of constrained trust calibration, emotional ambivalence, and institutional paralysis. We translate these findings into design concepts that make verification effort legible and treat emotional support as core function, as well as describing implications for institutional policy.

generative AI, trust appraisal, academic work, AI adoption, knowledge work, llm

1. INTRODUCTION

Academic work is increasingly characterised by intense workloads, administrative burden, and sustained pressure to produce high-quality research outputs (Naidoo-Chetty and Du Plessis 2021). These demands have been linked to rising levels of stress, burnout, and reduced wellbeing among academics (Singh and Strzelecki 2026). In parallel, large language models (LLMs) have rapidly emerged as tools that could alleviate aspects of this workload (Mella et al. 2024) by supporting tasks such as writing, summarisation, and even research question generation (Liu et al. 2023). While students have widely adopted these tools, academic researchers have been comparatively more selective in their uptake (Mhaidli and Roemmich 2024; Ahmetoglu et al. 2026). Recent survey data suggest that only around half of academics regularly use LLMs for research and teaching tasks (Arroyo-Machado et al. 2025). This raises a question that existing literature has not adequately answered: why do academics adopt LLMs so selectively, and what shapes those decisions?

Many models exist to help predict the adoption of technology within groups, notably the Technology Acceptance Model (TAM) (Davis 1989) and the Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh and Thong 2016), yet trust occupies an ambiguous position in both. Despite being a key factor in AI adoption (Choung et al. 2023a), UTAUT does not include it as an explicit construct, absorbing it instead through core constructs (Kenesei et al. 2025) — an approach that

© Chang et al., Published by BCS Learning and Development Ltd.

made sense for predictable, rule-based systems but breaks down for LLMs, which behave inconsistently and require continuous output evaluation (Lee et al. 2025; Galindez-Acosta and Giraldo-Huertas 2025). Models that do address trust (Gefen et al. 2003; Choung et al. 2023a) treat it as a general orientation toward a system rather than something renegotiated task-by-task, overlooking its cognitive and affective dimensions and its contextual, situated character (Bach et al. 2024). What is missing is an account of how trust is actually enacted in practice during academic work, at the level of granularity where adoption decisions are made.

This paper sets out to address this gap. We conducted a survey study with 18 academics across UK institutions, asking: **How do academics decide when and how to trust LLMs in their work?** Rather than treating trust as a prior attitude toward a system, we examine how it is evaluated, calibrated, and applied at the level of individual tasks. Our findings show that academics are not asking whether they trust LLMs in general, but whether a specific output can be trusted enough to act on, given the effort required to verify it, the consequences if it is wrong, and the institutional context in which it is being used. This reframes cautious attitudes toward LLMs not as a failure of adoption or a deficit of trust, but as a rational outcome of constrained trust calibration in a domain where the costs of being wrong are asymmetrically high.

Building on accounts of trust as dynamic and situated (Hoffman et al. 2013; Lee and See 2004), this work makes four contributions that together reframe how we understand selective LLM adoption in academia. Firstly, we provide support for the existing account that trust, rather than a global attitude toward a system, is a situated, task-level calculation. Secondly, we provide empirical grounding for task-level trust calibration in academic LLM use, revealing the mechanism that drives it: a verification burden that functions not merely as a design challenge (Nogueira et al. 2026) but as a structural explanation for why restraint is a reasonable response to a genuine cost-benefit imbalance. Thirdly, we show that this restraint coexists with a degree of emotional reliance, as academics simultaneously depend on LLMs for cognitive and psychological relief while experiencing them as a threat to the intellectual process that makes academic work meaningful. Finally, we show that even where individuals are willing to engage, responsible adoption is structurally constrained by ambiguous, unaligned institutional conditions. We conceptualise this as *institutional paralysis*: a systemic constraint that shifts the locus of the problem from the individual to the organisation.

2. RELATED WORK

2.1. Technology Adoption and Trust in AI

When information technology was introduced into the workplace in the late 1980s, early research focused on understanding why users were hesitant to adopt systems despite their potential efficiency gains (Davis 1989). TAM explained adoption through perceived usefulness and ease of use, and was later extended into the UTAUT, which incorporates social influence and facilitating conditions (Venkatesh and Thong 2016). While UTAUT remains a dominant model for workplace technology adoption (Lee et al. 2025), research on AI highlights trust as a critical additional factor, particularly for systems operating under uncertainty (Choudhary and Yamuna 2026). Although some extensions of TAM incorporate trust (Gefen et al. 2003), it is not included as an explicit construct in UTAUT, where trust-related concerns are integrated within performance expectancy and facilitating conditions (Kenesei et al. 2025). This omission mattered little for predictable, rule-based systems, but AI systems differ in that their outputs are probabilistic rather than deterministic, potentially requiring users to continuously calibrate trust based on performance and reliability over time.

Trust in computer systems, defined as the extent to which users believe an AI system will reliably and accurately support their goals in conditions of uncertainty, helps determine whether academics integrate LLMs into their work. Prior work establishes trust as a multifaceted construct that enables users to act under uncertainty, and that fluctuates in response to system performance and situational pressures (Lee and See 2004; Mayer et al. 1995; Hoffman et al. 2013). Yet despite this body of work, existing accounts of trust in AI operate at the level of a user's general orientation toward a system, leaving largely unexamined how trust is negotiated on a task-by-task basis within a single workflow.

The adoption of AI technologies introduces distinct trust-related challenges due to the opacity of machine learning systems and the increasingly anthropomorphic presentation of tools such as LLMs. As outputs from systems like ChatGPT are fluent and human-like, trust-in-human scales can better predict trust in such technologies compared to trust-in-machine models (Lankton et al. 2015). This human-like behaviour of AI can elevate user expectations, increasing the risk of overtrust (Glikson and Woolley 2020). At the same time, trust becomes especially critical in contexts characterised by high uncertainty and limited user knowledge (Nienaber and Schewe 2014) — conditions that characterise LLM use particularly well, given that users often cannot assess the accuracy of outputs in domains

where they lack expertise. However, misalignment between user trust and system capability—often referred to as miscalibration—can lead to both overreliance and underuse (Hoffman et al. 2013).

2.2. The Value of LLMs for Academics

In recent decades, the university landscape has seen a significant shift, with some arguing it has moved from a culture of collegiality toward a more bureaucratic, managerially-driven model (Trow 1994). To quantify success, bureaucratic processes are established, propelling longer work hours, as well as the publish or perish culture (Miller et al. 2011). These pressures have transformed academic work, with many academics working over 50 hours per week and experiencing poor health and well-being (Smith 2024; Sang et al. 2015). This is increasingly burdensome for academics who often face conflicting demands from teaching, research, and administrative tasks, leading to role overload and burnout (Doherty and Manfredi 2006; Jensen and Olsen 2023).

LLMs offer academics potential relief from heavy workloads by supporting tasks across the research process and improving work–life balance. Prior work suggests that LLMs can assist at all stages of research, including idea generation, literature synthesis, data preparation, analysis, and writing (Cargill and O'Connor 2009). For example, tools like CoQuest demonstrate how LLMs can support research question development through human–AI co-creation (Liu et al. 2023). LLMs are also widely used for writing support, particularly benefiting non-native English speakers as an accessible alternative to paid editing services (Kim 2023). Beyond writing, they can assist with tasks such as deductive coding, significantly reducing manual effort (Chew et al. 2023). Together, these applications highlight the potential of LLMs to streamline academic work and enhance research productivity.

2.3. Concerns and Barriers to LLM Adoption in the Academy

However, the utility of LLMs across the full research process remains highly debated. While some studies find that tools like ChatGPT are well-suited for early-stage idea generation, leveraging their ability to explore and connect broad concepts, they struggle with more specialised research tasks such as literature review and framework testing, where coherent and rigorous reasoning is required (Dowling and Lucey 2023). Their use has also raised concerns around authorship, with some researchers controversially listing ChatGPT as a co-author (Rahman et al. 2023). At the same time, there remains a lack of clear institutional policies to guide their responsible adoption. While some argue that generative AI should be embraced across research contexts, the growing ecosystem of LLM tools places increasing responsibility on institutions to support researchers through guidance, resources, and frameworks for evaluating and using these systems effectively (Liu and Jagadish 2024; Ahmetoglu et al. 2026).

Perhaps most critically, the reliability of AI-generated outputs remains a central concern in academic contexts. Prior work highlights that some users often accept LLM-generated content without thorough verification due to fluency and automation biases, increasing the risk of factual inaccuracies entering downstream work (Gould et al. 2024; Thompson et al. 2013). This issue is further exacerbated by the tendency of LLMs to hallucinate, which can lead to the unintentional propagation of misinformation and, in turn, erode trust among stakeholders (Joshi 2025). Yet how researchers navigate these reliability concerns in practice, and at what cost, remains poorly understood.

2.4. Summary

Prior work highlights trust as central to AI adoption but treats it as a generalised belief rather than an ongoing process enacted in practice. Moreover, existing context-dependent accounts of trust lack explanatory power in academic settings. Thus, rather than asking *whether* academics trust LLMs, we ask how they decide *when and how* these tools can be relied upon in specific tasks. This paper supports a context-dependent account of AI trust and identifies its mechanism (verification burden, including its expertise dependency) and its conditions (institutional paralysis) in academic knowledge work.

3. METHOD

3.1. Participants

Participants were recruited through a UK university mailing list and social media between June and July 2024. To be eligible, participants had to be at least 18 years old, a research student, researcher, or member of staff at a university or research institution, and have prior LLM experience, which they self-reported during screening. All participants received a £10 Amazon voucher for their time.

A total of 57 responses were excluded from the final dataset due to factors including: duplicates, not registering with an email affiliated with an academic institution and insufficient survey responses. This resulted in a final sample of 18 academics (11 males, seven females) for data analysis. The mean

age was 35.3 years (*SD* = 12, range = 24-65). Participants were drawn from eight UK institutions, predominantly UCL (n = 10), and spanned a range of career stages from PhD students (n = 8) to professor (n = 1). Disciplines were similarly varied, covering psychology and neuroscience (n = 7), computer science (n = 4), biological sciences (n = 2), English literature (n = 2), and several others. The majority held combined research and teaching responsibilities, with four focused on research only. Refer to Appendix C for the specific demographic information.

3.2. Materials

The survey captured quantitative and qualitative data on participants' experiences and attitudes toward LLM use. Written survey responses were chosen over synchronous interviews since academics are time-poor, and asynchronous responses reduce scheduling burden while allowing reflection without the pressure of live conversation. A 200-character minimum per question ensured substantive responses.

Participants completed the Information Sheet and consent form, then provided demographic information and rated eight statements on a six-point Likert scale capturing practical integration, individual attitudes, and institutional roles (See Appendix C and D). They then completed 13 open-ended questions prompting reflection on their experiences, decision-making, and areas for improvement (Appendix A), covering topics such as using LLM tools despite expecting low-quality output, how tools could better support service, administrative, and teaching tasks, and which work tasks they currently used LLM tools for and how effective these were.

3.3. Procedure

Ethical approval for this study was obtained from UCL's Ethics Committee, ensuring that all procedures adhered to institutional guidelines for research involving human participants. Participants began the study by reviewing the information sheet and consent form. Once they provided informed consent, they were redirected to the Qualtrics survey via a secure link. Upon completing the survey, participants received a debriefing message summarizing the study's purpose. Once data collection was finalised, researchers contacted participants via email to distribute Amazon voucher codes as compensation for their time.

3.4. Data Analysis

All survey data were anonymised prior to analysis. Demographic and Likert-scale responses were analysed quantitatively in Excel to produce descriptive statistics summarizing participant characteristics and attitudes toward LLM use.

The open-ended responses were analysed qualitatively using Braun and Clarke's six-step thematic analysis methodology, following a hybrid inductive-deductive approach (Braun and Clarke 2023). This process involved 1) familiarization with the data, 2) generation of initial codes, 3) searching for themes, 4) reviewing themes, 5) defining and naming themes, and 6) producing the final thematic report. Steps one and two were completed independently by the first author. The theme names and definitions were created first independently by the first author and then collaboratively iterated on with the remaining authors. Consistent with Braun and Clarke's position that inter-rater reliability is not an appropriate quality criterion for reflexive thematic analysis, validity was instead sought through researcher reflexivity and collaborative review of theme definitions across the author team.

The hybrid analysis approach involved two distinct phases of data coding (Swain 2018). This hybrid design was chosen because, while established frameworks exist for trust in AI and workplace technology, studies specifically examining LLM trust in academic contexts remain scarce. Combining both approaches allowed for data-driven insights alongside connection to existing theory.

In the first phase, an inductive approach was used, with codes and themes developed from close engagement with participant responses rather than derived from a predetermined framework. In the second phase, a deductive approach was applied, drawing on established frameworks for trust in AI and automation (Lee and See 2004; Hoffman et al. 2013; Mayer et al. 1995). These frameworks were used to situate inductively generated codes within existing theory, particularly around the dynamic and situated character of trust and the conditions under which it breaks down. After completing both phases of coding, results were synthesised to identify overarching themes aligned with the study's research objectives.

4. RESULTS

All participants had prior experience of using LLMs for work tasks. Our survey results (See Appendix D) indicate that nine participants regularly used LLM tools for work and felt confident doing so, describing them as valuable for saving time and alleviating workload pressures. However, the majority (n=12) expressed a desire for more institutional guidance on what is permitted, and many expressed concern

about how these tools might shape the future of academic work.

Regarding the thematic analysis, the first theme suggests that trust in LLMs is experienced as a situated judgment remade at the task level. The second theme illustrates that verifiability is the key mechanism underlying these situated trust judgements. The third theme explores how academics extend a form of trust to LLMs as emotional and cognitive scaffolds, even as this creates tensions with academic identity. Finally, we show how institutional ambiguity undermines the conditions necessary for trust decisions to be made at all, creating structural paralysis that constrains adoption even among academics otherwise willing to engage.

4.1. Theme 1: Trust as a Task-by-Task Calculation, Not a Global Attitude

Academics in our study did not hold one single settled view of whether LLMs could be trusted. Instead, participants described their trust calibration independently for each task, shaped by what was at stake, how easily the output could be checked, and how much the academic already knew about the subject. For example, the same participant who trusted an LLM to draft an email would not trust it for planning or teaching activities:

> *"I trust LLMs to draft emails and come up with plans and templates. This is because I know that I can review the information within and make adjustments where needed. I would not trust them to come up with more concrete plans or strategic objectives or learning outcomes."* (P7)

Even participants who had barely used LLMs employed the same logic. P8 articulated the underlying principle concisely: *"I'd be fine with using them for things I can check, as this requires less trust."* P13 expressed similar logic: *"only when I already know the answer so I can verify whether the response is accurate. I don't trust these tools in any other scenario."* This pattern appeared not only among light users but also among those with more established practices, where calibration operated through domain expertise rather than verifiability alone. P2, a professor with detailed subject knowledge, described using LLMs specifically on questions where they already knew the terrain: ''*I use ChatGPT when I already know the answer and wish to check if the LLM shows up anything extra I may have forgotten."* The tool is used only where the academic can evaluate the output.

Participants' accounts described a process of active calibration rather than a simple shift in trust over time. Several narrowed their use to specific tasks after early experiences revealed the tools' limitations, and some developed strategies to constrain outputs when they overstepped. P6 described adding a second prompt when an initial response changed their writing in ways that did not make sense, redirecting the tool to show them what could be improved rather than doing it for them — a practice that, over time, they felt had genuinely helped develop their own writing. P14 described setting clear boundaries after consulting colleagues: *"I decided to limit my use of LLM tools to specific, well-defined tasks."* P2 arrived at the same position through direct testing: *"I would only consult an LLM on questions where I feel knowledgeable enough to judge the veracity of what is produced."*

Even participants who quantified their trust explicitly framed it in task-relative terms: P17 estimated their trust at *"50-60 per cent,"* but qualified this immediately by noting it depended entirely on the type of task, with proofreading and error-checking sitting at one end and generating new text from scratch at the other. P5 captured the resulting disposition succinctly:

> *"I trust them about as much as I'd trust a keen but inexperienced intern. The variants that supply supporting references with their output are slightly more trustworthy, though even then I'm never sure if it's missing something that a well-trained professional human would not."* (P5)

The analogy does more than signal caution. It frames LLMs as fallible collaborators whose outputs warrant oversight proportional to the task's stakes, given work where the cost of checking is manageable, not where it isn't. That this emerged unprompted suggests academics are already operating with an informal but sophisticated model of calibrated trust.

4.2. Theme 2: The Verification Burden as a Core Trust Barrier

The most consistent factor participants identified when deciding whether to extend trust was whether they could independently verify the output. What emerged from participants' accounts was a precise insight: in academic work, the cost of verifying LLM outputs can exceed the cost of not using the LLM at all. P4 stated this directly:

> *"The least helpful aspect is not knowing if the response is accurate or not. To assess that, I would need to read the original text myself, therefore it would be pointless."* (P4)

Notably, P4 also articulated what a trustworthy output would look like from the other side: a response that presented key quotes with "*minimal editorialising*", so that the source material could be

located and checked without the academic having to reconstruct what the tool had done.

P5 identified source traceability as the condition that allows this trap to be avoided: *"If I can go to the source text and find the origin of those points, then it could be suitable for academic work, but if those are obscure, then it could be just more noise."* P13's formative experience illustrates the cost of this failing: *"I tried to retrieve some quotations from an author but the tool hallucinated the passages. I wasted an entire morning trying to verify the relevant quotations in the primary source and then I realised the tool had made it up."*

This experience is qualitatively different from verification failure in other professional contexts. A developer can run code and know within seconds whether it works. An academic verifying a synthesised literature claim may need to read multiple primary sources over hours or days. But the asymmetry runs deeper than effort alone. P2, a professor with substantial domain expertise, described their early experience of consulting LLMs on topics they knew well and finding *"important knowledge gaps and misleading statements"* in the replies. This points to a second dimension of the burden: not just the time required to verify, but the level of prior expertise needed to detect that something has gone wrong at all. Where that expertise is absent, the failure may not be visible until significant downstream work has already been built on it.

Conversely, when outputs could be verified, trust extended even for tasks where quality was uncertain. P3 trusted LLMs for rephrasing because the source material was their own: *"I trust it for rephrasing words because I feel confident I know where the information is coming from (i.e. me)."* P16 described a consistent practice of maintaining final authority: *"I wouldn't take any LLM output at face value. I will always have the final say on it."* This maintenance of authorial control was not a burden but a precondition of trust. The tasks academics trust LLMs for are systematically those where verification costs are low; those they avoid are where verification costs are high.

4.3. Theme 3: LLMs as Emotional Infrastructure: Between Relief and Identity

For several participants, LLMs were trusted as a form of support in relieving the psychological burden of academic work. This was most visible around writing, where the blank page carries a weight in academia that goes beyond practical inconvenience. P3 described the relief of having something to edit rather than nothing to start from:

> *"Sometimes using the AI is easier because even though I still have to do lots of editing, it's a different task and less emotionally stressful compared to when I'm editing my own writing and trying to write from scratch when I just get completely stuck and more frustrated with myself."* (P3)

P6 described the same dynamic from a slightly different angle, noting that once there was something on the page to respond to, their thinking could start:

> *"...even though [the outputs] are low-quality... my brain starts to work and think about what I actually want to write ... it easier for me to start...because the task is so complex!"* (P6)

P9 described the same pattern, adding that critiquing a poor draft built their own confidence: *"A lot of the time I experience writer's block and LLM can help with this even if the output is of low quality. It also increases my own confidence to be able to distinguish when LLM is doing a suboptimal job."*

This emotional function extended beyond the practical task of writing to encompass professional identity and self-confidence. P17, a non-native English speaker, described using LLMs to manage the ongoing insecurity of professional communication:

> *"I try to use it as a pilot testing tool for my word choices, and to make sure that my language reflects my personal beliefs or principles. As a non-native English speaker I feel always a bit insecure, and I believe ChatGPT is very helpful for that emotional support."* (P17)

The emotional function was tied not to productivity but to belonging and professional confidence in a language not their own. However, this emotional relationship was fragile. When LLMs failed to meet expectations, the response was pronounced. P3 described feeling *"emotionally deflated"* when outputs felt unnatural or failed to capture what they were trying to say. This fragility points to a deeper tension: the question of what academic writing is for. Several participants expressed principled resistance to delegating writing that was not about output quality but about intellectual responsibility. P13 articulated this directly: *"Writing is essentially an extension of your thinking capability. It is necessary to go down the hard path and produce all the work, even the boring bits, by oneself."*

P4 described discomfort with the creative displacement: ''*I did not enjoy the process since it reduced my own creative input."* P17, despite being among the most enthusiastic users, drew a firm boundary: ''*I have never used them for summarising papers because I believe that I won't learn anything from that."*

What emerges is a genuine tension when using LLMs: they can serve as a coping tool for the psychological demands of academic work, but threaten the intellectual process that makes academic work meaningful. Crucially, this tension is not evenly distributed. For some participants, particularly those navigating the additional pressures of writing in a second language or early in their careers, LLMs function primarily as a confidence scaffold rather than a threat to identity. For others, especially those with a more settled sense of what academic work is for, the same tools feel like a displacement of the thing that makes the work worthwhile. This tension does not have a straightforward parallel in most professional contexts, where the value of work is more likely to be assessed by its output rather than the process of producing it, though the degree to which this holds varies significantly across occupations as AI reshapes knowledge work more broadly (Constantinides and Quercia 2025).

4.4. Theme 4: Institutional Ambiguity as Structural Paralysis

A fourth pattern concerned the institutional conditions in which academics were making trust decisions. Three distinct forms of ambiguity emerged, each constraining adoption in different ways. Crucially, none could be resolved by individual academics acting alone.

The first was data security and privacy. Several participants expressed reluctance to share institutional or research information with LLM tools due to ambiguous or absent guidelines. P7 described this as a hard constraint: *"I would not trust them with sensitive data about courses, people, and institutions, as you can't be certain where the data will go."* This forced participants to use generic rather than specific prompts, reducing usefulness.

P6 extended the concern to intellectual property: *"I have concerns that it'll give away my thinking, especially new ideas, because I don't have control of the data being trained into the models."* What is notable here is the specificity of the mechanism: the concern was not generalised distrust of AI but a concrete worry about training data, about the possibility that novel ideas shared in a prompt could surface in outputs generated for others. P7 made the structural implication explicit, describing site-level deployment as the condition under which this constraint could be lifted: *"Until such a time that these models could be implemented securely at a site-level, I would not feel trust in supplying such information."* The barrier, in other words, was not individual reluctance but the absence of infrastructure that would make responsible use possible.

The second was ethical ambiguity about what is professionally acceptable. P17 described uncertainty not about privacy but about the norms of co-authorship: *"There is a lot of conflict and debate around what constitutes model-researcher co-creation, and what is just letting the model do the work for you when you shouldn't."* These participants were not constrained by absent policy but by the absence of shared professional norms.

The third, and perhaps most quietly significant, was the absence of institutional encouragement. P4 described having no strong objection to LLM use but no impetus either: ''*There are no current major barriers but I am not being encouraged to use AI either."* The same participant went further, noting that they would probably use a reliable, university-approved tool if one existed, suggesting that willingness was present but contingent on institutional endorsement. P8 described a parallel situation: *"I don't really know how LLM tools could help me, and I don't feel I have the time to just play around with them to find out."* Without institutional guidance or dedicated time to explore these tools, familiarity cannot develop. The barrier here is not resistance but the absence of the conditions that would make adoption feel worth attempting.

These accounts describe academics who are not blocked by prohibition but who lack the institutional scaffolding that would make responsible adoption possible. Taken together, these three sub-patterns constitute what we term *institutional paralysis*: a condition in which responsible adoption decisions cannot be made by individual academics alone, because the institutional conditions that would make responsible adoption possible have not been established. This shifts the locus of the problem from individual trust attitudes to the systemic conditions that determine whether trust decisions can be made at all.

5. DISCUSSION

This study set out to examine how academics decide when and how to trust LLMs in their work. Rather than treating trust as a fixed orientation academics hold toward a technology, our findings reveal it as a situated, iterative calculation, primarily anchored on verifiability, made at the level of individual tasks. Below we discuss the implications of each theme and connect them to broader questions of design and policy.

5.1. Trust as Situated Decision-Making, Not a Global Attitude

Our first finding reframes trust in LLMs as a type of sociotechnical byproduct rather than an individual disposition. Prior work on trust in AI, including frameworks derived from automation trust research, acknowledges that trust can fluctuate in response to system performance and situational pressures (Hoff and Bashir 2015; Choung et al. 2023b), but continues to model it at the level of a user's general orientation toward a system overall. Participants' accounts consistently described trust as a task-contingent judgment rather than a settled disposition, suggesting that whatever trust calibration occurs, it operates at the level of individual tasks rather than systems.

This finding has implications for how we theorise trust in sociotechnical AI systems more broadly. Trust here is not located in the user, the system, or even the user-system intersection, but in the ongoing negotiation between a task's stakes, the available means of verification, and the institutional context in which the decision is made. This positions trust as an emergent property of a sociotechnical situation rather than a trait. The result aligns with calls in CSCW and HCI for more contextually grounded accounts of how users engage with AI systems (Bach et al. 2024), and suggests that survey instruments designed to capture a single trust score may systematically misrepresent how trust actually operates in knowledge work.

5.2. The Verifiability Burden: Why Academia is Different

The mechanism underlying academics' trust calibration is the verification burden: the point at which the cost of checking an LLM output exceeds the benefit of having used it. Prior work has identified verification as a design challenge for LLM-assisted tasks (Nogueira et al. 2026), proposing design goals aimed at reducing verification friction through citation-grounded outputs, human-in-the-loop validation, and guided transparency. Our findings extend this in an important direction: verification burden is not only a design problem to be mitigated but a structural feature of academic knowledge work that partly explains academics' restraint in adopting LLMs.

A dimension Nogueira et al.'s framing does not fully address is the expertise dependency of verification itself. Detecting an error requires enough domain knowledge to recognise the failure in the first place. Without that expertise, errors may go undetected until they have already shaped subsequent work (Zhang et al. 2020). This points to a ceiling on what design interventions alone can achieve: source-traceable outputs reduce verification costs but do not eliminate the fundamental asymmetry between producing a fluent, plausible output and establishing whether it is trustworthy.

This asymmetry is compounded by a mismatch between commercial tool design and academic epistemic standards. Across industries, organisations are actively endorsing AI adoption, with the majority of knowledge workers now using AI in some capacity (Constantinides and Quercia 2025). Academic work shares many of the same characteristics, yet institutional endorsement (discussed more in section 5.4) has largely not materialised. Commercial tools are optimised for fluency and productivity rather than verifiability, and in the absence of institutional guidance, academics are left to navigate this landscape without the infrastructure that workers in other sectors increasingly have.

5.3. Emotional Labour and LLMs in High-Pressure Academic Environments

While the psychological demands of academic work are well documented, existing design guidance for academic AI tools has focused predominantly on accuracy and workflow integration. Recent work positions LLMs as a potential resource for alleviating workload pressure (Zhang 2025), but our findings extend this by showing how this function operates among working academics rather than students, and by revealing the tensions it produces.

Participants described using LLMs not primarily for efficiency but as a scaffold for managing the affective demands of their work: the blank page, writer's block, and the anxiety of professional communication in a second language. This is consistent with broader accounts of the invisible emotional labour embedded in academic work (Mella et al. 2024), and suggests that LLMs are being recruited into this affective dimension in ways that adoption models centred on productivity do not capture.

This emotional reliance is uneven, however, and sits in tension with a deeper concern about what academic work is for. For some academics, particularly those writing in a second language or early in their careers, LLMs function as a confidence scaffold without threatening intellectual ownership. For others, the same tools feel like a displacement: writing is not a means to an output but a mode of thinking, and delegating it means losing something that cannot be recovered at the editing stage (Tang 2025; Zhai et al. 2024). The same tool that supports one academic's wellbeing may undermine another's relationship with their own thinking.

5.4. Institutional Paralysis: Ambiguity as a Collective Action Problem

Our fourth finding extends institutional paralysis accounts: responsible adoption cannot be established by individual academics alone. Participants were not held back by technophobia; many were willing to engage with LLMs but lacked the shared norms or guidance responsible use requires.

This vacuum operates at multiple levels (Liu and Jagadish 2024). Institutional policies vary widely, and journal policies are equally inconsistent, ranging from prohibition to permissive disclosure, most emerging only after 2023. Despite roughly 70% of journals now having some AI guidance, this does not reliably shape behaviour, since requirements differ so substantially across venues (He and Bu 2026): one journal may require declaring all AI assistance, another none (Ahmetoglu et al. 2026; Xia et al. 2026). Faced with this inconsistency, restraint is the safest response, precisely the pattern observed.

Individual academics cannot resolve this alone. Coordinated efforts are underway: the International Science Council, COPE, STM, and the Global Young Academy are developing a Global Reporting Standard for AI Disclosure in Research, due by the end of 2026 (International Science Council 2026). Success depends on translating such standards into institutional practice: clear data security guidance, secure institution-managed LLM environments, and clarity on what is permitted. For funders, they could require evidence of institutional AI governance, clear policies, approved tools, staff training, before funding is awarded. Without this, the burden falls on researchers who lack the authority to resolve it.

5.5. Implications for Design

We propose two design implications for academic AI tools based on our findings. We instantiate each as a low-fidelity interface concept: a Verification Panel that addresses the verification burden (Figure 1) and a Writing Support Panel that addresses the emotional burden of starting a piece of work (Figure 2).

5.5.1. Design for Task-Level Verifiability

Trust in our study operated at the level of individual tasks, not systems. Consequently, improving the reliability of LLMs will likely have a limited effect if verification at the point of use stays costly. Hence, the lever is legibility, instantiated in the Verification Panel (Figure 1) through three features. Firstly, provenance badges tag sources by type and publication status (peer-reviewed, preprint, community, etc.) allowing users to calibrate their trust accordingly. This contrasts with tools such as scite that classify citation intent rather than source type and cover only published literature. Secondly, quote grounding expands a claim to the exact source passage and page, letting users compare the LLM-generated paraphrase against the original text without opening the document. Thirdly, a verification cost indicator signals how much effort each output takes to check, from low (direct quotation, single source, exact page) to high (speculative synthesis with no supporting quotation), estimating the cost to the human of establishing correctness. The verification cost indicator has a limit: it reduces the effort needed to verify an output, but not the underlying expertise gap, since a user without domain knowledge may not recognise an error even when checking is easy. Such interventions are therefore most valuable where verification burden is highest and domain expertise is lowest.

5.5.2. Design for Emotional Function

Participants used LLMs not only for efficiency but for psychological relief, from blank-page paralysis to second-language insecurity, suggesting emotional utility is a legitimate design target. We instantiate it in the Writing Support Panel (Figure 2) through three features. Firstly, a writer's block mode treats the affective barrier as the entry point, asking what is stopping the user and routing each answer to a different form of support, rather than generating a draft as generic LLM tools, such as ChatGPT, do unless otherwise prompted. Secondly, adaptive assistance adapts to the state the user selects rather than to an inferred level of expertise. For example, a user with too many ideas is offered clustering and outlining whereas a user writing in a second language is offered fluency support that preserves intended meaning. Thirdly, low-autonomy generation deliberately produces rough, incomplete material: bullet points, directions, and open questions rather than polished prose. As a mechanism, deliberate incompleteness is established in HCI, but motivated as a way to preserve cognition and critical thinking, through provocations (e.g. (Drosos et al. 2025)), cognitive forcing functions (e.g. (Buçinca et al. 2021)), and Socratic scaffolding (e.g. (Hilliger et al. 2026)). Our findings recruit the same mechanism for a different end. Here, deliberate incompleteness gives academics a rough scaffold that lowers the affective cost of starting. Framing outputs as prompts for thinking rather than drafts for editing thus serves an emotional as much as an epistemic function.

5.6. Limitations

This study has several notable limitations. Our sample of 18 academics, though sufficient for qualitative thematic analysis, was drawn primarily from UK institutions and especially UCL. The overrepresentation of UCL participants and PhD students may have

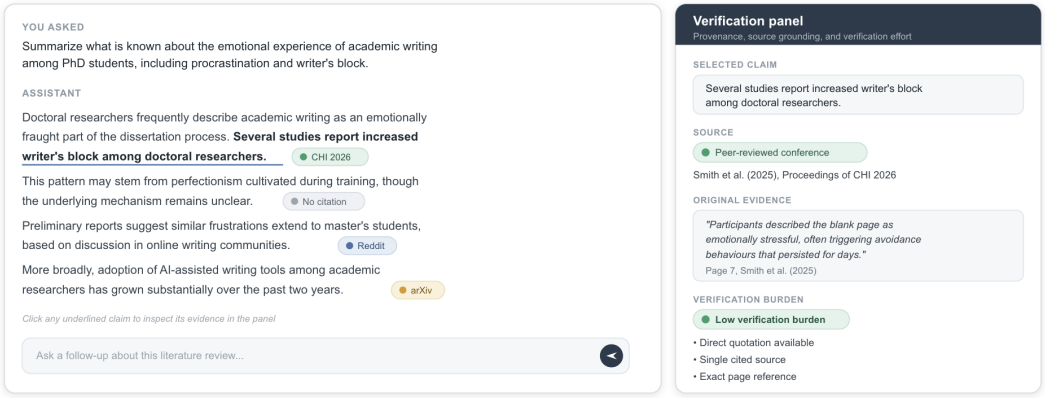


Figure 1: *Each claim in output carries a provenance badge signaling source type rather than quality, and expands into an evidence card showing the original quotation alongside a verification-burden rating (low, medium, high).*

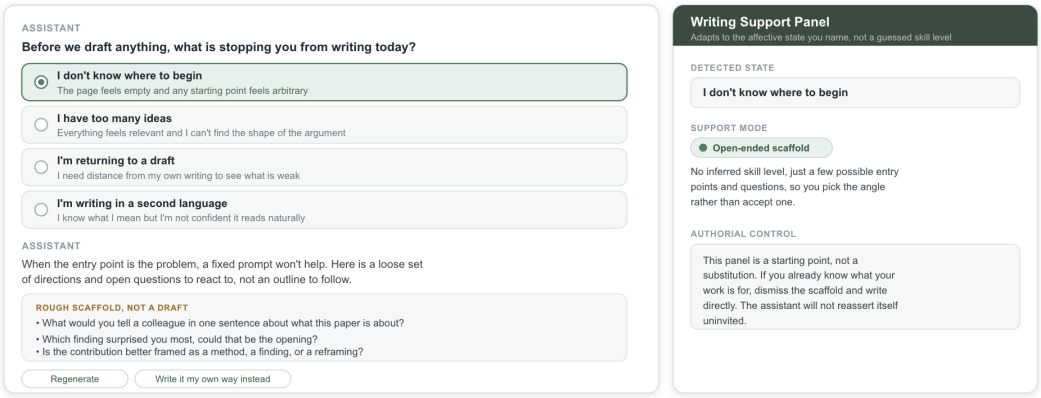


Figure 2: *Writing Support Mode asks what is stopping the user and routes to matched support (clustering, critique, fluency) via a rough scaffold rather than polished prose. Regenerate and bypass controls keep the writer in control.*

inflected the institutional paralysis theme, since both are likely to feel the absence of institutional guidance more acutely than senior academics at institutions with developed AI governance, and prior work finds that LLM engagement varies across disciplines (Ahmetoglu et al. 2026). Because responses were asynchronous and written, they capture reflective accounts rather than trust calibration in the moment, and the task-level logic participants describe may be a post-hoc rationalisation; experience sampling or think-aloud methods during actual use would test this directly. Participants were also self-selected prior LLM users, so we cannot speak to those with no experience of these tools. Finally, the study was conducted in mid-2024, and the tool and policy landscape is shifting quickly, so the specific tools and norms described may already look different. Future work should examine longitudinal patterns in trust calibration and whether the institutional paralysis we describe shifts as policies develop.

6. CONCLUSION

Academic restraint toward LLMs is a nuanced response to a structural asymmetry: in academic work, the cost of verifying outputs can exceed the benefit of using them, and the expertise needed to detect errors is often absent precisely where the tools are most needed. Our findings reframe trust appraisal from the stable, global orientation toward a system assumed in prior theoretical accounts. Instead, academics' trust in LLMs primarily reflects individual and task-level calculations shaped by verifiability, stakes, and further constrained by institutional context. These findings illustrate the importance of legibility when designing LLMs for use by academics: tools should expose source material, and flag verification costs to reduce the burden that currently governs where LLMs are and are not trusted. Beyond design, institutional paralysis is not an individual failure but a systemic one that individuals cannot resolve alone. What is needed is not more guidance documents but the infrastructure, both technical and normative, that makes responsible adoption possible in the first place.

REFERENCES

Ahmetoglu, Y., M. Constantinides, and A. Cox (2026). AI Disclosure with DAISY.

Arroyo-Machado, W., J. Ma, T. Chen, T. P. Johnson, S. Islam, L. Michalegko, and E. Welch (2025, August). Generative AI and academic scientists in US universities: Perception, experience, and adoption intentions. *PLOS One 20*(8), e0330416.

Bach, T. A., A. Khan, H. Hallock, G. Beltrão, and S. Sousa (2024, March). A Systematic Literature Review of User Trust in AI-Enabled Systems: An HCI Perspective. *International Journal of Human–Computer Interaction 40*(5), 1251–1266.

Braun, V. and V. Clarke (2023). *Thematic analysis.* American Psychological Association.

Buçinca, Z., M. B. Malaya, and K. Z. Gajos (2021, April). To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction 5*(CSCW1), 1–21.

Cargill, M. and P. O'Connor (2009). *Writing scientific research articles: strategy and steps*. Chichester, UK: Wiley-Blackwell.

Chew, R., J. Bollenbacher, M. Wenger, J. Speer, and A. Kim (2023). LLM-Assisted Content Analysis: Using Large Language Models to Support Deductive Coding.

Choudhary, G. S. and G. Yamuna (2026). Extending AI acceptance theory: Trust and psychological enablers in faculty adoption of ChatGPT. *IUP Journal of Information Technology 22*(1), 7.

Choung, H., P. David, and A. Ross (2023a). Trust in AI and its role in adoption: A systematic review. *International Journal of Human–Computer Interaction 39*(9), 1727–1743.

Choung, H., P. David, and A. Ross (2023b, May). Trust in AI and Its Role in the Acceptance of AI Technologies. *International Journal of Human–Computer Interaction 39*(9), 1727–1739.

Constantinides, M. and D. Quercia (2025, June). AI, Jobs, and the Automation Trap: Where Is HCI? In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*, Amsterdam Netherlands, pp. 1–8. ACM.

Davis, F. D. (1989, September). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly 13*(3), 319.

Doherty, L. and S. Manfredi (2006, April). Action research to develop work-life balance in a UK university. *Women in Management Review 21*(3), 241–259.

Dowling, M. M. and B. M. Lucey (2023). ChatGPT for (Finance) Research: The Bananarama Conjecture. *SSRN Electronic Journal*.

Drosos, I., A. Sarkar, Xiaotong, Xu, and N. Toronto (2025). "It makes you think": Provocations Help Restore Critical Thinking to AI-Assisted Knowledge Work. Version Number: 1.

Galindez-Acosta, J. S. and J. J. Giraldo-Huertas (2025, November). Trust in AI emerges from distrust in humans: A machine learning study on decision-making guidance.

Gefen, Karahanna, and Straub (2003). Trust and TAM in Online Shopping: An Integrated Model. *MIS Quarterly 27*(1), 51.

Gefen, D., E. Karahanna, and D. Straub (2003, August). Inexperience and experience with online stores: The importance of tam and trust. *IEEE Transactions on Engineering Management 50*(3), 307–321.

Glikson, E. and A. W. Woolley (2020). The human side of artificial intelligence: A conceptual framework and research agenda. *Academy of Management Annals 14*(2), 627–660.

Gould, S. J. J., D. P. Brumby, and A. L. Cox (2024, May). ChatTL;DR – You Really Ought to Check What the LLM Said on Your Behalf. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA, pp. 1–7. ACM.

He, Y. and Y. Bu (2026, March). Academic journals' AI policies fail to curb the surge in AI-assisted academic writing. *Proceedings of the National Academy of Sciences 123*(9), e2526734123.

Hilliger, I., M. Pérez-Sanagustín, E. Villalobos, R. Ferreira, and C. Gonzalez (2026). Scaffolding learning scenarios with a socratic chatbot: Insights from educators. *Lecture Notes in Computer Science 16064 LNCS*, 155 – 161.

Hoff, K. A. and M. Bashir (2015, May). Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust. *Human Factors: The Journal of the Human Factors and Ergonomics Society 57*(3), 407–434.

Hoffman, R. R., M. Johnson, J. M. Bradshaw, and A. Underbrink (2013, January). Trust in Automation. *IEEE Intelligent Systems 28*(1), 84–88.

International Science Council (2026). AI disclosure in research: Towards a global reporting standard.

Jensen, M. T. and E. Olsen (2023). Academic burnout: Causes and consequences. In *Practicing responsibility in business schools*, pp. 164–181. Edward Elgar Publishing.

Joshi, S. (2025). Comprehensive review of AI hallucinations: impacts and mitigation strategies for financial and business applications.

Kenesei, Z., L. Kökény, K. Ásványi, and M. Jászberényi (2025, February). The central role of trust and perceived risk in the acceptance of autonomous vehicles in an integrated UTAUT model. *European Transport Research Review 17*(1), 8.

Kim, S.-G. (2023, March). Using ChatGPT for language editing in scientific articles. *Maxillofacial Plastic and Reconstructive Surgery 45*(1), 13, s40902–023–00381–x.

Lankton, N. K., D. H. McKnight, and J. Tripp (2015, October). Technology, Humanness, and Trust: Rethinking Trust in Technology. *Journal of the Association for Information Systems 16*(10), 880–918.

Lee, A. T., R. K. Ramasamy, and A. Subbarao (2025, January). Understanding Psychosocial Barriers to Healthcare Technology Adoption: A Review of TAM Technology Acceptance Model and Unified Theory of Acceptance and Use of Technology and UTAUT Frameworks. *Healthcare 13*(3), 250.

Lee, J. D. and K. A. See (2004, January). Trust in Automation: Designing for Appropriate Reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society 46*(1), 50–80.

Liu, J. and H. V. Jagadish (2024, February). Institutional Efforts to Help Academic Researchers Implement Generative AI in Research. *Harvard Data Science Review* (Special Issue 4).

Liu, Y., S. Chen, H. Cheng, M. Yu, X. Ran, A. Mo, Y. Tang, and Y. Huang (2023). CoQuest: Exploring Research Question Co-Creation with an LLM-based Agent.

Mayer, R. C., J. H. Davis, and F. D. Schoorman (1995, July). An Integrative Model of Organizational Trust. *The Academy of Management Review 20*(3), 709.

Mella, J., S. Frampton, and A. L. Cox (2024). Unseen Work: Leveraging Generative AI for Invisible Academic Labour. . *September*.

Mhaidli, A. and K. Roemmich (2024, May). Overworking in HCI: A Reflection on Why We Are Burned Out, Stressed, and Out of Control; and What We Can Do About It. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA, pp. 1–10. ACM.

Miller, A. N., S. G. Taylor, and A. G. Bedeian (2011, September). Publish or perish: academic life as management faculty live it. *Career Development International 16*(5), 422–445.

Naidoo-Chetty, M. and M. Du Plessis (2021, June). Job Demands and Job Resources of Academics in Higher Education. *Frontiers in Psychology 12*, 631171.

Nienaber, A.-M. and G. Schewe (2014, February). ENHANCING TRUST OR REDUCING PERCEIVED RISK, WHAT MATTERS MORE WHEN LAUNCHING A NEW PRODUCT? *International Journal of Innovation Management 18*(01), 1450005.

Nogueira, B., W. Geyer, A. A. Anderson, T. J.-J. Li, D. Kim, N. Moniz, and N. V. Chawla (2026, March). From Verification Burden to Trusted Collaboration: Design Goals for LLM-Assisted Literature Reviews. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces*, Paphos Cyprus, pp. 61–65. ACM.

Rahman, M., H. J. R. Terano, N. Rahman, A. Salamzadeh, and S. Rahaman (2023, March). ChatGPT and Academic Research: A Review and Recommendations Based on Practical Examples. *Journal of Education, Management and Development Studies 3*(1), 1–12.

Sang, K., A. Powell, R. Finkel, and J. Richards (2015, July). *'Being an academic is not a 9–5 job* ': long working hours and the 'ideal worker' in UK academia. *Labour & Industry: a journal of the social and economic relations of work 25*(3), 235–249.

Singh, S. and A. Strzelecki (2026, January). Academics as adopters of generative AI: an application of diffusion of innovations theory. *Education and Information Technologies 31*(2), 621–645.

Smith, A. M. (2024, April). "Unmeasured" and never-ending working hours in UK higher education: *time* for workers, unions and employers to reinterpret the law? *Journal of Further and Higher Education 48*(4), 331–343.

Swain, J. (2018). *A Hybrid Approach to Thematic Analysis in Qualitative Research: Using a Practical Example*. 1 Oliver's Yard, 55 City Road, London EC1Y 1SP United Kingdom: SAGE Publications Ltd.

Tang, B. L. (2025, December). The Epistemic Downside of Using LLM-Based Generative AI in Academic Writing. *Publications 13*(4), 63.

Thompson, V. A., J. A. P. Turner, G. Pennycook, L. J. Ball, H. Brack, Y. Ophir, and R. Ackerman (2013, August). The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. *Cognition 128*(2), 237–251.

Trow, M. (1994, June). Managerialism and the Academic Profession: The Case of England. *Higher Education Policy 7*(2), 11–18.

Venkatesh, V. and J. Thong (2016). Unified theory of acceptance and use of technology: A synthesis and the road ahead. *Journal of the association for Information Systems*.

Xia, Q. N., M. Constantinides, A. Sarkar, D. Brumby, and A. Cox (2026). "If You're Very Clever, No One Knows You've Used It": The Social Dynamics of Developing Generative AI Literacy in the Workplace. tex.copyright: Creative Commons Attribution 4.0 International.

Zhai, C., S. Wibowo, and L. D. Li (2024, June). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: a systematic review. *Smart Learning Environments 11*(1), 28.

Zhang, M. (2025, October). Optimizing academic engagement and mental health through AI: an experimental study on LLM integration in higher education. *Frontiers in Psychology 16*, 1641212.

Zhang, Y., Q. V. Liao, and R. K. E. Bellamy (2020, January). Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, Barcelona Spain, pp. 295–305. ACM.

APPENDIX A: OPEN-ENDED SURVEY QUESTIONS

Q	Question
1	To what extent do you trust LLM tools when it comes to supporting your work tasks? Are there specific tasks for which you trust these tools more than others?
2	Can you describe the first time you used an LLM tool for work? Tell us about this experience. What worked for you and what didn't? If you haven't used an LLM tool for work, how would you imagine yourself using this tool?
3	Has your trust in LLM tools changed since you first started using them? If so, can you explain why, using specific events which changed your trust in these tools? If not, then why do you think it has remained the same?
4	Do you recall any instance where you have used LLM tools despite knowing their outputs may be of low quality? If so, can you explain a bit about why you used them regardless? If not, in what circumstances would you do this?
5	For which of your work tasks do you use LLM tools right now? Why these tasks specifically and how is this working for you? If you haven't used an LLM tool for work tasks, what work tasks could you imagine using these tools for? Why these tasks specifically and how do you think this would work for you?
6	Which of your work tasks are not well supported by LLM tools right now? What improvements could be made to these tools so that they could better support these tasks? If you haven't used an LLM tool for work tasks, what work tasks do you imagine would not be well supported by these tools? What improvements could be made?
7	LLM tools give you some form of response (e.g., a summary of a paper, an answer to a question, a draft of an abstract). What makes a response helpful when using these tools for your work tasks? What makes these responses not helpful?
8	Imagine a dream scenario where LLM tools can support every aspect of your academic work. What would you like these tools to do for you? Consider the potential for AI to integrate with programmes you already use, and provide answers beyond text, such as images, video, and voice.
9	Academics typically have responsibilities related to service, administration and teaching. Do you currently use LLM tools to support these tasks? If yes, please provide specific examples. If no, how would you imagine yourself using LLM tools for these kind of tasks?
10	How could LLM tools be improved to better help you deal with service, administration and teaching tasks?
11	Academia can also require work which involves thoughts and emotions, rather than tangible outputs – we term this 'behind-the-scenes' work (e.g., planning and strategising, monitoring and anticipating the needs of others, developing philosophies which guide your decision-making, managing your behaviours and emotions for the benefit of others). Do you currently use LLM tools to support these tasks? If yes, please provide specific examples. If no, how would you imagine yourself using LLM tools for these kind of tasks?
12	How could LLM tools be improved to better help you deal with these 'behind-the-scenes' academic tasks?
13	Do you experience any barriers preventing you from using LLM tools to support service, administration, teaching or 'behind-the-scenes' academic work? If yes, please explain these barriers.

APPENDIX B: AI USAGE DECLARATION

The following disclosure statement was created using DAISY (available at: https://dryoanaahmetoglu.github.io/daisy-disclosure/) (Ahmetoglu et al. 2026):

The authors used generative AI tools (ChatGPT 5, Claude Sonnet 4.6/5, Claude Opus 4.6/4.7/4.8/5) to improve language and readability in some parts of the manuscript, support ideation and brainstorming occasionally, draft or rewrite text in some parts of the manuscript and generate or revise a small number of figures or tables. Inputs provided to these tools: Draft notes, manuscript excerpts. Safeguards: Identifiers removed. All outputs were reviewed and verified by the authors (fact-checking, verifying citations). Authorship boundaries: Author retains full responsibility for ideas, interpretations and claims.

APPENDIX C: PARTICIPANT DEMOGRAPHIC INFORMATION

Table 1: *Participant Demographics. *UEA = University of East Anglia; UCD = University College Dublin*

	Gender	Age	Institution	Role	Discipline	Work Focus	Mode
P1	M	24	UCL	PhD Student	Psych, Psychiatry & Neuro	Research + Teaching	Office
P2	M	65	UCL	Professor	Biological Sciences	Research + Teaching	Office
P3	F	24	UCL	PhD Student	Computer Science	Research + Teaching	Hybrid
P4	M	35	UCL	Lecturer	English Lang & Lit	Teaching + Research	Remote
P5	M	58	UCL	Lecturer	Business & Management	Teaching + Research	Hybrid
P6	F	25	UCL	PhD Student	Psych, Psychiatry & Neuro	Research only	Hybrid
P7	M	28	Chester	PhD Student	Psych, Psychiatry & Neuro	Research only	Remote
P8	F	51	UEA*	Assoc. Professor	Chemistry	Teaching + Research	Office
P9	F	28	UCL	PhD Student	Psych, Psychiatry & Neuro	Research + Teaching	Hybrid
P10	M	26	Surrey	PhD Student	Psych, Psychiatry & Neuro	Research + Teaching	Hybrid
P11	M	33	Cardiff	PhD Student	Allied Health	Research + Teaching	Remote
P12	M	32	Chester	Research Asst.	Public Health	Research + Teaching	Remote
P13	F	27	UCL	PhD Student	English Lang & Lit	Research + Teaching	Remote
P14	M	34	UCD*	Research Fellow	Computer Science	Research + Teaching	Hybrid
P15	F	50	UCL	Sr. Researcher	Psych, Psychiatry & Neuro	Research + Teaching	Remote
P16	M	36	UCL	Lecturer	Computer Science	Research + Teaching	Office
P17	F	28	UCD*	Research Asst.	Computer Science	Research only	Remote
P18	M	31	Imperial	Research Asst.	Biological Sciences	Research only	Remote

APPENDIX D: ATTITUDES AND PERCEPTIONS OF LLM TOOL USE AT WORK

