

How do you ChatGPT

What OpenAI's Usage Report tells us about How to ChatGPT



Dr Neil Hopkin
Director of Education
Fortes Education



When OpenAI released ChatGPT in late 2022, the speed of adoption startled even the most jaded observers of digital culture. Within months, a tool that had once seemed the preserve of researchers and hobbyists became an everyday utility. By mid-2025, an estimated ten percent of the world's adult population was using it, a rate of diffusion that eclipses the rise of the internet, smartphones, and social media (Agarwal et al., 2025). Never before has an advisory technology, capable of producing fluent text and simulated conversation, found its way into so many hands with such velocity. It is in that unprecedented pace that the educational stakes lie, because children now grow up in a world where consulting an algorithm is as natural as checking a dictionary once was.

The recent National Bureau of Economic Research working paper by Agarwal and colleagues (2025) offers the most comprehensive look so far at what this adoption means. They sifted through billions of anonymised interactions to understand what people actually do with ChatGPT. Their findings complicate the narrative of generative AI as merely a “homework machine” or a “code generator.” The majority of interactions, around seventy percent, are non-work in character: individuals seeking advice on everyday dilemmas, asking for information, experimenting with creative writing, or simply conversing for amusement. Even in professional contexts, writing dominates: emails, reports, and communications more than programming or analysis. And most striking of all, almost half of all user messages are categorised as asking: a search for guidance or interpretation. People appear to value ChatGPT not only as a producer of words, but as an advisor, an external voice in the room (Agarwal et al., 2025).

The paper paints a picture of a technology that is far less about automation than it is about cognition: less robot worker, more counsellor. That framing has immediate resonance for educators. If adults are reshaping their intellectual habits around an “advisor mode” tool, then children, whose learning environments are more formative, are likely to internalise these habits in even more profound ways. What happens to the development of memory when students know that an omnipresent conversational partner will remind them of key facts? Psychologists already speak of the “Google effect,” the tendency to forget information more quickly when we believe it will be accessible later (Sparrow, Liu and Wegner, 2011). What happens to the writing voice of adolescents when their drafts are repeatedly polished by a model trained to smooth out idiosyncrasies? Scholars of retrieval practice remind us that durable knowledge requires effortful recall and self-testing, not frictionless retrieval (Roediger and Karpicke, 2006). And what happens to the culture of classroom help-seeking when asking a chatbot is easier than engaging in productive struggle (Aleven et al., 2003; Kapur, 2008)?

The implications extend beyond cognition. They touch upon equity and access. Agarwal et al. (2025) note that adoption rates are fastest in lower-income countries. For many students, ChatGPT is their first experience of personalised, on-demand tutoring. This suggests extraordinary potential for narrowing resource gaps in education (Warschauer, 2003). Yet as Selwyn (2016) has long cautioned, digital technologies often reproduce inequalities unless carefully curated. Language barriers, cultural mismatches, and the uneven capacity of teachers to integrate AI into their practice all risk widening divides rather than closing them.

The contrast with earlier waves of technology is instructive. Smartphones brought distraction into classrooms; social media reshaped attention spans. Both provoked reactive bans before more nuanced approaches emerged. ChatGPT is different because it insinuates itself into the very heart of schooling: reading, writing, reasoning, and deciding. Banning it risks irrelevance; uncritical embrace risks erosion of core skills. What is needed is intentional design, guardrails that ensure children use AI in ways that reinforce rather than undermine the fundamentals of learning.

That is the question this article takes up. Its title asks in two halves: How do you ChatGPT?—a descriptive question about how people already interact with the tool, answered in large part by the NBER study. And then: How to ChatGPT?—a prescriptive question about how schools must shape this interaction, ensuring that young learners do not simply become passive consumers of generated text but active agents in their own cognitive development. The journey through these pages will begin by laying out what we have learned from the usage data, then turn to the heart of the matter: the guardrails and designs that can make ChatGPT a learning amplifier rather than a shortcut. The stakes could hardly be higher. For the first time, education faces not just a new medium, but a conversational partner that millions are already treating as a second mind.

What We Have Learnt from Usage

The NBER working paper by Agarwal and colleagues (2025) is invaluable not because it confirms what we suspected, but because it challenges some of the prevailing clichés about generative AI. The early commentary framed ChatGPT as a homework machine, a plagiarism engine, or a universal code assistant. Yet when billions of real-world interactions are analysed, a far subtler pattern emerges. People are not primarily asking it to take over labour. They are not, for the most part, handing off tasks wholesale. Instead, they are reaching for advice, clarification, and textual refinement. They are turning to it as an interlocutor.

The scale of this shift should not be underestimated. In the dataset, nearly half of all user messages were classified as asking: questions, queries, prompts for explanation (Agarwal et al., 2025). This matters because it reveals a particular psychological positioning: users are treating ChatGPT less like a tool and more like an advisor. Where previous technologies mechanised processes (spreadsheets, calculators, search engines) this one embeds itself in our decision-making and sense-making. If this is how adults are shaping their cognitive environment, then the implications for children, whose intellectual habits are still malleable, are enormous.

The dominance of asking over doing points to a fundamental reconfiguration of help-seeking. In classic educational psychology, help-seeking is viewed as a metacognitive strategy: learners recognise a gap in their knowledge, then seek targeted support (Aleven et al., 2003). Done productively, it fosters autonomy; done passively, it risks dependency. ChatGPT supercharges this dynamic. A child can bypass the effortful struggle of retrieval and move straight to external guidance. Research on “productive failure” shows that grappling with problems before receiving instruction leads to deeper learning (Kapur, 2008). Yet the allure of an ever-available advisor may tilt the balance toward expedience rather than persistence.

For teachers, this presents a paradox. On the one hand, the advisory mode is precisely where the technology is most powerful. It can scaffold reasoning, suggest alternative perspectives, and widen access to expert-like feedback. On the other, if left unchecked, it risks short-circuiting the very cognitive processes (struggle, retrieval, and synthesis) that education is designed to strengthen. The lesson from usage patterns is not to discourage asking, but to structure it. Students must be guided to ask better questions, and to see the AI’s answer not as an endpoint but as raw material for further reasoning.

The dataset also shows that writing is the most common work-related use, eclipsing programming and analysis (Agarwal et al., 2025). This aligns with what teachers have observed in classrooms: students turning to ChatGPT to draft essays, smooth grammar, and generate templates. On one level, this is unsurprising. Writing is cognitively demanding and socially high-stakes. It is also the one area where small improvements in clarity and fluency produce outsized returns in school and work contexts. In effect, the technology has inserted itself into the bottleneck of human productivity.

But the educational implication is subtle. If AI writing becomes the default, students risk outsourcing not only mechanics but also voice. Already, scholars of composition warn that overreliance on algorithmic editors can produce a “flattening effect,” where idiosyncratic phrasing and authentic style are smoothed into homogeneity (Perkins, 2023). The craft of writing is not only about conveying information but about persuading, evoking, and asserting identity. If every teenage essay starts to sound like ChatGPT, we risk narrowing the diversity of expression. This is not a reason to exclude AI from writing instruction. It is a reason to integrate it deliberately, ensuring students compare drafts, reflect on revisions, and retain ownership of stylistic choices.

A third dominant category in the usage data is information-seeking. This is unsurprising when viewed against the backdrop of a decade-long shift from “search” to “answer.” For younger generations, the logic of a conversational agent is natural: why sift through links when a neatly summarised answer is instantly available? Yet here too, cognitive science offers a caution. When people expect information to be easily retrievable, they are less likely to encode it into memory (Sparrow, Liu and Wegner, 2011). The where of knowledge replaces the what. In education, this

risks a hollowing out of long-term knowledge structures, the very schema upon which critical thinking depends (Bjork and Bjork, 2011).

The educational response cannot be to pretend that students will stop using AI for information. Rather, it must be to design tasks that harness the efficiency of AI while reinforcing retrieval. For instance, students could use ChatGPT to gather initial perspectives, then close the laptop and reconstruct the key points from memory. Or they could generate competing answers and weigh their credibility. The key is to move from passive consumption of answers to active evaluation of evidence.

Equity and access shifts

Perhaps the most striking demographic finding in the NBER paper is that adoption rates are fastest in lower-income countries (Agarwal et al., 2025). This contradicts the assumption that generative AI is a luxury of the affluent. In contexts where teacher-student ratios are high and resources thin, ChatGPT may serve as the first experience of personalised tutoring. For education systems long constrained by scarcity, this is a radical possibility.

Yet, as Selwyn (2016) reminds us, technologies are rarely neutral. They are shaped by cultural contexts, linguistic accessibility, and institutional capacity. A child in Lagos or Dhaka using ChatGPT in English is receiving not only personalised tutoring but also cultural framing rooted in Anglo-American norms. The risk is a subtle form of epistemic colonialism, where local knowledges are crowded out by globalised text generation. Equally, the advantage may accrue disproportionately to students whose parents and schools can curate AI usage, leaving others exposed to misinformation and shallow engagement. Warschauer (2003) made a similar observation in the early internet era: access is not enough; meaningful use is the differentiator.

Cognitive implications from usage patterns

When one steps back from the granular categories of the NBER dataset, a broader pattern emerges. The way people are using ChatGPT maps directly onto core cognitive processes in education:

- Asking reshapes help-seeking and metacognition.
- Doing overlaps with practice, rehearsal, and refinement.
- Expressing intersects with creativity and identity.

This triad mirrors the architecture of schooling itself. It also highlights where the risks lie. If asking becomes expedient rather than exploratory, retrieval weakens. If doing becomes automating rather than rehearsing, skill acquisition stagnates. If expressing becomes mediated rather than authored, voice attenuates. The data shows us not only what people are doing, but what could be lost if those patterns are simply replicated in classrooms without adaptation.

The key lesson is not to ban, nor to celebrate uncritically, but to observe usage honestly and respond with design. Schools must assume that students will engage with AI for writing, information, and advice. They must therefore build guardrails that make these engagements serve learning rather than erode it. This means embedding retrieval before AI assistance, structuring prompts that require reasoning, and foregrounding the human dimensions of writing. It also means recognising the equity stakes: ensuring that access does not widen gaps, and that cultural frames are not uncritically imported.

In short, what we have learnt from usage is that ChatGPT is not primarily a machine of labour replacement but of cognitive partnership. Adults are already shaping their intellectual environment around it. For children, the consequences will be amplified. If education does not act deliberately, the patterns of use described in the NBER paper will quietly reshape how an entire generation thinks, writes, and decides.

Why Schools Cannot Ignore This

There is a familiar rhythm in the history of technology and schooling. A new tool enters the cultural mainstream. Teachers and parents worry about distraction, dilution of skills, or outright cheating. Bans are imposed, on calculators, then on mobile phones, later on social media. Slowly, nuance creeps in, and the technology is accommodated in a managed form. At first glance, ChatGPT might seem destined for the same cycle. But this time is different. What distinguishes generative AI is not only its pervasiveness but its intimate overlap with the very skills schools are charged with cultivating: literacy, reasoning, and judgment. To ignore it is to misread the terrain.

The first reason schools cannot turn away is that children are already using AI, often in ways invisible to adults. Surveys in the United States and Europe show that a significant minority of secondary students report using ChatGPT for homework help, essay drafting, or explanations of difficult concepts (Pew Research Center, 2023). These are not marginal behaviours. They are becoming part of the texture of learning. The NBER paper confirms that even among adults, the majority of usage is non-work related (Agarwal et al., 2025). For students, whose “work” is schoolwork, the boundary between play, exploration, and formal learning is porous. A child who uses ChatGPT to explain a maths concept for fun on a Saturday may use it again under pressure on a Sunday night homework deadline. The habits formed in one context bleed into the other.

The second reason is more subtle. If students come to rely on ChatGPT without guidance, they may internalise what educational theorists call a “hidden curriculum of expedience” (Jackson, 1968). That is, they learn that the goal of education is not to master concepts but to produce acceptable outputs quickly. When AI can generate plausible answers, the temptation is to measure success by surface fluency rather than depth. Studies of writing with algorithmic support already suggest that students who rely heavily on AI feedback produce smoother prose but retain weaker long-term mastery of grammar and structure (Perkins, 2023). The risk is that cognitive shortcuts become cultural defaults. This is not simply a question of cheating. It is a question of epistemology: what it means to know. If every query can be resolved by externalising memory into an algorithm, then the value of internal knowledge, what cognitive psychologists describe as “retrieval strength” (Bjork and Bjork, 2011), may erode. Education has always balanced between external supports (textbooks, notes, peers) and internal mastery. ChatGPT alters the balance decisively toward the external.

Third, there is the teacher dimension. Generative AI offers remarkable productivity gains for teachers, particularly in lesson planning, report writing, and administrative communication. Randomised control trials in other sectors show that less experienced workers benefit the most from AI support: their productivity gains are larger, and their output quality rises disproportionately (Noy and Zhang, 2023; Brynjolfsson et al., 2023). The same pattern is likely in schools. Novice teachers can lean on ChatGPT to produce exemplars and planning templates that would take hours to create unaided. For overburdened staff, this is a lifeline. Yet the unevenness matters. Experienced teachers may use AI as a complement, enhancing creativity and tailoring. Novices may use it as a crutch, limiting their growth. Just as with students, the question is one of design: whether AI is used to accelerate expertise or to replace it. A school that bans ChatGPT risks denying teachers a powerful support; a school that allows unfettered use risks weakening professional judgment. In either case, the stakes are institutional, not just individual.

Beyond workload, the risks lie in what cannot be immediately measured. When students use ChatGPT to write essays, the outputs look competent, sometimes more polished than unaided work. Teachers may be relieved, even impressed. But the unseen erosion may lie in critical faculties. Kuhn (2019) reminds us that the ability to weigh evidence, consider counterarguments, and justify claims is the bedrock of reasoning. If AI provides not just information but also ready-made argument structures, students may bypass the very struggle that builds reasoning muscles. The danger is not that children will stop writing essays, but that they will stop thinking through them. Similarly, the mechanics of reading may shift. If a child uses ChatGPT to summarise novels or articles, they may develop a habit of consulting “digests” rather than wrestling with the text itself. This is not unprecedented: CliffsNotes and SparkNotes did much the same. But the ease and ubiquity of AI summaries multiply the effect. Reading for gist may crowd out reading for nuance. Over time, the cognitive stamina required for deep reading, the slow, effortful processing of complex syntax and layered meaning, could decline (Wolf, 2018).

Finally, schools cannot ignore AI because of the equity paradox it presents. On one hand, ChatGPT could be the most democratising technology since the school library: a 24/7 tutor, translator, and editor, available at no cost. In lower-resource contexts, this is revolutionary. Agarwal et al. (2025) note that adoption rates are highest in lower-income countries. On the other hand, access without guidance risks reinforcing inequalities. Students from advantaged backgrounds may learn to use AI strategically as a thinking partner, as a mirror for revision, as a source of critique. Students from disadvantaged backgrounds may use it uncritically, copying outputs wholesale. The same tool could narrow or widen gaps depending on the scaffolding provided. Warschauer (2003) argued this two decades ago about the internet; the principle is unchanged. For international schools, especially in contexts like Dubai where linguistic and cultural diversity is extreme, the challenge is sharper still. AI trained primarily on Anglo-American corpora may marginalise local histories, literatures, and identities. Unless teachers intervene, students may unconsciously absorb a globalised but homogenised intellectual diet. The result could be a quiet displacement of cultural authenticity by algorithmic convenience.

Taken together, these reasons converge on a single conclusion: ChatGPT cannot be treated as just another classroom distraction to be managed. It cuts too close to the marrow of education. Literacy, reasoning, memory, and equity - the pillars of schooling - are all directly touched. Banning it may win temporary reprieve, but at the cost of irrelevance. Students will use it anyway, outside of school, in ways invisible to adults. To abdicate responsibility is to leave children alone with a technology that shapes how they think. The only viable path is intentional integration, with guardrails that ensure learning is amplified rather than eroded. So how schools can design those guardrails. The evidence from usage is clear: people treat ChatGPT as an advisor, lean on it for writing and information, and adopt it rapidly across socio-economic contexts. The task now is to translate that descriptive reality into prescriptive guidance. How do you ChatGPT? We know the answer. How to ChatGPT? That is what schools must now decide.

Guardrails and Designs: From Insights to Action

The findings of Agarwal et al. (2025) set the stage, but they do not provide the script. Knowing that most people use ChatGPT to ask questions, seek information, and write gives us a descriptive map. What schools now need is a prescriptive compass: a way to turn usage patterns into learning designs. If ChatGPT is to become a learning amplifier rather than a shortcut, education must embed guardrails: structures that direct how the tool is used, not whether it is used. What follows is a framework for those guardrails, grounded in cognitive science, pedagogy, and equity.

1. Balance Asking, Doing, Expressing

The NBER taxonomy of ChatGPT interactions, Asking, Doing, and Expressing, is not just an analytical convenience. It is a blueprint for pedagogy. These categories align closely with three educational traditions: inquiry, practice, and creativity. But they are uneven in the wild. Almost half of all messages are asking, while expressing accounts for little more than a tenth (Agarwal et al., 2025). Left unchecked, classroom usage will replicate this imbalance: lots of queries, fewer acts of creation.

A guardrail is therefore to design tasks that deliberately cycle through all three modes. For instance, in a history class, students might first ask ChatGPT for competing interpretations of an event; then do the work of structuring those interpretations into an essay plan; and finally express by composing a reflective narrative or speech in their own voice. By scaffolding the full cycle, teachers ensure that AI does not simply become a question-answer machine but a partner in the wider arc of thinking. This echoes Bruner's (1960) classic call for a "spiral curriculum," where concepts are revisited at increasing levels of complexity. Asking, doing, and expressing can become the spiral of AI-mediated learning: inquiry, rehearsal, creation. Without this guardrail, expedient asking will dominate, and expressive voice will wither.

2. Retrieval Before Revision

One of the most consistent findings in learning science is the power of retrieval practice. Students remember more when they are forced to recall information rather than re-read it (Roediger and Karpicke, 2006). Yet ChatGPT's convenience undermines this very mechanism. If an answer is always available, the struggle to recall weakens. Sparrow, Liu and Wegner (2011) showed that people remember where to find information rather than the information itself when they believe it will be accessible later.

A guardrail here is temporal: retrieval first, AI assistance second. Students could be asked to write a free-response outline of an essay, or to solve a set of problems unaided, before consulting ChatGPT for feedback or comparison. Teachers can operationalise this as the “two-pass rule”: a human-only first pass, then an AI-assisted revision. Another variation is the “closed-tab check”: after using ChatGPT, students close the screen and write down the three most important points from memory.

The principle is simple: effortful recall must precede AI refinement. Without this guardrail, ChatGPT risks becoming a cognitive prosthesis that erodes memory strength (Bjork and Bjork, 2011). With it, the tool becomes a mirror against which retrieval is sharpened.

3. AI as Mirror, Not Mask

Writing is the epicentre of AI usage. The temptation for students is to let ChatGPT mask their weaknesses by producing flawless drafts. But the real value lies in using it as a mirror. Research on metacognitive reflection shows that learning deepens when students evaluate their own work against exemplars (Nicol and Macfarlane-Dick, 2006). ChatGPT can serve as a generator of such exemplars, but only if students are asked to compare and critique.

A guardrail here is annotation. When students submit AI-assisted work, they should be required to highlight what the AI changed and explain why they kept or rejected each alteration. This transforms ChatGPT from a ghost-writer into a reflective partner. It also guards against the “flattening effect” noted by Perkins (2023), where AI smooths idiosyncratic voice into homogeneity. By foregrounding annotation, schools can preserve individuality while still harnessing AI's fluency.

4. Decision-Making Rubrics

The NBER paper highlights that the most valued function of ChatGPT is as an advisor (Agarwal et al., 2025). But advisory power without critical scrutiny is dangerous. Children may accept algorithmic suggestions uncritically, mistaking fluency for truth. Kuhn (1991; 2019) argues that the essence of reasoning is not the production of claims but the weighing of evidence and counterarguments.

A guardrail here is to embed decision-quality rubrics in AI-supported tasks. Students should not only present an answer but also rate its credibility, identify alternative perspectives, and justify why they selected one pathway over another. In practice, this could mean that when ChatGPT provides a solution to a science problem, students must annotate which assumptions the model made, and what counter-evidence might exist.

This guardrail transforms AI from an oracle into a participant in argumentation. It reinforces what Osborne et al. (2016) describe as the “epistemic aims of science education”: to evaluate knowledge claims, not just consume them.

5. Help-Seeking Guardrails

Help-seeking is one of the most delicate skills in education. Done well, it builds autonomy; done poorly, it fosters dependence (Aleven et al., 2003). The NBER paper shows that most ChatGPT interactions fall into this category of asking for help. The danger is that students learn to skip the stage of productive struggle, short-circuiting the very process that leads to durable learning. Kapur (2008) demonstrated that “productive failure”, wrestling with problems before receiving instruction, produces superior understanding.

A guardrail here is to script the form of questions. Students might be trained to ask: “What is my next step?” rather than “What is the answer?” Or: “What assumptions am I missing?” rather than “Write me the essay.” By reframing help-seeking, the AI becomes a coach rather than a crutch.

Teachers can reinforce this by building “prompt libraries” that emphasise process-oriented questions. Over time, students internalise these habits, treating AI not as a shortcut but as a scaffold.

6. Equity by Design

Agarwal et al. (2025) found that adoption rates are fastest in lower-income countries, suggesting extraordinary potential for narrowing resource gaps. But as Selwyn (2016) and Warschauer (2003) remind us, access alone does not equal equity. Without scaffolding, AI may widen gaps: advantaged students use it strategically, disadvantaged students use it uncritically.

Guardrails for equity include school-managed access (so that AI use is not restricted to those with home connectivity), bilingual prompt packs (to ensure relevance in multilingual settings), and offline fallback materials (to prevent learning loss when connectivity fails). In contexts such as Dubai, where classrooms are linguistically and culturally diverse, this also means curating examples that reflect local histories and narratives, not only Anglo-American defaults.

Equity by design acknowledges that AI is not neutral. It must be situated in context. Without this, the promise of democratisation may turn into the reality of stratification.

7. Teacher Development as Critical Infrastructure

The role of teachers cannot be reduced to gatekeeping. They are the designers of learning environments, and their own relationship with AI will set the tone for students. Evidence from workplace studies shows that AI assistance produces the greatest gains for novices (Brynjolfsson et al., 2023; Noy and Zhang, 2023). In schools, this means early-career teachers stand to benefit most, but also risk becoming overly dependent.

Guardrails here take the form of professional development. Schools must provide structured opportunities for teachers to experiment with AI in lesson planning, feedback, and assessment design. Crucially, PD should focus not on gimmicks (“ten great prompts”) but on pedagogical patterns: Socratic questioning, compare–contrast tasks, error analysis. Mentored adoption, pairing novice teachers with experienced colleagues, can help ensure that AI strengthens rather than dilutes professional growth.

This is not a marginal add-on. It is critical infrastructure. A school that fails to support teacher development in AI use will see uneven uptake, with some staff innovating and others resisting. The result will be inconsistency for students.

8. Assessment Boundaries

One of the thorniest questions is how to integrate AI into assessment. If students can use ChatGPT to draft essays or solve problems, how can schools ensure authentic measurement of learning? The guardrail here is clarity. There must be explicit AI-free checkpoints, quizzes, in-class essays, oral examinations, alongside AI-permitted drafts with attribution.

Some universities have already adopted an “AI declaration” model, where students must state what AI tools they used, for what purpose, and with what effect (Cotton et al., 2023). Schools can adapt this practice to younger learners. The point is not to punish AI use but to normalise transparency. By making AI usage visible, schools can both support learning and uphold integrity.

9. Monitoring and Metrics

Guardrails are only as good as their enforcement. Schools must therefore establish lightweight metrics to monitor AI use. These might include:

- Mode proportions: What proportion of tasks are in asking, doing, and expressing? Aim for balance.
- Retrieval gains: Do students perform better on AI-free exit tickets after AI-assisted tasks?
- Decision quality: Are students improving in their ability to evaluate sources and justify claims?
- Teacher workload: How much time is being saved in planning, and is quality rising?

These metrics need not be bureaucratic. They can be embedded into normal routines. But without monitoring, guardrails remain aspirational rather than practical.

10. Cultural Framing

Finally, there is the question of culture. ChatGPT is not a neutral machine; it is trained on a vast corpus shaped by particular linguistic and cultural biases (Bender et al., 2021). Schools must therefore act as cultural mediators. This means designing prompts that foreground local histories, literatures, and identities. It means asking students to critique AI outputs for cultural framing: Whose perspective is privileged here? Whose voice is missing?

Such guardrails ensure that AI does not simply become a vehicle for epistemic homogenisation. Instead, it becomes a tool for cultivating cultural intelligence, a capacity as important as literacy in a globalised world. Taken together, these guardrails form a design philosophy. They accept that ChatGPT is here to stay, but they insist that how it is used matters. Retrieval before revision; asking as inquiry, not expedience; AI as mirror, not mask; decision-making rubrics; equity by design; teacher development; assessment clarity; cultural mediation.

Each of these guardrails responds directly to the patterns documented in the NBER paper. If people are asking more than doing, then schools must channel that asking into deeper reasoning. If writing dominates, then schools must ensure that writing remains a site of voice, not just fluency. If adoption is fastest in lower-income countries, then equity must be curated, not assumed. The challenge for schools is not whether students will ChatGPT. They already do. The challenge is how. And the answer lies not in bans or laissez-faire acceptance but in intentional designs that turn a generative model into a generative pedagogy.

Unintended Consequences to Monitor

Every technological intervention in education comes with side effects. Calculators freed students from manual arithmetic but triggered decades of anxiety about numeracy. The internet opened worlds of knowledge but spawned a generation adept at browsing yet shallow in comprehension. With ChatGPT, the risks are subtler, less visible in the short term, but potentially more corrosive over time. Guardrails, however well designed, must be accompanied by vigilance. Schools need to anticipate the unintended consequences of integration and monitor them carefully.

The first concern is memory. Sparrow, Liu and Wegner (2011) demonstrated that when people expect information to be retrievable later, they are less likely to store it internally. This “Google effect” is amplified in conversational AI, where the barrier to access is lower still. A student who asks ChatGPT for definitions or formulae may quickly learn the habit of externalising recall. The danger is not ignorance per se, after all, reference materials have always existed, but the erosion of retrieval strength. As Bjork and Bjork (2011) argue, durable learning depends on the act of retrieval, not just exposure. If left unchecked, ChatGPT could normalise a culture of “knowing where” rather than “knowing what.” This would compromise the very schemas that underpin problem-solving and transfer. Schools must therefore monitor whether AI-supported students retain core knowledge over time. Quick retrieval tests, oral checks, and AI-free assessments become not just evaluative tools but diagnostic ones. Without such vigilance, memory could silently atrophy.

The second risk is stylistic homogenisation. Early analyses of AI-assisted student writing suggest that while grammar improves, individual voice diminishes (Perkins, 2023). Adolescents in particular need space to experiment with style, awkward metaphors, eccentric phrasing, bold digressions. These are not flaws but the raw materials of identity formation. If every draft is polished by ChatGPT, the rough edges that give writing its texture may be sanded away.

The unintended consequence here is cultural as much as cognitive. Schools could end up producing cohorts of students whose prose is competent but indistinguishable, voices tuned to the statistical average. Teachers must therefore monitor not only accuracy but originality. Rubrics should reward risk-taking in expression, not just fluency. Without that corrective, AI could create an age of stylistic conformity.

A third concern is epistemic shallowness. ChatGPT excels at producing summaries and digests. For many students, this is a gift. But the unintended effect may be a decline in the stamina required for deep reading. Wolf (2018) warns of the dangers of a “skimming culture,” where the neurological circuits for sustained comprehension weaken through disuse. If students rely on AI-generated synopses, they may bypass the slow wrestling with complex syntax and layered meaning that builds cognitive depth. This risk is difficult to detect because AI-generated summaries feel efficient and harmless. Yet over time, students may lose the ability—or the patience—to engage with unmediated texts. Schools must monitor not only whether students can answer questions about content but whether they can navigate original sources unaided. Otherwise, the discipline of reading could be supplanted by the convenience of digesting.

Another unintended consequence is epistemic overconfidence. ChatGPT produces fluent answers even when uncertain. Students may mistake fluency for truth, a phenomenon well-documented in cognitive psychology (Oppenheimer, 2008). Without explicit training in epistemic vigilance, children may absorb misinformation without recognising it. The danger here is not just factual error but the erosion of critical faculties. If students grow accustomed to accepting AI outputs at face value, their capacity for scepticism—central to scientific and civic reasoning—may decline. Schools must therefore monitor not only accuracy but epistemic behaviour: Are students checking sources? Are they recognising uncertainty? Without such guardrails, misinformation becomes a hidden curriculum.

Finally, there is the equity paradox. As noted, ChatGPT adoption is fastest in lower-income countries (Agarwal et al., 2025). This presents extraordinary opportunities but also hidden risks. Students with curated guidance may use AI strategically, while those without may become passive consumers. The unintended consequence could be widening rather than narrowing gaps.

Moreover, AI outputs reflect cultural and linguistic biases in their training data (Bender et al., 2021). For students in multilingual contexts, this may mean subtle displacement of local narratives by globalised ones. Without vigilance, classrooms could become sites of epistemic homogenisation, where Anglo-American frames dominate. Monitoring must therefore include cultural audits: Are local voices represented? Are students encouraged to critique the cultural framing of AI outputs?

What unites these risks is their invisibility. Offloaded memory, flattened voice, shallow reading, misplaced confidence, hidden inequities—none show up immediately in test scores. They emerge slowly, shaping habits and dispositions. This makes vigilance itself a pedagogical act. Teachers must not only deliver content but also observe how students are using AI, what skills are being reinforced, and which ones are quietly eroding. The challenge is not to eliminate unintended consequences—no technology has ever done so—but to anticipate and manage them. By naming the risks, schools can build them into design. Retrieval checks guard against memory erosion. Rubrics for originality counteract flattening. Sustained reading tasks preserve stamina. Source evaluation exercises strengthen vigilance. Cultural critique ensures plurality. In each case, monitoring is not an afterthought but a central pillar of intentional integration.

The story of ChatGPT in education is not about whether students will use it. They already do. It is about how schools respond to the patterns revealed by usage data and the cognitive science that helps us interpret them. The NBER study (Agarwal et al., 2025) shows us a world where people treat AI as an advisor, where writing and information-seeking dominate, and where adoption is rapid even in low-resource contexts. For education, this is both a warning and an invitation.

The warning is that without intentional guardrails, memory may erode, voices may flatten, reading stamina may decline, and inequalities may widen. The invitation is that with thoughtful design, ChatGPT can become a scaffold for retrieval, a mirror for reflection, a catalyst for argumentation, and a partner in creativity. The question, then, is twofold. How do you ChatGPT? The answer lies in usage: as an advisor, a writer's aide, a source of quick information. How to ChatGPT? That is the challenge before schools: to design environments where AI strengthens rather than supplants the fundamentals of learning. The choice is not between adoption and resistance, but between erosion and amplification.

About the Author

Dr. Neil Hopkin is a globally recognised thought leader in international K-12 education, and serves as the Director of Education at Fortes Education.

His extensive academic background includes advising UK government bodies and spearheading significant educational initiatives, particularly with the EdTech, Early Years, Higher Education and Teacher Professional Development fields, equipping him with invaluable insights and expertise. As the head of Fortes' Academic Leadership Team, Dr. Hopkin is responsible for overseeing academic performance, operational efficiency, curriculum development, and staff professional development across Fortes Education institutions.



For more information contact Dr Neil Hopkin at:

www.sunmarke.com

www.risdubai.com



Bibliography

Agarwal, R., Gans, J., Goldfarb, A., Gopal, R., Kramer, H., Stern, S. & Tambe, P. (2025) The Usage and Impact of ChatGPT. NBER Working Paper No. 34255. National Bureau of Economic Research.

Aleven, V., Stahl, E., Schworm, S., Fischer, F. & Wallace, R. (2003) 'Help seeking and help design in interactive learning environments', *Review of Educational Research*, 73(3), pp. 277–320.

Bender, E.M., Gebru, T., McMillan-Major, A. & Shmitchell, S. (2021) 'On the dangers of stochastic parrots: Can language models be too big?', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York: ACM, pp. 610–623.

Bjork, R.A. & Bjork, E.L. (2011) 'Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning', in Gernsbacher, M.A., Pew, R.W., Hough, L.M. & Pomerantz, J.R. (eds.) *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*. New York: Worth, pp. 56–64.

Bruner, J. (1960) *The Process of Education*. Cambridge, MA: Harvard University Press.

Brynjolfsson, E., Li, D. & Raymond, L.R. (2023) 'Generative AI at work', NBER Working Paper Series, No. 31161. National Bureau of Economic Research.

Cotton, D.R.E., Cotton, P.A. & Shipway, J.R. (2023) 'Chatting and cheating: Ensuring academic integrity in the era of ChatGPT', *Innovations in Education and Teaching International*, 60(2), pp. 176–185.

Jackson, P.W. (1968) *Life in Classrooms*. New York: Holt, Rinehart & Winston.

Kapur, M. (2008) 'Productive failure', *Cognition and Instruction*, 26(3), pp. 379–424.

Kuhn, D. (1991) *The Skills of Argument*. Cambridge: Cambridge University Press.

Kuhn, D. (2019) *Arguing to Learn*. New York: Routledge.

Nicol, D. & Macfarlane-Dick, D. (2006) 'Formative assessment and self-regulated learning: A model and seven principles of good feedback practice', *Studies in Higher Education*, 31(2), pp. 199–218.

Noy, S. & Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187–192.

Oppenheimer, D.M. (2008) 'The secret life of fluency', *Trends in Cognitive Sciences*, 12(6), pp. 237–241.

Osborne, J., Simon, S., Christodoulou, A., Howell-Richardson, C. & Richardson, K. (2016) 'Learning to argue: A study of four schools and their attempt to develop the use of argumentation as a common instructional practice and its impact on students', *Journal of Research in Science Teaching*, 50(3), pp. 315–347.

Perkins, R. (2023) 'AI and the flattening of student voice: Risks and responses in writing education', *English Journal*, 112(5), pp. 35–43.

Pew Research Center (2023) *Teens, Social Media and Technology 2023*. Washington, DC: Pew Research Center.

Roediger, H.L. & Karpicke, J.D. (2006) 'Test-enhanced learning: Taking memory tests improves long-term retention', *Psychological Science*, 17(3), pp. 249–255.



Selwyn, N. (2016) *Education and Technology: Key Issues and Debates*. 2nd edn. London: Bloomsbury.

Sparrow, B., Liu, J. & Wegner, D.M. (2011) 'Google effects on memory: Cognitive consequences of having information at our fingertips', *Science*, 333(6043), pp. 776–778.

Warschauer, M. (2003) *Technology and Social Inclusion: Rethinking the Digital Divide*. Cambridge, MA: MIT Press.

Wolf, M. (2018) *Reader, Come Home: The Reading Brain in a Digital World*. New York: HarperCollins.