Check for updates

Essay

Journal of Management Education
1–33
© The Author(s) 2024
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/10525629241261313
journals.sagepub.com/home/jme
Sage

Exploring the Impact of ChatGPT on Business School Education: Prospects, Boundaries, and Paradoxes

Sorin Valcea[1], Maria Riaz Hamdani[2], and Shuai Wang[1]

Abstract

This essay explores the nuanced impact of generative AI technologies on management and business education, framed through three paradoxes: the *Expertise Paradox* suggests that AI's adequate performance at lower-level tasks may weaken students' development of higher-level thinking; the *Innovation Paradox* states that AI's creativity aid could stifle original thinking; and the *Equity Paradox* highlights AI's potential to provide immense gains to experts but disproportionately harm novices. We take the position that without "sensible" AI use guidelines in management education, AI is likely to have a net-negative effect on learning. This stance is based on our trials with ChatGPT on various cognitive tasks organized around the revised Bloom's Taxonomy of learning. We identify areas where AI tools can enhance learning, such as comprehending established subject domains, as well as areas where they exhibit significant limitations, such as logical reasoning and critical thinking. We caution against the potential deskilling in critical thinking due to students' overreliance on AI for basic tasks. To alleviate these challenges, we recommend sensible AI uses by students that

[1]Cleveland State University, OH, USA
[2]The University of Akron, OH, USA

Corresponding Author:
Maria Riaz Hamdani, College of Business, The University of Akron, 259 South Broadway Street, Akron, OH 44325, USA.
Email: hamdani@uakron.edu

support skill development without fostering overreliance. We also suggest how faculty, administrators, and employers may support students in getting the most out of this new tool.

Keywords
Generative AI, andragogy, deskilling, Bloom's taxonomy, creativity, insights, sensible AI use

Introduction

On November 30th, 2022, OpenAI (2022) released the beta version of its ChatGPT chatbot to the public for free. The press release described ChatGPT to be an AI chatbot built on top of OpenAI's generative pre-trained transformer (GPT) models GPT-3.5 and GPT-4 families of large language models (LLMs). As a result, ChatGPT can mimic human's natural, fluid, back, and forth conversations in response to user prompts. ChatGPT's release was a watershed moment in the history of AI, evoking a mix of excitement and caution among people, as they considered the potential impact of this new technology. Since then, public access to myriad generative AI tools is paving the way for large-scale assimilation of this technology into daily activities, including creating content, translating languages, providing detailed answers to queries across many domains of knowledge, writing code, generating reports, drafting letters, and automating customer service (McKinsey & Company, 2023).

While such integration of generative AI into business processes promises efficiency and innovation, it also introduces complex challenges, such as the deskilling and reskilling of workforce, threats of job eliminations, misuse of intellectual property, data privacy, and unauthorized AI use. Management education finds itself at a pivotal juncture, where it must gear to equip future managers with the necessary knowledge and skills to navigate these challenges. While some have raised alarms about the potentially large AI-exposure of many business-related professions (e.g., Felten et al., 2021), we focus in this essay on how overreliance on AI may hinder business students' development of critical thinking and creativity.

There is a diversity of opinion regarding challenges brought on by AI and how to address them. A minority of experts tend to be dismissive of generative AI. For example, Naom Chomsky called ChatGPT a high-tech tool for plagiarism (Mazhari, 2023), and stated: "Given the amorality, faux science and linguistic incompetence of these systems, we can only laugh or cry at

their popularity" (Chomski et al., 2023). Overall, most opinions appear to gravitate either toward AI-as-savior or AI-as-destroyer.

In the former camp, Eric Schmidt, former CEO of Google, stated: "What is arriving is a kind of intelligence that's different. It comes to answers differently than we do. It seems to have hidden understanding and meaning that we don't understand today. It discovers things that we've never known. We don't know how far this goes (. . .). We have not figured out how powerful this technology's going to be yet (. . .). It will be able to generate insights and ideas that we as humans have not had. Think of them as existing as savants" (Schmidt, 2023).

In contrast, others have voiced concerns that AI brings existential threats to humanity, and should receive at least the same attention as the non-proliferation of nuclear weapons or climate change (Safe AI, n.d.). These claims are further magnified by the media's breathless coverage that often runs ahead of the facts (Vinsel, 2023). Yet there are some who reason that the hyperbolic claims about AI's benefits and threats can be traced to commercial interests surrounding the technology (Leaver & Srdarov, 2023). Despite this, they remain cautiously optimistic about the benefits of AI. We agree with this cautiously optimistic approach.

In this essay, we opine that generative AI is neither a savior nor a destroyer. Paradoxically, although generative AI can certainly help with one's learning and work, the very way in which it helps can end up harming the individual's long-term growth, when used indiscriminately. Our position is that, without guidelines for what we later describe as "sensible" AI use in education, AI is likely to have a net-negative effect on learning. This is because more cognitive skills and processes, which are foundational for learning, risk being "outsourced" to the technology.

Ironically, AI's intermediate capabilities (positioned between being a savior/destroyer and a mere tool for plagiarism) heightens the likelihood of negative impact. If AI were significantly more powerful, it would likely face greater resistance and scrutiny from labor and governments due to the potential for substantial economic disruption. Conversely, if AI were less capable, organizations would have minimal incentive to invest in its deployment. As it currently stands, the lure of short-term productivity gains and the lack of an immediate threat are likely to encourage more rapid and widespread adoption. Ultimately, the cumulative societal effects depend on how potential short-term benefits and long-term harms are managed through effective policies and practices in education and businesses. Management education can play a critical role in preparing students to address these issues. In doing so, the first step is to understand how this technology is affecting management education.

In this essay, we use the framework of three paradoxes of expertise, innovation, and equity, to encapsulate the central challenges and opportunities of generative AI for management education. The Expertise Paradox highlights the tension between the convenience of using generative AI tools to produce eloquent (but shallow) content and the disincentive to invest effort in developing more complex critical thinking skills. The Innovation Paradox examines the delicate balance between adopting technological innovations and preserving human creativity. The Equity Paradox confronts the issue of generative AI allowing experts to fully benefit from the power of AI tools, while at the same time exposing novices to significant risk to their expertise development.

We focus on the multifaceted role of management education in preparing leaders for a future where AI plays a central role in business strategy and operations. We argue for a systematic approach to business education that addresses the three paradoxes we identified by fostering in students a sensible use of AI, expanding management curriculum to include AI and management implication topics, and partnering with the business community to find innovative ways to upskill students. Students should be trained to recognize sensible uses of generative AI tools that allow them to upskill rather than deskill. This means teaching students that indiscriminate reliance on ChatGPT might help them turn in assignments faster (and cheat without detection) but will also likely stagnate their skill development. Coursework should help students navigate myriad issues that surround the implementation of generative AI in work, such as data privacy issues and security, biases and fairness in generative AI output, regulatory compliance, costs and returns in investment, as well as the ethical and societal implications of its deployment. Management educators must cultivate a nuanced understanding of how generative AI technologies can impact one's ability to think critically, innovate, and create equity.

Our arguments are particularly relevant for the traditional undergraduate student with limited expertise and job experience. We start by differentiating between human learning and machine learning. We use the revised Bloom's taxonomy (Anderson et al., 2001) as a framework to understand learning processes and different types of knowledge that could be negatively impacted by unchecked generative AI usage. To elucidate the paradoxes, we share our experience with ChatGPT on solving various knowledge and reasoning tasks. We end by making recommendations for sensible AI use by management and business students. We also discuss how faculty, administrators, and employers can support students' learning in the age of generative AI. We focus primarily on ChatGPT because of the rapid growth in its user base and the

intense attention it has received in both media and academia; nevertheless, issues around the use of ChatGPT may be illustrative for the broader context of generative AI use in education.

Adult Learning Models in Management Education

Large Language Models (LLMs), such as ChatGPT, belong to a broader category of AI algorithms called machine learning (ML). Unlike traditional algorithms that follow programmer-given instructions, ML algorithms operate backwards, analyzing vast data to find patterns and use these patterns to generate outputs. LLMs generate coherent, grammatically accurate sentences by learning to predict words based on the statistical representation of the co-occurrence patterns in human-produced text (Brynjolfsson et al., 2023).

This fluency creates a compelling illusion of conversing with an intelligent entity, possibly because coherent language outputs invoke deeply embedded mental models of sharing thoughts, feelings, and companionship (Roche et al., 2002). To a casual user, the AI appears to exercise some sort of machine wisdom to access, comprehend, evaluate, and synthesize collective human knowledge, and even learn from the interlocutor. However, this is not the case: LLMs do not actually understand words or context (Weidinger et al., 2021). Machine outputs are coherent without any "receptive wisdom." It is this non-receptive coherence that is so spectacular but at the same time misleading for a human brain.

In contrast, human sense-making and learning are deeply contextual, experiential, personal, and complex (Schwandt, 2005), as emphasized in many theories of andragogy invoked in management education. For example, Fink's (2013) significant learning model contextualizes learning around a learner's history, and focuses on deepening a learner's interest in the topic by creating new feelings, interests, and values. D. A. Kolb's (1984) cycle of experiential learning (A. Y. Kolb & Kolb, 2005) assumes a similar constructivist approach where knowledge comes from the integrated functioning of total human experience (thoughts, feelings, abstractions, and actions). These models underscore learning as a deeply personal process, where context and personal meaning lead to insights and discoveries, highlighting the contrast with AI's inability to understand the context of its outputs. For humans, understanding and adapting to the context is often the whole point of learning (A. Y. Kolb & Kolb, 2005).

In addition, learning theories also suggest a hierarchical dependency between various levels of learning. The original and revised Bloom's

taxonomy (Anderson et al., 2001; Bloom et al., 1956) focus on six hierarchical cognitive processes or thinking skills in learning: remember, understand, apply, analyze, evaluate, and create. The higher-level cognitive processes are built on lower-level processes of remember, understand, and apply (Anderson et al., 2001). For example, one cannot evaluate a hypothesis if one cannot recall and apply foundational knowledge in the subject domain. Bloom's taxonomy has guided many instructional innovations in management education to teach increasingly complex thinking and reasoning skills, such as critical thinking, logical reasoning, and moral reasoning (e.g., Athanassiou et al., 2003; Hillmann et al., 2018; Lovelace et al., 2016; Reeves, 1990).

The revised Bloom's Taxonomy (Anderson et al., 2001) integrates six cognitive processes with four knowledge types: factual (basic elements like terms and facts), conceptual (connections among elements within a framework), procedural (how to perform tasks and methods), and metacognitive (awareness of one's cognitive processes). Some theorists (Polanyi, 1966) have also delineated between explicit (formalized, codified, and transferable) and implicit/tacit (informal, personal, and vague) knowledge, indicating that the creation of new knowledge often results from the conversion of tacit to explicit. As we detail later, generative AI can affect the way students acquire and apply different forms of knowledge.

Dependencies between lower and higher cognitive skills are also evident from neurosciences. The brain's temporal and parietal regions store various knowledge types, which the prefrontal cortex then accesses for higher-level thinking and decision-making (Dietrich, 2004; Dietrich & Kanso, 2010; Heilman et al., 2003). Managerial skills like creative problem-solving depend on this stored knowledge (Heilman et al., 2003; Shen et al., 2013). Relying on AI tools for quick and easy information access, instead of retaining and comprehending knowledge may decrease creative insights, as we discuss later.

These distinctions between machine learning and human learning are important because, as our trials with ChatGPT demonstrate, LLMs appear to simulate foundational knowledge and, we argue later, may create disincentives for learners to invest time and effort in acquiring such knowledge. To gain insight on how generative AI tools may affect student learning, we explored ChatGPT's capability on business curricula. We provide a more detailed discussion of these trials with ChatGPT and their results in Supplemental Appendices A to C. Next, we present the three paradoxes that become apparent from the AI's performance on these tests, and reference some of our trials for illustrative purposes.

Three Paradoxes of Generative AI Tools—Expertise, Innovation, and Equity

We use the term paradox to mean a statement or situation that seems to be self-contradicting, and yet may still be true or possible. We argue that text-generating AIs can fit this definition because the way they help learners in the short-term may hurt their long-term development. Reflecting on the impact of generative AIs on literature and writing, Vara (2023), a journalist and an author, asked a thought-provoking question: if AI is good at writing fiction, should writers use it to assist with their own writing. Later, she raises another equally challenging question: if the use of AI for creative writing becomes common, would human writers still retain their creative ability to write independently of any assistive technology? Our trials with ChatGPT suggest that there may be similar concerns for students. Below, we argue that over-relying on AI help can be harmful in the long run by deskilling the very act of critical thinking. Later we discuss how to avoid this and leverage AI for learning and development.

The Paradox of Expertise

> *The paradox of expertise is that while generative AI assists in executing tasks requiring lower-level thinking skills, this advantage paradoxically hinders a learner's ability to develop higher-level thinking skills required for complex tasks. This happens because the use of generative AI by novices weakens their learning of foundational skills that are required for higher-level thinking.*

The first and, in our opinion, the most important paradox of AI is its effect on the development of expertise. On one side of the AI hype, proponents claim that AI will greatly augment human productivity, by allowing workers to focus on the more challenging, creative, and meaningful aspects of their work, and replacing their effort in other more repetitive and laborious tasks (Tschang & Almirall, 2021). Implicit to this idea is that users of AI are experts whose efforts are better utilized on complex tasks. On the other end of the hype, "detractors" warn that AIs are close to being so advanced as to replace even the skills of experts, and thus cause enormous disruption to the economic livelihood of individuals.

Our own trials with ChatGPT, Bard (now Gemini), and Bing Chat (now Bing AI Copilot) suggest a middle reality, but one nonetheless rife with challenges that need to be overcome, if we are to fully realize the opportunities from AI. Like the proponents of AI suggest, we found that generative chatbots performed well on several understanding and application tasks (second

and third level of Bloom's taxonomy) concerning factual, procedural, and conceptual knowledge, as illustrated in trials summarized in Supplemental Appendix A and more detailed in Supplemental Appendix C. For example, AIs could summarize theoretical frameworks and models (comprehension of conceptual knowledge in the domains tested), accurately classify and exemplify concepts (comprehension of factual knowledge), carry out procedures, such as adverse impact test, record accounting transactions, and solve a variety of finance problems, including computing expected portfolio return, account margin and margin call scenarios, stock relative strength, and payoff and profits of call and put options (application of procedural knowledge). To sum, it performed reasonably well on remember, understand, and apply levels of Bloom's taxonomy across various forms of knowledge.

Contradicting the more alarmist arguments, we found that the AIs performed less well on tasks from the higher levels of Bloom's taxonomy. For example, we prompted the AIs with an application problem which stated that a study found a correlation corrected for measurement and sampling error between job satisfaction and affective commitment of .91 (Le et al., 2010). We then asked if this correlation supports the idea that job satisfaction and affective commitment are distinct constructs, and we received an affirmative answer (incorrect answer). The AI continued to provide affirmative answers even when prompted to first discuss discriminant and convergent validity, and apply the discussion to the interpretation of the correlation.

Nevertheless, we argue that the ability of AIs to effectively replace human skills for lower-level tasks sets the stage for a stagnation in expertise development. The reason is that, while experts can safely replace their efforts on lower-level tasks and refocus on more complex work, novices (i.e., learners) who do the same jeopardize their development, because those tasks often provide the conceptual and experiential basis for more advanced work. In other words, one cannot attempt the complex without first mastering the basics. However, the easy availability of AI-generated answers for recall, comprehension, and application questions may reduce students' perceived need to acquire knowledge and promote overreliance on the AI tools.

Anecdotal evidence from faculty across disciplines, a few studies on AI, as well as studies on other technologies cumulatively suggest that AI will be used, at least by some, as a means of replacing rather than augmenting effort. Social media is abuzz with reports from faculty on the surge of homework submissions that appear to have been generated by AIs. The problem is significant enough that some have asked for outright bans of AIs on campuses, fearing an increase in student cheating (Clercq, 2023). Recently, a Pew Research Center study found that one in five US teens who have heard of ChatGPT have used it for schoolwork (Sidoti & Gottfried, 2023). Clearly,

many students are already taking advantage of the new technology to reduce their time commitment, effort, and, unfortunately, development opportunity.

Elsewhere, a recent study found that the use of AI tools improved the productivity of employees across skill levels and increased the output quality of lower skilled employees on writing tasks, such as press releases and official letters (Noy & Zhang, 2023). Interestingly, employees preferred to use the AI to substitute for their effort (i.e., by often submitting ChatGPT created drafts without any edits), rather than complement their effort (i.e., by cycling through AI-created versions and selecting the one they liked most, or by improving upon the AI draft). In learning environments, such effort-substituting shortcuts can greatly diminish learning by reducing the tacit procedural knowledge that is acquired when personally doing things.

Furthermore, Wegner and Ward (2013) report on a series of experiments which suggest that the easy access to information via technology reduced cognitive effort in participants. In one of their experiments, individuals were less likely to commit a fact to memory (i.e., decreased factual knowledge), even if explicitly instructed to do so, when they had the option to store that information on a computer. In another experiment, individuals who had relied on Google to answer trivia questions felt higher cognitive self-esteem, hinting that the speed with which search engines answered questions created an illusion that the user's own mind had produced the results. The authors concluded that, ironically, people who relied on the internet for information felt they knew more and more, while knowing less and less. In other words, technology replaced the effort, but created the illusion of only augmenting the user's own capability, and distorting the learner's metacognitive knowledge.

Generative AIs are likely to exacerbate distortions in metacognitive knowledge for several reasons. First, they are even better and faster at summarizing internet information than traditional search engines. Second, because an answer is generated (rather than merely retrieving a link), one gets the feeling of having produced work output. Finally, because the prompts given to a generative AI can be significantly more elaborate than the typical search words fed to a search engine, the user may feel a stronger involvement and ownership over the output. The recent interest in prompt engineering is certainly illustrative of this idea.

One of the reasons users of AI (including business students) may choose to use AI to replace effort is the perception that they would be outcompeted if they did not. While champions of AI typically paint a hopeful future, they do agree with the apocalyptic alarmists on one thing: workers who do not use AI will be replaced by those who do (Lakhani, 2023). Many students often must balance work, family, and academic responsibilities. Offloading some

of the effort onto an AI may appear as a reasonable short-term choice, even for those who may understand the detrimental long-term implications.

Students may also choose to use AI output instead of their own because it can often time be superior. This is not particularly surprising. On one hand, reasonably good composition skills and flawless grammar and spelling are some of the best documented strengths of text generating AIs like ChatGPT. On the other hand, novices in any field need to spend time and effort to attain that same level of proficiency. AI promises to help bridge that gap. For example, a study examining a corporate implementation of LLMs in customer services (Brynjolfsson et al., 2023) found that AI-assisted novices quickly achieved productivity levels similar to experts; experts, however, actually performed less well when they had to pay attention to AI suggestions. The previously referenced study on writing tasks (Noy & Zhang, 2023) also showed that AIs helped lower-skilled workers with their quality of work. Neither of the above-mentioned two studies tested whether quality and productivity gains were sustainable once the AI assistance was removed. This is an important unknown, since it is possible that the AI tools acted as a crutch, stifling the natural gain in skills and tacit knowledge that comes with repeated practice and experience.

The above evidence raises concerns that AI use by students will reduce learning of foundational material, which then will reduce development of more complex thinking skills that depend on this foundational material (Anderson et al., 2001). Classrooms can always benefit from more experiential and discursive pedagogies that AIs cannot replace, but one cannot have meaningful discussion of, or make informed decision on more complex issues without the mastery of the fundamentals. AI-caused deskilling may be further aggravated by corresponding phenomena in the labor market. We illustrated earlier that the AI performed well on several finance, accounting, and human resource problems. These are precisely the type of tasks an entry-level position may demand out of a college graduate. If an AI can greatly augment or even substitute human performance in these tasks, it stands to reason that fewer jobs of this kind will be necessary. If entry level positions are eliminated and organizations only hire experts, fresh graduates would have fewer opportunities to gain experience and grow professionally. Thus, pipelines for talent development will be bottlenecked.

Anecdotal evidence in management education and research indicates a shift toward increased ChatGPT reliance over graduate students for various tasks, from planning data collection, to writing statistical code, as shared by a colleague. This trend could lead to a loss of vital mentoring and developmental opportunities. Additionally, as these tools become more sophisticated

and accurate, their increasing adoption by various users might further curtail developmental opportunities for honing critical thinking skills.

The paradox of expertise sets the stage for the related but different issues of innovation and equity. To make this connection more apparent, we refer here to a study which tested whether non-experts and experts in a domain of knowledge were more likely to show scientific creativity in that particular domain (Acar & van den Ende, 2016). One of the core findings of the study was that both groups needed to invest substantial time and effort thinking and researching about the problem to be successful. This finding substantiates our concerns about over-reliance on AI reducing learning effort and hampering positive outcomes. Further down, we suggest it may also reduce creative effort, and hence creativity.

A second finding from the study was that experts and non-experts needed to employ different research strategies to be effective. Those who had expertise relevant to the problem used broad searches across wider domains of knowledge, identified common patterns across fields, and utilized these new thinking patterns to overcome their cognitive fixation. In contrast, non-experts innovated by conducting focused searches within the problem domain to familiarize themselves with the focal field. Thus, if students (i.e., non-experts) do not develop foundational knowledge to familiarize themselves with the domain because of over-reliance on AI, they may also limit their creative potential.

Interestingly, the difference in effective strategies between experts and non-experts also prefaces our subsequent arguments about the paradox of equity. Our experience with AI has taught us that it is particularly adept at summarizing well-established knowledge domains, but starts to falter when one delves deeper into those domains. On one hand, the type of broad and shallow access to knowledge that AI provides is precisely what benefits experts the most in creative problem solving (Acar & van den Ende, 2016). On the other hand, AI is not only less effective for the type of in-depth research that would benefit non-experts, but could also be potentially be harmful. It can produce incomplete, misleading, or outright incorrect outputs that non-experts might have difficulty recognizing as such. With this preamble, we proceed to discuss in detail the remaining two paradoxes.

The Paradox of Innovation

Generative AI has the potential to bolster innovation and creativity by helping with lower-level tasks so one can focus on higher-level creative tasks; however, generative AI may also discourage creative thought, by repackaging existing content as something new.

The revised Bloom's taxonomy ranks "Create" as the most complex cognitive process in learning. Creativity involves inductive reasoning and depends on the learner's ability to critically evaluate prior knowledge and integrate it to make a new whole (Anderson et al., 2001, p. 294). Creating is often the ultimate pedagogical goal, that is, to integrate previously acquired knowledge to construct a new idea or a solution. Accordingly, business and management courses often involve assignments that require students to demonstrate generating, planning, and implementing new solutions—the three sub-processes of create.

The arguments for generative AI bolstering creativity include the refocusing of human time and effort away from tedious tasks, that the AI would take over, and toward more creative work. For example, AIs could speed up software development by rapidly generating code for more routine functionalities and allowing the engineers to focus more on design. In the video game industry, generative AIs could quickly produce numerous graphics assets that would otherwise need to be painstakingly produced by humans. The question is whether this type of benefit can also apply to learning environments.

Inherent in the view that creativity is enhanced by using generative AI is the assumption that an individual possesses creative potential but may lack the necessary time or concentration. This assumption is often violated for students; as stated earlier, developing the cognitive skills to "create" (the highest level in Bloom's taxonomy) is often the pedagogical goal. Thus, like with expertise development, the learner must be careful not to fully replace their own effort in the creative process. In our own trials, the AI could create a standard job advertisement (a typical creative assignment for an undergraduate student in HRM) based on the information given, fully replicating, and even improving upon the student responses to this same assignment. While this approach could save time for a busy HR professional, allowing them to concentrate on more creative tasks, it might eliminate the much-needed cognitive effort in creating and learning by a student.

Unfortunately, anecdotal evidence is already accumulating that users are employing generative AI to replace creative effort. We mentioned earlier the concern of faculty with a surge of essays, homework assignments, and student projects that appear to be produced by ChatGPT. Similar trends are becoming apparent outside of the classroom. For example, at Upwork, a platform where freelancers can contract work in the Gig economy (Shrivastava, 2023), employers experienced a large increase in volume of sub-standard proposals, which they suspected had been prepared using ChatGPT. Presumably, freelancers were using the AI to produce more quantity, at the expense of insight and creativity. Thus, generative tools are increasing plagiarism and tempting workers to tradeoff volume over quality.

Contrasting the human mind with LLMs, Chomski et al. (2023) stated, "The human mind is not, like ChatGPT and its ilk, a lumbering statistical engine for pattern matching, gorging on hundreds of terabytes of data and extrapolating the most likely conversational response or most probable answer to a scientific question. On the contrary, the human mind is a surprisingly efficient and even elegant system that operates with small amounts of information; it seeks not to infer brute correlations among data points but to create explanations." These ideas align with research on insights and creativity which shows that innovative problem-solving is a complex process and resides in the unique connections that the human mind makes (Dietrich & Kanso, 2010). In contrast, access to AI tools creates the temptation to gain more with less cognitive effort. Reduced mental effort can directly interfere with creativity and insights, as these processes rely on extensive thinking on the problem at hand, working through multiple incorrect approaches, and often reaching an impasse before correctly restructuring the problem and finding a breakthrough solution (Dietrich & Kanso, 2010; Haavold & Sriraman, 2022).

In addition, we must remember that AI generates content, be it text, music, or images, by drawing on large datasets of human-created content. ChatGPT, for example, is quite good at writing about topics that human writers have written a lot about (e.g., a job ad), and not so good at topics that are more niche or unexplored by human authors. The resulting output is only new in the sense that it does not simply copy existing work. Otherwise, the output reflects the status-quo, because the "correct" or "best" output is one that more closely reproduces the patterns observed in the training dataset. Human creators also take inspiration from the works of others, and every creative field, from arts to sciences, necessarily includes a foundation period of making oneself aware of the works of others. However, the end goal is to push the boundaries of what is already known, rather than to reshuffle existing content. AIs, at least in their current state, do more of the latter. This is illustrated by a recent study on creative writing, which found that AI assistance improved the quality of stories, particularly for less creative writers. However, it also made the stories more homogenous, compared to the group that did not rely on AI assistance (Doshi & Hauser, 2023)

As users publish more AI-generated content, innovation may be hindered even further. The output from generative AI is reductive due to lossy compression of knowledge, a process similar to saving a digital image multiples times, where each iteration results in information loss leading to a low quality image (Chiang, 2023). At this time, generative AIs are trained based on human-generated content that is publicly available. However, as users will inevitably create and publish work with the help of AI, these works may end

up as part of future training sets for the AI. Unless efforts are made to clearly identify AI-generated content, such that the training process is not biased, we could witness further deterioration of the content produced by generative AIs, as each iteration is a poorer approximation of the original set of human content from which it was derived.

The Paradox of Equity

Big-tech companies claim AI will democratize access to information and provide more equal access to economic opportunities; however, AIs can exacerbate the existing economic gap by disproportionately deskilling novices while upskilling experts.

During Microsoft's launch of the AI-powered Bing, one of the key highlights was the superiority of AI-powered internet search (Mehdi, 2023). Users would no longer have to navigate through various weblinks and read unhelpful content to find the specific information they were looking for. Instead, the AI would directly and succinctly produce an answer to the question asked. The implication was not only that everyone would benefit greatly from the time savings, but also that access to information would become more democratic: even those who did not know how to search efficiently (e.g., by using advanced search features and the right keywords) could now get the right answers from a search, because the question could be formulated in natural language. For example, the AI could assist a farmer unfamiliar with the bureaucratic application process necessary to access government subsidies or development funds.

While there are clearly applications of generative AI which are useful for everyone, we believe ultimately experts are likely to reap most of the benefits while avoiding most of the limitations, for at least three reasons: (1) the expertise and innovation paradoxes presented earlier affect novices more than experts; (2) novices are more likely to suffer from inaccurate AI output; and (3) jobs typically held by novices are at a greater risk of being replaced by AI compared to the roles occupied by experts. We discuss these further below.

Different Exposure to Expertise and Innovation Paradoxes. We argued earlier that generative AI can result in a replacement of human effort that is detrimental for the development of expertise and creative thought. By definition, those who have already accumulated expertise are less affected by the expertise paradox. Experts are also at an advantage with respect to creativity because existing knowledge facilitates learning knew knowledge (e.g.,

absorptive capacity, Cohen & Levinthal, 1990), and because new knowledge often results from recombining existing knowledge (Fleming et al., 2007).

In addition, experts can leverage AI capabilities better, because the type of research that AI can do well—that is, broad, but shallow summarizing of knowledge domains—is better suited for experts than novices, as illustrated in a study (Acar & van den Ende, 2016) cited earlier. Notably, experts were more successful at generating innovative solutions to problems when they performed broad searches outside their areas of expertise. Outsiders were better served by focusing on understanding the new knowledge domain; this can be riskier when relying on AI, because generative AIs are not yet capable of handling complex cognitive work, as our trials illustrated. For example, it could not properly interpret the implications of a large correlation for construct validity, it confounded similarly-worded constructs from two distinct theoretical domains, and it was prone to conjunction fallacy (Kahneman, 2017) unless specifically prompted to incorporate that concept into the response.

Thus, more experienced workers will be less reliant on AI assistance for their core competences, while still being able to augment their core tasks with AI. Furthermore, experts can also leverage AI assistance on more peripheral work, previously handled by less experienced human assistants, further reducing developmental opportunities for the latter.

Different Exposure to AI Limitations. As we have illustrated with our trials of ChatGPT, if we remain in the well-structured, algorithmic realms of coding, accounting, and bureaucracies, AIs have the potential to assist with and augment human work. We also referenced earlier some studies that showed AI could help less experienced workers be as effective as more experienced workers. However, in our trials, in areas where complex thinking was required, these tools provided highly misguiding answers or drew the wrong conclusions. Even in the best scenarios, AI could provide relevant and correct answers only when prompted in ways that required subject matter expertise. For example, the AI was not able to produce the correct implications of a very large correlation (.91) until prompted to also incorporate the concept of discriminant validity into its answer.

Correct and incorrect AI-generated answers were equally convincing, and therefore a non-expert would have no warning sign when reading the wrong conclusion. Experts could recognize and correct such errors, allowing them to avoid any negative consequences and to capitalize on productivity gains from AI without loss of quality. Errors are more likely to affect the work of novices, with potential career and economic consequences. We discussed earlier how technologies that facilitate access to information may inadvertently

create a meta-cognitive illusion: the user thinks they know more than they do. Generative AI adds another problematic layer: it may create confidence in solutions and answers that are incorrect.

Particularly, young adults are at risk for incurring disproportional harm from generative AI tools. In a recent LinkedIn article, Penot (2023) argues that members of Generation Z are already at a disadvantage, compared to millennials or Generation X, when it comes to technological savviness. Contrary to the popular belief, surveys find that Generation Z can be understood better as consumers of technology than savvy users (Penot, 2023). This is because many grew up with access to user-friendly technologies without the need to learn back-end mechanics, or to troubleshoot technical problems. In contrast, millennials and Generation X grew up alongside the development of technology, and often had to figure out creative ways to make technology work for their needs. Accustomed to utilizing technology as is, Generation Z members may be less likely to fact-check output, experiment with different prompts, or maintain a healthy skepticism regarding the technology.

Different Exposure to AI Job Replacement. It is no secret that for corporations, one of the biggest allures of generative AI is the efficiency, cost savings, and increased revenue streams. As a result, many entry level jobs are on the chopping block with the influx of generative AI tools capable of basic language tasks, including writing marketing content, drafting letters, automating customer service, document analysis, and product development (McKinsey & Company, 2023). This certainly increases economic opportunities for those who can successfully deploy these tools to cut costs and increase revenue streams. However, it can worsen economic disparities for the rest, particularly students. This economic paradox undermines the idea that higher education serves as a gateway to improved economic prospects for its graduates.

Entry-level positions often provide inexperienced students with opportunities to learn on the job and get mentored for advanced level positions. As this part of the talent pipeline will be replaced by generative AI tools, in the long run this deskilling could cause a shortage of high-skilled labor (and an oversupply of low-skilled labor). According to economic theory, demand and supply curves shift to form a new equilibrium where wages increase for high-skilled labor and decrease for low-skilled labor; this would lead to a larger wage gap between high and low-skilled workers. This widening inequity and unemployment can lead to economic instability and higher crime rates (Becker, 1968).

Recommendations for AI-Ready Management and Business Education

Through this essay, we put forth our thesis that the paradoxical nature of generative AI arises from its ability to enhance short-term learning and productivity in a way that could hinder students' long-term gains in expertise, innovation, and equity. Here we propose possible pathways to mitigate these risks. We focus on four key stakeholders: management and business students, faculty, administrators, and employers. Table 1 outlines our strategy for navigating various paradoxes among these groups. While we have identified the paradoxes each stakeholder is best equipped to address, this does not preclude their involvement in other areas. Given the interconnected nature of these paradoxes, efforts to address one may also positively impact others. Our goal is to provide initial insights into this complex interplay.

Student Guidelines for Sensible AI Use

When it comes to resolving the expertise paradox, a student's role cannot be emphasized enough. Students are key stakeholders in the rapidly evolving AI-based future, with their learning and long-term career outcomes potentially at stake. As the use of generative AI tools increases, students should be encouraged to consider the complex consequences of AI use. At the heart of the expertise paradox is the myopic emphasis on the immediate benefits of convenient AI-generated output. This emphasis detracts from students' commitment to the long-term goals of learning and building upon foundational skills. Drawing on theories of goal-setting (Locke & Latham, 2019) and valued-driven action (Bond et al., 2006), we suggest that students can mitigate the expertise paradox by explicitly identifying the skills they need to develop and understanding why they value these skills. By aligning their daily efforts with their long-term skill goals and values through a clearly articulated action plan, students might be able to minimize the expertise paradox.

Students must develop self-discipline and cultivate a personal code of ethics in the sensible use of generative AI. Supplemental Appendix D contains a model worksheet we have designed to assist students in this endeavor. This worksheet guides students through five sections, enabling them to assess their competencies, set specific improvement goals, and thoughtfully consider generative AI's role in their academic and professional development.

The first section, Identifying Goals and Values, requires students to express their academic and career aspirations. The second section, Core Competencies Assessment, requires a self-evaluation of essential skills aligned with one's career goals and values. The third, Identifying Sensible AI

Table 1. Recommendations to Minimize Three Paradoxes of Generative AI for Management and Business Students.

	Expertise paradox	Innovation paradox	Equity paradox

Student	• Develop personal code of sensible AI use:
	  ○ Identify development goals.
	  ○ Identify augmenting AI uses.
	  ○ Identify hindering AI uses.
	  ○ Identify ethical AI uses for skill development.
	  ○ Ongoing monitoring of plan.
		[a]	[a]

Faculty	• Instructional Innovation
	• Robust Assessments
	• Use AI to create learning materials:
	  ○ Create study guides.
	  ○ Develop practice questions and test banks.
	  ○ Use it for tutoring support.
		• Identify innovation
		• Facilitate Innovation
		• Reward Innovation
			• Curriculum content to create awareness on human-centric AI principles and responsible AI Governance.

University	[a]	• Enable experiential learning by providing resources and support. • Provide faculty and students access to AI tools.	• AI usage policies that upskill students. • AI Governance mechanisms to ensure positive impact on students. • Invest in technologies that enables upskilling of students. • Invest in faculty training. • Educate stakeholders on Generative AI benefits and harms. • Support faculty research and in-house analysis to understand the educational impact of AI.
Businesses	• Partner with universities to support experiential learning, service projects, and internships.	[a]	• Invest in talent development.

[a]This table illustrates which paradoxes each stakeholder is best equipped to address. This does not preclude stakeholders' involvement in other areas. These paradoxes are interconnected, efforts to address one may also minimize other paradoxes.

Uses, encourages students to explore how AI can enhance or potentially diminish their skills, aiming for a balanced perspective on AI usage. The fourth, Personal AI Ethics Code Development, provides a structure for creating personal guidelines for transparent and fair AI use. The final section, Adaptation and Continuous Learning, motivates the ongoing evaluation of sensible AI use by identifying strategies and resources for staying updated.

This worksheet serves as a comprehensive tool to foster thoughtful consideration of AI's role and ethical implications in professional development. It aids students in identifying critical skills and distinguishing between constructive and obstructive uses of AI, while also promoting sensible engagement with AI in a manner that supports their learning efforts. Through this process, students can gain a better understanding of AI's functionalities and limitations, potentially reducing the paradoxical impact of AI on expertise development.

To help students further, we provide some suggestions below. We believe these suggestions are sensible because the AI is used to complement (rather than replace) the student's learning and creative effort, by providing information input, as well as feedback on the outcomes of learning. When combined with instructor feedback, these uses allow the student to "mess around" and "geek out" with AI (Leaver & Srdarov, 2023), which would improve students' grasp of its workings and limitations. Sensible usage is likely to minimize the paradoxical harm of AI on students' expertise.

Sensible AI use for Learning

Copy Editing. Generative AIs are great for laborious, copy-editing work. Traditional document editing software are already able to check spelling, grammar, and suggest minor phrasing improvements. In addition, generative AI can rephrase entire paragraphs and can make a paragraph more concise and focused. Since this type of use relies entirely on the student's own draft, it is akin to utilizing resources already provided by universities, such as feedback from a writing center or a professor. Our suggestion presupposes that an average university student has mastered English grammar and spelling during their earlier educational years, and AI is used primarily to make minor corrections and improvements. However, a college student with significant writing deficiencies would be more successful in the long run by seeking foundational remediation to first develop enough writing expertise instead.

Draft Reviewing and Feedback. As demonstrated, generative AIs are also good at summarizing documents, offering a valuable tool for students to check if they have effectively communicated their main thesis. By prompting an AI to summarize or extract the main ideas from their drafts, students

could verify if their drafts effectively communicate each intended point. This method can provide useful early feedback before more substantive feedback is provided by instructors. Application examples for business courses include: for Organizational Behavior and Human Resource Management domains, reviewing job advertisement drafts to ensure the use of inclusive language; within Strategy and Entrepreneurship domains, reviewing business plan drafts for completeness, including standard elements like strengths, weaknesses, opportunities, and threats (SWOT); and in Finance/Accounting areas, reviewing drafts of financial reports, equity analysis, and recommendations.

Exploration and Research. The earlier points concerned students' seeking AI-assistance on work that they have produced. Students need to exercise caution when they use AI to seek information, explore topics that they are unfamiliar with, conduct research, or produce an output. In particular, students need to exercise good judgement to not substitute AI-assistance for developing their own knowledge and skills. In addition, as we have shown in Supplemental Appendix C, generative AIs can make mistakes in providing the right answers on topics that are more niche.

Nevertheless, for most well-established theories and domains of knowledge, students can use the AI to simply learn more about a topic. The advantage, compared to a traditional textbook-based approach, is that the student can guide the "conversation" toward topics of particular interest, request examples, ask for simpler or more complex explanations of a phenomenon, and even test one's understanding by requesting a quiz on the topic. In our experience, ChatGPT was effective at explaining well established theories in management, such as goal setting, expectancy, the Big Five model of personality, and the resource-based view of the firm. It could also support students with explanations and examples for related knowledge domains, such as statistics and research methods.

Checking Problem-Solving. We have warned in our essay about the potential pitfall of substituting hands-on problem solving with AI-generated solutions. Nevertheless, there are some legitimate uses that can leverage AI strengths and bolster learning, by combining the AIs capabilities for topic exploration and draft feedback. For example, we have shown that the AI is good at providing step-by-step advice on how to solve undergraduate-level accounting and finance problems (procedural knowledge typically involved in applying learned concepts in Bloom's taxonomy) and will likely get better at more advanced problems in future updates. Thus, students could use AI as a form of tutoring. First, the student would see a demonstration of the problem.

Then the student could use the AI to generate new instances of the problem, solve them independently, and then check their solution with the AI.

Given that AIs still have the potential to make mistakes that students may not be able to detect, some level of instructor involvement may still be necessary for this type of use. For example, ChatGPT worked less well when problems required more complex formulas (e.g., annuities). Moreover, in our trials, ChatGPT was especially good at writing code for business applications (i.e., statistical models) in general programming languages, such as Python and R, but struggled somewhat with more specialized software such as SAS and SPSS.

Producing Output on Peripheral Task. As we have argued through our essay, the key principle in the sensible use of AI is to not replace one's own effort on core learning tasks, and not to use AI to mask proficiency deficit in core competencies. For example, if a Human Resource student is struggling with comprehending and applying the adverse impact formula, they could use AI to further review the concept and example applications, but should not rely on it to solve related assignments in their place. In general, students should follow this approach for all class assignments where the instructor prohibits the use of AI in output production.

However, there may be peripheral tasks for which the use of AI to generate output might be warranted. For example, if a student is not aiming to learn programming in Python, but needs to write a simple code for a task (e.g., formatting all files in a certain way in an assignment folder), it would be legitimate to seek AI assistance in this area. Obviously, this very action would be considered cheating if renaming the files was an assignment in a specific class. Other examples may include relying on AI for graphics, images, animations, or sounds to go with one's PowerPoint presentation in a business class.

Design of Management Curriculum and Learning Assessments by Faculty

Faculty can play a critical role in how students engage with new technologies for learning, through their curriculum design and instruction. Faculty can take several steps to minimize all three paradoxes.

Minimizing the Expertise Paradox Through Instructional Innovation and Assessments. To reduce the expertise paradox, the curriculum design needs to emphasize learning foundational skills, and implementing robust ways of assessment that consider students' access to AI tools. In this regard, one of the

biggest challenges is students' use of unfair means to circumvent assessments. Educators are concerned about undetectable AI-assisted plagiarism (Cotton et al., 2024; Zhu et al., 2023). However, the interactive and generative nature of these tools may blur the line between cheating and legitimate use. Some have asked for outright bans of AIs on campuses, fearing an increase in student cheating (Clercq, 2023). Others have suggested we focus more on in-person, experiential, or meaningful learning experiences that are not as easily substituted by generative AIs (Geerling et al., 2023). Academia is shifting from banning to constructively incorporating AI into education, recognizing the need for responsible integration (Leaver & Srdarov, 2023).

The pedagogical approaches that are more resilient to student's access to generative AI tools support the augmentation view of AI (Tschang & Almirall, 2021), such as flipping the classroom by emphasizing in-person, experiential and service learning, less replaceable by generative AI (Geerling et al., 2023). In our trials, we found some support for the augmentation view. For example, ChatGPT could quickly teach one how to operate a financial calculator, saving precious time browsing the internet, reading a technical manual, or seeing a classroom demonstration for the same instructions. Business classes are particularly suited for such augmentation implementations, as they often include applications and exercises requiring students' mastery of technology and tools (e.g., Excel, PowerPoint, and statistical software). These tools, taught in various classes, may require refreshing and updating to align with the content currently being taught.

However, even with the use of instructional innovations, it is crucial that assessments incorporate foundational skills to mitigate the expertise paradox. Specifically, assessments should evaluate whether students understand and can independently apply discipline-specific concepts, without relying on AI tools. This may require more frequent use of proctored exams and oral examinations where appropriate (Cotton et al., 2024). Additionally, faculty members must ensure that their assignments and assessments emphasize critical thinking. This could involve in-class activities and learning labs where students apply critical thinking skills for problem-solving, with supervised access to AI.

For instance, faculty can adopt a two-phased approach. Initially, students should independently propose a solution to a problem, and then augment their responses with AI in the second phase. When used with well-established domains, this method can help students identify knowledge gaps or uncover blind spots. In our own teaching, we have used emerging AI research (e.g., Zhang & Parker, 2023) to help HRM students solve job design issues with AI support. Under the guidance of a subject matter expert, students can learn how to effectively prompt AI, tapping into a type of crowd wisdom, and

develop checklists of known causes for troubleshooting problems in established domains.

In addition to AI-resilient instructional innovations and robust assessments, faculty can also rely on generative AI to help students in learning foundational skills. We propose three ways that faculty can augment their efforts through AI tools.

Creating Study Guides. To facilitate student's mastery of foundational content faculty can use AI to help students engage with the course content by generating study guides and test bank questions. Generative AIs can be used to create study guides with relative ease. As experts, faculty are in a great position to evaluate the accuracy and completeness of such guides, supplement them with additional information that the AI may have missed, or recommend further readings. Used this way, AI may remove the need for a textbook, and thus help reduce the overall cost of education for students. In addition, it would support a flipped classroom by providing necessary materials for studying content at home.

Generating Practice Banks and Test Items. Another interesting use is in the generation of practice tests, questions, problems, or short cases. These can further support a flipped classroom with needed content to practice at home, but also with activities to be performed in the classroom.

Providing Tutoring Support. Many subjects involve the use of technology or software that requires additional training time in the classroom. Faculty can offload some of that training to AI, which can provide step-by-step instructions (procedural knowledge) on the use of financial calculators, Excel functions, statistical software, and general programming software. While students usually receive training on some of these tools in other classes, the instruction often occurs in isolation. AI can be utilized to refresh students' knowledge of various technology and tools, and to better integrate technological knowhow into the specific content domains taught (e.g., accounting and HRM).

Minimizing the Innovation Paradox by Identifying, Facilitating, and Rewarding Original Thinking. The curriculum design and assessments can also play a major role in reducing the innovation paradox. This paradox stems from using AI-generated output without investing necessary effort in creative problem-solving. Ultimately, this can lead to detrimental uniformity in thinking and outputs at a collective level. To counteract this pattern, the curriculum design and assessment needs to identify, facilitate, and reward original thinking and problem-solving.

Identify Original Thinking. Faculty must discern between instances when students are engaging in creative problem-solving and when they are over-relying on AI-generated outputs. As mentioned earlier, detecting plagiarism is hard with generative AI tools. However, running new assignments through AI tools might provide some idea of what a plagiarized assignment might look like. In developing guidelines for discouraging AI-assisted cheating, teaching assistants might be important conversation partners, as they navigate the dual role of students and teachers.

Facilitating Original Thinking. Perhaps one of the best ways of reducing the innovation paradox is to highlight the AI limitations in the classroom. Faculty should prepare AI-generated answers for complex questions and projects, encouraging students to compare these with their own solutions. This approach can foster a healthy skepticism toward AI-generated solutions by allowing students to observe firsthand the distinctions between AI-produced outcomes and those derived from human expertise. In addition, faculty should develop a portfolio of exemplary student solutions to showcase in class, offering a to AI-generated solutions. This approach can aid in understanding AI limitations in tackling creative tasks. The literature on fostering creativity in management education can be informative on designing curriculum that is robust to the innovation paradox (e.g., Schmidt-Wilk, 2011).

Rewarding Original Thinking. Related to the previous point, ChatGPT assignments may serve as a baseline "crowd wisdom" for students. Course assessments should emphasize originality to earn higher scores. A "do better than this" motto can direct students to think how to improve or to independently consider unique aspects that AI output may have overlooked. The grading rubrics should include criteria related to how similar or dissimilar a response is compared to AI-generated answers. Faculty can highlight the most creative answers to the class, both to acknowledge creative students and to further motivate creativity in all students.

Minimizing the Equity Paradox Through Curriculum Content. Finally, the equity paradox arises when novices compromise their long-term professional development and earning potential by prematurely relying on generative AI. To mitigate these risks, curriculum content can be pivotal. Ultimately, this discussion relates to the ethical and human-centric design and deployment of AI. As such, it is important for faculty to cover topics that deepen students' understanding of AI ecosystems. This is crucial, as business and management are critical human activity domains which are poised for substantial changes due to AI (Garibay et al., 2023).

Business education should create awareness about emerging management practices in the use AI of tools, including evaluating ground truth about the accuracy of AI tools (Lebovitz et al., 2023), exercising organized skepticism in AI implementation decisions (Vinsel, 2023), understanding data privacy risks, seeking informed consent when using AI tools, and being transparent when augmenting work with AI assistance (Zhu et al., 2023). While these issues are universally relevant, they are especially relevant for business programs. Business graduates may later be responsible for AI implementations in various organizations. Hence, educating students about organized skepticism (Vinsel, 2023) and the commercial roots of AI-hype is likely to cultivate a more balanced approach in adopting AI. From our experience, incorporating these topics into the curriculum has motivated students to reflect on their role in AI implementation. For instance, one undergraduate student noted in their assignment "I will do my own research before using any AI system, I don't want to support AI developers who abuse the power of AI and are not using human-centric design principles."

In addition, management curriculum should cover topics about responsible and ethical AI governance, principles of designing responsible AI business systems, frameworks for human supervision in AI decision-making, environmental costs of these models, and managing occupational impact through stakeholder perspective (Bubeck et al., 2023; Garibay et al., 2023; Leaver & Srdarov, 2023; Tschang & Almirall, 2021). Curriculum content can help students appreciate that generative AI is a complex and evolving technology with significant ethical, social, economic, and environmental implications (Leaver & Srdarov, 2023).

University Administrators Shaping the Broader Context for Deeper Learning

The university administration can play a critical role in minimizing the innovation and equity paradox of generative AI. The advent of generative AI unfolds in already challenging times for US higher education. Universities have drifted away from training students to be thinkers and becoming better persons, and not merely technical experts with degrees (Bosch, 2018). Instead of fostering collaboration among institutions to promote education, curriculum design is excessively driven by competition and consumerism (Christensen & Eyring, 2011; Lewis, 2006; Lund Dean et al., 2020). Amid these concerns, students now have access to AI tools, which are here to stay. If higher education does not cultivate deep thinking skills among students, they risk facing limited prospects and growing inequities due to deskilling.

To reduce the innovation paradox, administrators need to consider whether certain trends in higher education, such as accelerated courses and large class sizes, are still feasible for the effective assessment of learning. Higher education faces many external pressures, including limited state funding (Mintz, 2021). In addition, these institutions face increased competition for students and faculty, technological advancements requiring costly updates, and a growing demand for online and flexible learning options (Christensen & Eyring, 2011). These factors collectively strain budgets, impacting the quality and accessibility of education offered (Carnevale et al., 2020).

Time is pressing us to confront difficult questions: Are business school programs emphasizing experiential learning, critical thinking, and significant learning? Does the value proposition of our curriculum go beyond the first three levels of Bloom's taxonomy (the very levels at which generative AI tools seem to do at par with students, and may potentially make cheating rewarding)? How do we develop a learning environment that encourages students to be deep divers rather than jet skiers (Carr, 2010) when learning? Hopping from topic to topic, simply consuming information, rather than spending time analyzing and interpreting it, are some of the habits that our curriculum needs to challenge. We ought to create learning communities where students appreciate that education is not just simply power browsing, skim reading, data processing and predicting, but that contemplation, reflection, sense-making, and deliberation are key learning processes in andragogy (Schwandt, 2005).

Management scholars have been advocating for various pedagogical approaches (e.g., experiential learning) to develop moral reasoning and critical thinking skills (Bagley et al., 2019; Bevan & Kipka, 2012; Ghoshal, 2005; Statler, 2014). The advent of generative AIs and the paradoxes they bring for education makes it more urgent to identify and remove any systemic barriers to a wider implementation of effective pedagogical approaches. Addressing these issues will be essential in resolving the paradox of innovation and may increase the possibility that AI will augment human work, rather than deskill the workforce and exacerbate inequity.

Curriculum design that emphasizes critical thinking, deeper learning, and reflection requires that faculty give considerable individualized attention to each learner. However, increased class sizes have consistently pushed classroom instruction in the opposite direction. Moreover, in the last twenty years, the internet has promoted the fast and shallow consumption of information rather than the deep study of a topic (Carr, 2010, 2020). A recent popular post on LinkedIn, suggesting effective ChatGPT prompts, reflects this mindset: "I want to learn about [insert topic]. Identify and share the most important 20% of learnings from this topic that will help me understand 80% of it." Thus,

there are significant structural and attitudinal constraints to implementing a deeper learning curriculum that can minimize paradoxes of innovation and inequity. These challenges to student learning must be addressed by administrators and other key decision-makers at the state and federal level.

In addition to supporting faculty in developing and utilizing pedagogical approaches that are robust to the arrival of AI, administrators should also consider how to effectively incorporate this new technology into the classroom. At a minimum, universities need to ensure faculty and students have equitable campus access to the latest AI technologies, especially when some of these tools are locked behind paywalls. Administrators should keep abreast of relevant developments in the field, including on potential education applications, and disseminate information to faculty and students. Finally, universities should ensure they have both the technological base and intellectual capital needed to teach new applications of AI.

One of our key recommendations for avoiding the AI paradoxes was to develop sensible-use guidelines. Many organizations are developing governance policies and appointing personnel to oversee the responsible use of AI. Universities must develop similar institutional mechanisms (e.g., hiring staff with expertise on AI governance and instructional technology) to create uniform guidelines and practices with the mission to upskill students, while allowing the flexibility required to experiment with AI in the classroom.

Furthermore, the potential of AI in education, particularly for personalized feedback and lesson plans, must be carefully assessed for its impact on student learning. Universities must monitor trends in assessment of learning statistics and adjust guidelines appropriately, if these statistics indicate any of the issues we raised in our essay. In this respect, administrators should also encourage and support faculty research about the impact of AI on student learning.

Employers Investing in Talent Development

Reactions to the release of generative AI often compare this technology to the introduction of automation in manufacturing: employers will benefit, because of lower labor costs and greater productivity, but many employees will be negatively affected because of job cuts. We argued in this essay that there may also be an unrecognized cost to employers. Unlike automation in manufacturing, AI can replace skills and jobs that are foundational for the development of other skills and jobs which AI is not yet or may never be able to replace. Thus, employers should carefully consider how AI deployment may affect the development of core competences in employees.

As scholarly work creates awareness regarding the potential harm of over-reliance on AI, we must call for businesses to take a more active role in talent development by supporting experiential learning, field projects, service learning, and internships in colleges. This requires a culture of collaborative learning and university-industry partnership. Perhaps the corporate social responsibility ratings should include parameters gauging firm partnerships with universities, and maybe even the responsible use of AI in talent development. These factors are likely to incline reputable and profitable companies to comply and promote a cooperative talent development culture.

Such measures are likely to support the upskilling of students as they enter the job market. These steps, when combined with continued investment in developing skills at entry level jobs (rather than myopic replacement of raw talent with AI) are likely to minimize the disproportional economic impact on novices. Thus, through strategic talent development and partnerships, employers can help reduce the expertise and equity paradoxes.

In conclusion, we discussed paradoxes of an AI-facilitated world, where without concerted effort by all stakeholders, our students risk deskilling. Moving forward, we recommend that business schools focus on exploring innovative ways to integrate technology and value-consciousness into management education. We live in a world supersaturated with information and content. Helping students differentiate between consuming content versus meaningfully engaging with it has become ever more important in the AI-powered world. The classroom experience should facilitate students' climb up the learning ladder, where data and information serve as tools to value-driven action. With each ascending step the learner develops a deeper appreciation that data is not information, information is not knowledge, knowledge is not action, and action is not purpose. Technological progress in generative AI should prioritize and facilitate student learning and purposeful living.

Acknowledgments

The authors thank their editor John Stark for his exceptional guidance and the two anonymous reviewers for their constructive feedback.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Maria Riaz Hamdani https://orcid.org/0000-0002-9926-9776

Supplemental Material

Supplemental material for this article is available online.

References

Acar, O. A., & van den Ende, J. (2016). Knowledge distance, cognitive-search processes, and creativity: The making of winning solutions in science contests. *Psychological Science*, *27*(5), 692–699.

Anderson, L. W., Krathwohl, D. R., Airasian, P. W., Cruikshank, K. A., Mayer, R. E., Pintrich, P. R., Raths, J., & Wittrock, M. C. (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.

Athanassiou, N., McNett, J. M., & Harvey, C. (2003). Critical thinking in the management classroom: Bloom's taxonomy as a learning tool. *Journal of Management Education*, *27*(5), 533–555.

Bagley, C. E., Sulkowski, A. J., Nelson, J. S., Waddock, S., & Shrivastava, P. (2019). A path to developing more insightful business school graduates: A systems-based, experimental approach to integrating law, strategy, and sustainability. *Academy of Management Learning & Education*, *19*(4), 541–568. https://doi.org/10.5465/amle.2018.0036

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, *76*(2), 169–217.

Bevan, D., & Kipka, C. (2012). Experiential learning and management education. *Journal of Management Development*, *31*(3), 193–197.

Bloom, B. S., Engelhart, M. D., Furst, E., Hill, W. H., & Krathwohl, D. R. (1956). *Handbook I: Cognitive domain*. David McKay.

Bond, F. W., Hayes, S. C., & Barnes-Holmes, D. (2006). Psychological flexibility, ACT, and organizational behavior. *Journal of Organizational Behavior Management*, *26*(1–2), 25–54.

Bosch, G. (2018). Train PhD students to be thinkers not just specialists. *Nature*, *554*, 277.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work*. NBER.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., & Lundberg, S. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.

Carnevale, A. P., Schmidt, P., & Strohl, J. (2020). *The merit myth: How our colleges favor the rich and divide America*. The New Press.

Carr, N. (2010, June 05). Does the Internet make you dumber. *Wall Street Journal*. https://www.wsj.com/articles/SB10001424052748704025304575284981644790098

Carr, N. (2020). *The shallows: What the Internet is doing to our brains*. WW Norton & Company.

Chiang, T. (2023, February 9). ChatGPT is a blurry JPEG of the web. *The NewYorker*. https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web

Chomski, N., Roberts, I., & Watumull, a. J. (2023, March 8). Guest essay: The false promise of ChatGPT. *International New York Times*. https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

Christensen, C. M., & Eyring, H. J. (2011). *The innovative university: Changing the DNA of higher education from the inside out*. John Wiley & Sons.

Clercq, G. D. (2023, January 27). Top French university bans use of ChatGPT to prevent plagiarism. *Reuters*. https://www.reuters.com/technology/top-french-university-bans-use-chatgpt-prevent-plagiarism-2023-01-27/

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, *35*, 128–152.

Cotton, D. R., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*, 228–239.

Dietrich, A. (2004). The cognitive neuroscience of creativity. *Psychonomic Bulletin & Review*, *11*(6), 1011–1026. https://doi.org/10.3758/BF03196731

Dietrich, A., & Kanso, R. (2010). A review of EEG, ERP, and neuroimaging studies of creativity and insight. *Psychological Bulletin*, *136*(5), 822–848. https://doi.org/10.1037/a0019749

Doshi, A. R., & Hauser, O. (2023). *Generative artificial intelligence enhances creativity*. SSRN. https://doi.org/10.2139/ssrn.4535536

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217.

Fink, L. D. (2013). *Creating significant learning experiences: An integrated approach to designing college courses*. John Wiley & Sons.

Fleming, L., Mingo, S., & Chen, D. (2007). Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly*, *52*(3), 443–475.

Garibay, O. O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., Falco, G., Fiore, S. M., Garibay, I., Grieman, K., Havens, J. C., Jirotka, M., Kacorri, H., Karwowski, W., Kider, J., Konstan, J., Koon, S., Lopez-Gonzalez, M., Maifeld-Carucci, I., . . . Xu, W. (2023). Six human-centered artificial intelligence grand challenges. *International Journal of Human–Computer Interaction*, *39*(3), 391–437. https://doi.org/10.1080/10447318.2022.2153320

Geerling, W., Mateer, G. D., Wooten, J., & Damodaran, N. (2023). Is ChatGPT smarter than a student in principles of economics? *SSRN Electronic Journal*.

Ghoshal, S. (2005). Bad management theories are destroying good management practices. *Academy of Management Learning & Education*, *4*(1), 75–92.

Haavold, P. Ø., & Sriraman, B. (2022). Creativity in problem solving: Integrating two different views of insight. *ZDM – Mathematics Education*, *54*(1), 83–96. https://doi.org/10.1007/s11858-021-01304-8

Heilman, K. M., Nadeau, S. E., & Beversdorf, D. O. (2003). Creative innovation: Possible brain mechanisms. *Neurocase*, *9*(5), 369–379. https://doi.org/10.1076/neur.9.5.369.16553

Hillmann, J., Duchek, S., Meyr, J., & Guenther, E. (2018). Educating future managers for developing resilient organizations: The role of scenario planning. *Journal of Management Education*, *42*(4), 461–495.

Kahneman, D. (2017). *Thinking, fast and slow*. Macmillan.

Kolb, A. Y., & Kolb, D. A. (2005). Learning styles and learning spaces: Enhancing experiential learning in higher education. *Academy of management learning & education*, *4*(2), 193–212.

Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice Hall.

Lakhani, K. (2023). *AI won't replace humans: But humans with ai will replace humans without AI* [Interview]. HBR. https://hbr.org/2023/08/ai-wont-replace-humans-but-humans-with-ai-will-replace-humans-without-ai

Le, H., Schmidt, F. L., Harter, J. K., & Lauver, K. J. (2010). The problem of empirical redundancy of constructs in organizational research: An empirical investigation. *Organizational Behavior and Human Decision Processes*, *112*(2), 112–125.

Leaver, T., & Srdarov, S. (2023). ChatGPT isn't magic: The hype and hypocrisy of generative artificial intelligence (AI) rhetoric. *M/C Journal*, *26*(5), 5. https://doi.org/10.5204/mcj.3004

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2023). The no. 1 question to ask when evaluating AI tools. *MIT Sloan Management Review*, *64*(3), 27–30.

Lewis, H. (2006). *Excellence without a soul: How a great university forgot education*. Publication Affairs U.S.

Locke, E. A., & Latham, G. P. (2019). The development of goal setting theory: A half century retrospective. *Motivation Science*, *5*(2), 93–105. https://doi.org/10.1037/mot0000127

Lovelace, K. J., Eggers, F., & Dyck, L. R. (2016). I do and I understand: Assessing the utility of web-based management simulations to develop critical thinking skills. *Academy of Management Learning & Education*, *15*(1), 100–121.

Lund Dean, K., Fornaciari, C., Bento, R., & Asarta, C. (2020). Premises, promises, and perils of the Academic Potemkin Village. *Management Learning*, *51*(4), 491–510. https://doi.org/10.1177/1350507620903171

Mazhari, M. (2023, March 24). Chomsky says AI advancements have paved the way for high-tech plagiarism. *Tehran Times*. https://www.tehrantimes.com/news/483116/Exclusive-Chomsky-says-AI-advancements-have-paved-the-way-for

McKinsey & Company. (2023). *The state of AI in 2023: Generative AI's breakout year*. https://McKinsey.com

Mehdi, Y. (2023). *Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web*. https://blogs.microsoft.com/blog/2023/02/07/

reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web

Mintz, B. (2021). Neoliberalism and the crisis in higher education: The cost of ideology. *American Journal of Economics and Sociology*, *80*(1), 79–112.

Noy, S., & Zhang, W. (2023). *Experimental evidence on the productivity effects of generative artificial intelligence* (SSRN 4375283). SSRN.

OpenAI. (2022). *Introducing ChatGPT*. https://openai.com/blog/chatgpt#OpenAI

Penot, M. (2023). Gen Z is not tech-savvy. https://www.linkedin.com/pulse/gen-z-tech-savvy-mathieu-penot%3FtrackingId=YOerK0KkR0uOZbIzjFsdRA%253D%253D/?trackingId=lH75I1x%2BTJuT2zimHK4PUQ%3D%3D

Polanyi, M. (1966). *The tacit dimension*. University of Chicago Press.

Reeves, M. F. (1990). An application of Bloom's taxonomy to the teaching of business ethics. *Journal of Business Ethics*, *9*, 609–616.

Roche, B., Barnes-Holmes, Y., Barnes-Holmes, D., Stewart, I., & O'Hora, D. (2002). Relational frame theory: A new paradigm for the analysis of social behavior. *The Behavior Analyst*, *25*, 75–91.

Schmidt-Wilk, J. (2011). Fostering management students' creativity. *Journal of Management Education*, *35*(6), 775–778.

Schmidt, E. (2023). *Fmr. Google CEO eric Schmidt on the AI revolution* [Interview]. Public Broadcast Radio. https://www.pbs.org/video/fmr-google-ceo-eric-schmidt-ai-revolution-2mknqb/

Schwandt, D. R. (2005). When managers become philosophers: Integrating learning with sensemaking. *Academy of Management Learning & Education*, *4*(2), 176–192. https://doi.org/10.5465/AMLE.2005.17268565

Shen, W., Luo, J., Liu, C., & Yuan, Y. (2013). New advances in the neural correlates of insight: A decade in review of the insightful brain. *Chinese Science Bulletin*, *58*(13), 1497–1511. https://doi.org/10.1007/s11434-012-5565-5

Shrivastava, R. (2023, April 20). 'I've never hired A writer better than ChatGPT': How AI is upending the freelance world. *Forbes*. https://www.forbes.com/sites/rashishrivastava/2023/04/20/ive-never-hired-a-writer-better-than-chatgpt-how-ai-is-upending-the-freelance-world/?sh=72f91ace62be

Sidoti, O., & Gottfried, J. (2023). *About 1 in 5 US teens who've heard of ChatGPT have used it for schoolwork*. Pew Research Center. https://policycommons.net/artifacts/8245911/about-1-in-5-us/9162789/

Hinton, G., Bengio, Y., Hassabis, D., Altman, S., Amodei, D., Song, D., Lieu, T., Gates, B., Zhang, Y., Sutskever, I., Babuschkin, I., Legg, S., Hellman, M., Manyika, J., Zeng, Y., Zhan, X., Efimov, A., Graylin, A. W., Zhang, J., . . . *Statement on AI Risk—Open Letter*. (n.d.). https://www.safe.ai/statement-on-ai-risk#open-letter

Statler, M. (2014). Developing wisdom in a business school? Critical reflections on pedagogical practice. *Management Learning*, *45*(4), 397–417.

Tschang, F. T., & Almirall, E. (2021). Artificial intelligence as augmenting automation: Implications for employment. *Academy of Management Perspectives*, *35*(4), 642–659. https://doi.org/10.5465/amp.2019.0062

Vara, V. (2023, September 21). Confessions of a viral AI writer. *Wired*. https://www.wired.com/story/confessions-viral-ai-writer-chatgpt/

Vinsel, L. (2023). Don't get distracted by the hype around generative AI. *MIT Sloan Management Review*, *64*(3), 1–3.

Wegner, D. M., & Ward, A. F. (2013). How Google is changing your brain. *Scientific American*, *309*(6), 58–61.

Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Biles, C., Birhane, A., Haas, J., Rimell, L., Hendricks, L. A., . . . Gabriel, I. (2021). Ethical and social risks of harm from language models. *arXiv*. https://doi.org/10.48550/arXiv.2112.04359

Zhang, F., & Parker, S. K. (2023). How ChatGPT can and can't help managers design better job roles. *MIT Sloan Management Review*, *65*(1), 1–5.

Zhu, C., Sun, M., Luo, J., Li, T., & Wang, M. (2023). How to harness the potential of ChatGPT in education? *Knowledge Management & E-Learning*, *15*(2), 133–152. https://doi.org/10.34105/j.kmel.2023.15.008