


**Harvard
Business
Review**

Technology And Analytics

How Leading Consultancies Can Better Manage AI Risk

by Brian Spisak

November 06, 2023



Claude Garipey/Getty Images

Summary. In an age when AI-powered tools are reshaping industries, consultancies are embracing the potential of AI copilots to revolutionize their services. From PwC's ChatPwC to McKinsey's Lilli, KPMG's KymChat, and more, the surge of AI integration echoes historical inflection points like the internet's global connectivity and the industrial revolution's manufacturing shifts. As AI copilots redefine work processes and business competition, questions arise about responsible innovation. Will this AI arms race drive equitable practices or lead to hasty shortcuts? The answers lie in collaborative initiatives, much like Big Tech's voluntary AI commitments, which offer a roadmap for consultancies to harness AI's potential while ensuring ethical conduct. Through transparent reporting, collaborative risk management, and shared best practices, the consultancies can navigate the AI revolution responsibly, shaping the trajectory of business and society in a meaningful and lasting way. [close](#)

Generative AI is reshaping the consulting industry. From PwC's ChatPwC to McKinsey's Lilli to KPMG's KymChat, and more, AI copilots are redefining work processes and business competition. But will this AI arms race drive equitable practices or lead to hasty shortcuts? The answers lie in collaborative initiatives, much like Big Tech's voluntary AI commitments, which offer a roadmap for consultancies to harness AI's potential while ensuring ethical conduct. Through transparent reporting, collaborative risk management, and shared best practices, the consultancies can navigate the AI revolution responsibly, shaping the trajectory of business and society in a meaningful and lasting way.

ChatPwC was designed to assist employees in their work by allowing them to rapidly find the right people for a specific task, leverage more of the company's massive knowledge base, draft stronger proposals, and ultimately deliver next-level service. McKinsey & Company and KPMG launched Lilli and KymChat, respectively, for similar purposes. EY is working on an "Intelligent Payroll Chatbot," and Deloitte has formed a strategic alliance with Google to broadly advance generative AI capabilities.

Given the size and reach of these prominent consultancies, the idea of enhancing work processes through their AI-powered tools holds immense significance. These initiatives not only reflect the evolving landscape of business operations but also set a precedent for leveraging cutting-edge technology to drive efficiency, innovation, and superior service delivery across all industries and facets of business. In other words, these initiatives have the potential to reshape how all businesses operate, transcending sectors and reshaping the way we work, collaborate, and innovate.

Why This Matters

The direct, explicit, and frequent interaction with generative AI at scale is important for many reasons. It's a fundamental transformation of what workers do and how they accomplish tasks. The emergence of AI copilots marks a significant shift in how we conduct and compete for business. Imagine a scenario where a team of consultants is preparing a comprehensive market analysis and strategy for a client. In the traditional approach, this

process would involve hours of research, data collection, and analysis. With the integration of AI copilots, the workflow changes significantly.

The consultants can now engage with the AI copilot to quickly gather vast amounts of relevant market data, industry trends, and competitor analysis. The AI copilot can process and distill this information into concise summaries, visualizations, and even preliminary insights. This rapid data processing frees up the consultants' time, allowing them to focus on higher-level strategic thinking and creative problem-solving.

In this way, AI copilots shift how business is conducted by expediting data-driven decision-making, fostering quicker and more informed strategy development, and augmenting the creative and communicative aspects of consulting work. The interaction between consultants and AI copilots becomes a collaborative partnership that leverages the strengths of both human expertise and AI capabilities, ultimately leading to more efficient, effective, and innovative business practices.

However, avoiding over-reliance on generative AI in content development is crucial for preserving quality. Excessive dependence leads to repetitive, formulaic content. Essentially, it becomes a useless photocopier churning out copies of copies. Generative AI, as remarkable as it is, requires input from people for optimal performance. Without fresh real-world data, it will produce homogenized output, leading to a decline in quality and

even, as recent research from scientists at Rice and Stanford suggests, the failure of the AI model itself. Striking a balance between human creativity and AI capabilities ensures a harmonious and sustainable collaboration that leverages the best of both worlds.

The AI revolution is also a test of responsibility. The fierce competition between companies developing AI-powered business copilots poses a crucial test between speed versus safety. Will the battle of the consulting bots, for example, lead to reckless shortcuts and AI elitism, or will it steer towards democratized and equitable practices, enhancing the performance and well-being of workers globally? Will this AI arms race foster unparalleled innovation or result in a reckless scramble to the bottom, where the pursuit of AI supremacy takes precedence over responsible development?

How these companies tackle challenges related to data bias, diversity, and quality, and how they mitigate inherent risks tied to large language models (LLMs), such as hallucinations and prompt hacking will shape the trajectory of AI adoption and determine whether the benefits can be harnessed without exacerbating existing inequalities or compromising ethical business practices. This is immediately significant because businesses worldwide are increasingly turning to consultants for guidance on AI adoption. In the U.S., for instance, IT consulting revenue is projected to grow by 4% year-over-year, a notable increase compared to the

annual average growth rate of 2% over the past five years. Simply put, the solutions they develop have a far-reaching impact, affecting us all.

The Trajectory of AI

Ideally, an open-source collaboration between major consultancies will emerge to share best practices in constructing secure and responsible AI-powered platforms. They can compete for clients based on product quality and service delivery while working together to address the significant challenges posed by LLMs.

Opting for a cooperative path is both prudent and ultimately easier in the long term relative to short-term competitive isolation, as it allows for the shared mitigation of challenges, the pooling of resources, and the establishment of industry-wide ethical standards that benefit all stakeholders. An open-source collaboration between major consultancies in the development of AI-powered copilots and platforms delivers several key benefits:

Shared Best Practices: An open-source collaboration would facilitate the sharing of best practices in AI development, implementation, and ethical considerations. This would lead to the creation of more secure and reliable AI-powered platforms.

Ethical Frameworks: Working together would allow the consulting firms to collectively establish comprehensive ethical guidelines for AI usage. This could include guidelines to ensure fairness, privacy, and mitigation of biases in AI algorithms.

Risk Mitigation: Collaboratively addressing challenges related to data quality, security, and the risks associated with large language models would reduce the potential negative impacts of AI technologies. This could lead to more responsible and reliable AI solutions.

Industry Reputation: An open-source collaboration focused on responsible AI development could enhance the overall reputation of the consulting industry. Clients and stakeholders would perceive these firms as leaders in ethical AI adoption, which could attract more business and partnerships.

Transparency: An open-source approach encourages transparency in AI development. By openly sharing code, algorithms, and methodologies, consulting firms can build trust with clients and the wider public.

The 8 Steps

To reap these collective benefits, large consultancies developing AI-powered business tools should consider adopting a framework similar to Big Tech's voluntary AI commitments. While the sustainability of Big Tech's commitments has been called into

question, such a proactive step is essential for building a foundation of trust and long-term viability in the rapidly evolving AI landscape. Without it, I suspect regulatory interventions—which will stifle innovation—are inevitable.

To avoid this outcome, I devised a set of voluntary commitments adapted from Big Tech’s self-imposed obligations to manage AI risk:

1. **Pre-Deployment AI Quality Assurance:** Implementing thorough pre-deployment AI testing procedures that are specifically designed to align with the consultancies’ services, ensuring security, functionality, and ethical standards.
2. **Collaborative AI Risk Management:** Collaborating extensively within the consultancies to share knowledge and strategies for effectively managing AI-related risks, strengthening the collective ability to mitigate potential challenges.
3. **Strategic Investments in AI Security Measures:** Allocating resources to develop and apply cybersecurity measures customized for the consultancies, ensuring the protection of internal systems and client interactions.
4. **Exploration of AI Vulnerability Programs:** Exploring the viability of “bug bounty” programs tailored to the consultancies’ AI applications, encouraging external experts to identify and report vulnerabilities responsibly.

5. **Innovations in AI Authenticity Indicators:** Researching and developing technical solutions like watermarking systems to allow clients to identify AI-generated content, ensuring transparent communication in consultancy services.
6. **Transparent Reporting of AI Capabilities:** Committing to transparently communicate the abilities and limitations of AI systems used in consultancy services, enabling clients to make informed decisions.
7. **Societal Impact Research:** Conducting research on the specific societal impacts of AI within the context of consultancy services, addressing concerns such as AI bias and privacy violations and proposing suitable solutions.
8. **Elevated Integration of Advanced AI Systems in Operations:** Expanding the integration of advanced AI systems across consultancy functions to improve efficiency, precision, and innovation, ultimately enhancing the safety and quality of services provided by the consultancies.

By embracing these commitments, these companies are charting a course toward responsible AI adoption, emphasizing ethical considerations, transparency, and client-centricity. In an era where AI transformation is inevitable, these commitments will guide them in shaping the future of AI-powered business practices with integrity and foresight. It's now up to them to make the right decision.

Brian R. Spisak is a Research Associate at Harvard's National Preparedness Leadership Initiative and an independent consultant. He is the author of ***Computational Leadership: Connecting Behavioral Science and Technology to Optimize Decision-Making and Increase Profits*** (Wiley, 2023).

Recommended For You

Life's Work: An Interview with Patti Smith



Life's Work: An Interview with Jacques Pépin



Life's Work: An Interview with Chita Rivera



PODCAST
How the United States Air Force Accelerated AI Adoption

