


BEYOND ATTRIBUTION

The Ethics of Pretending — AI-Generated Thought Leadership and the Authenticity Question

A Practitioner's Commentary on What Happens When the Thinking Is Outsourced

Giovanni Leonardi

September 2025



The Ethics of Pretending — AI-Generated Thought Leadership and the Authenticity Question

Leadership · No. 119
September 2025

“The uncomfortable question is not whether AI can write thought leadership but whether thought leadership written by AI is still thought leadership at all.”

The Quiet Shift

Something is happening to thought leadership that the profession has not yet reckoned with. Across consulting firms, advisory practices, and individual professional brands, an increasing proportion of the white papers, articles, and opinion pieces published under a practitioner's name are being generated, in whole or in substantial part, by AI. The tools are capable enough that the output is plausible. The economics are compelling enough that the practice is spreading. And the disclosure is, in almost every case, absent.

This is not a technology problem. It is an honesty problem. And it matters because thought leadership — the practice of sharing hard-won professional insight to establish credibility, advance a discipline, and contribute to serious debate — depends on a contract between author and reader that AI-generated content quietly violates.

The Contract

The contract is implicit but real. When a senior practitioner publishes a paper on programme governance or transformation leadership, the reader makes a set of assumptions: that the ideas reflect the author's own experience, that the arguments have been shaped by years of practice, and that the perspective carries the weight of someone who has actually done the work. The reader may disagree with the analysis, but they engage with it because they believe it comes from a place of genuine professional knowledge.

AI-generated thought leadership breaks this contract in a specific way. The output may be well-structured, articulate, and even insightful in a generic sense. But it does not come from experience. It comes from pattern recognition across training data. It cannot tell the reader what it was like to navigate a failing programme, because it has never navigated one. It cannot offer the kind of earned wisdom that comes from watching the same mistake play out across five organisations over a decade, because it has not watched

anything. It can simulate the form of practitioner insight without possessing the substance.

The defence most commonly offered is that the AI is merely a drafting tool — the human provides the ideas, the AI provides the prose. In some cases, this is true and unproblematic. A practitioner who uses AI to polish their own thinking, to structure an argument they have already developed, or to overcome the blank-page problem is using a tool in service of genuine thought. The ideas are theirs; the execution is assisted.

But in my observation, this is increasingly not what is happening. What is happening is that practitioners are providing a topic and a tone and allowing the AI to generate the argument itself — the analysis, the framework, the recommendations. The human contribution is editorial rather than intellectual. The name on the paper lends authority to ideas the named author did not have.

Why It Matters

The immediate harm is to the reader, who is making decisions — about which consultant to hire, which approach to adopt, which risks to take seriously — on the basis of perceived expertise that may not exist. A well-written AI paper on transformation governance may look identical to one written by a practitioner with twenty years of relevant experience, but its reliability as a guide to action is fundamentally different.

The deeper harm is to the discipline itself. Thought leadership has always been uneven — some of it is brilliant, much of it is mediocre, and a fair amount is self-serving marketing dressed in analytical clothing. But even at its most imperfect, it has served as a mechanism for practitioners to share what they have learned, to challenge each other's assumptions, and to advance collective understanding. When the signal is diluted by machine-generated content that mimics the form without contributing genuine insight, the mechanism degrades. Readers learn to trust less. Genuine contributors are drowned out by volume. The discourse becomes shallower precisely when the challenges facing organisations demand depth.

There is also a professional integrity dimension that is uncomfortable to name but important to face. A practitioner who publishes AI-generated content under their own name, without disclosure, is engaged in a form of misrepresentation. The severity varies — a LinkedIn post drafted with AI assistance is different from a white paper that forms the basis of a consulting engagement — but the direction is the same. The practitioner is claiming intellectual authorship of work they did not do.

What Honesty Requires

The path forward is not to prohibit AI from the writing process. That would be both impractical and unnecessary — AI is a legitimate tool for improving the quality and efficiency of genuine human thought. The path forward is honesty about the boundary between human thinking and AI production.

This means, at minimum, three things. First, practitioners must be honest with themselves about whether they are using AI to express their own ideas or to generate ideas they do not have. The distinction is real and they know which side they are on. Second, organisations that publish thought leadership should develop clear standards about AI involvement and disclosure — not to police the use of AI, but to protect the credibility that makes thought leadership worth publishing. Third, readers should become more demanding: asking not just whether a piece is well-written but whether there is genuine experience behind it, and learning to distinguish between the patterns of real insight and the patterns of sophisticated imitation.

The uncomfortable question is not whether AI can write thought leadership. It clearly can, in the narrow sense of producing text that looks and reads like thought leadership. The question is whether thought leadership written by AI is still thought leadership at all — or whether it is something else entirely, wearing a borrowed authority it has not earned.