


BEYOND THE BYLINE

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

A Practitioner's Commentary on the Growing Gap Between Authorship and Attribution

Giovanni Leonardi

September 2025



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

Leadership · No. 106
September 2025

"The test is not 'did you write every word?' but 'would you defend every claim in a room full of people who know the subject as well as you do?'"

The Quiet Fraud

There is a genre of content proliferating across professional networks and industry publications that presents itself as the considered opinion of a named author but is, in substance, the output of a large language model lightly edited for tone. The named author may have provided a topic, perhaps a few bullet points of direction. The thinking, the structure, the argument, and very often the prose itself were produced by an AI. The author's contribution is their name and their willingness to claim it.

This is not, strictly speaking, new. Ghostwriting has existed for as long as publishing. Executives have always had speeches written for them. The consulting partner whose name appears on the thought-leadership paper rarely wrote more than the covering email. The profession has long accepted a degree of performance in its intellectual output.

But there is a difference of kind, not merely degree, in what is happening now — and the profession has been remarkably reluctant to examine it.

What Has Changed

The difference is threefold.

- **Volume.** A ghostwriter produces a considered piece over days or weeks. An AI agent produces one in minutes. The economics of production have shifted so dramatically that what was once an occasional practice — publishing under your name something you did not write — has become a continuous content strategy. Some practitioners are publishing daily, across multiple platforms, with a consistency and range that would be physically impossible without AI assistance. The audience may not yet have noticed, but the suspicion is growing.
- **Depth of detachment.** A ghostwriter who works with a client over months absorbs their thinking, their vocabulary, their positions. The output, while not authored by the named individual, is at least derived from their actual views. An AI agent has no

such relationship. It produces plausible-sounding content that could reflect the named author's views but just as easily might not. The author may not even have read the output carefully enough to know whether they agree with it.

- The erosion of trust. Thought leadership has always served a specific function in professional services and consulting: it signals that this person has thought deeply about this problem. A client reading an insightful paper on transformation governance makes an inference — this person understands our challenges, they have seen this before, they can help. When the paper was actually produced by an AI agent and the author's contribution was pressing "generate," that inference is fraudulent. The signal has been decoupled from the substance.

The Profession's Complicity

The uncomfortable truth is that the profession has created the conditions for this. The relentless pressure to "build a personal brand," to "demonstrate thought leadership," to "publish or perish" in a corporate context has turned intellectual contribution into a marketing activity. When the goal is visibility rather than insight, the incentive to use whatever tools produce the most content at the lowest cost is overwhelming.

The ethical question is not whether AI can produce convincing thought leadership. It plainly can. The question is whether a profession built on the promise of human judgement and experience can survive when its public intellectual output is neither human nor experiential.

What Practitioners Owe Their Audience

The answer is not to ban AI from the writing process. That ship has sailed, and the demand is too strong. Nor is it to require disclosure on every piece of content — though transparency about process would be healthier than the current pretence.

The answer, I believe, is simpler and harder: practitioners must only publish what they have genuinely thought. An AI agent can draft, structure, and refine — but the argument, the position, the insight must originate in the author's actual experience and actual reflection. The test is not "did you write every word?" but "would you defend every claim in a room full of people who know the subject as well as you do?"

The practitioners who meet this test — who use AI as a drafting tool but remain the source of the thinking — will, over time, be distinguishable from those who do not. The quality of the argument, the specificity of the examples, the willingness to take positions that are genuinely controversial rather than safely provocative — these are the markers that audiences will learn to look for.

The profession's credibility depends on it.