

AI-Assisted B2B Content: Give the Buyer a Reason to Care

Juan Beltran | Revised 2 October 2026

Interchangeable copy fails the reader before a style detector enters the conversation. Build B2B content around a useful decision, applicable evidence and accountable judgment.

Three posts with the same first line

I read three LinkedIn posts in a row that opened with “In today’s fast-paced world.” One came from a CEO, one from a new founder and one from a competitor. They shared the opening line, a three-bullet middle and an emoji closer.

The pattern did not tell me who wrote them. It did make the writing interchangeable. Nothing in the opening explained why I should give this particular person my attention.

That is the problem I would ask a B2B marketing leader to address. AI can help produce more material at lower effort. The reader still needs a reason to care: an observation that fits their situation, evidence they can examine, or a judgment that helps them make a decision.

“Delve” is a real word. An em dash can be useful. Removing both does not turn a generic claim into an interesting one. A page can pass every surface-style rule and still leave the buyer with nothing they could not find elsewhere.

Write for a decision the buyer recognises

“Our innovative platform helps forward-thinking teams unlock efficiency” makes the reader do the work. Which team? What becomes easier? Under which constraints? What evidence should a buyer believe?

A useful asset answers those questions at the level appropriate to its audience. A maintenance manager may need to know how a service handles an unresolved fault. Procurement may need a comparison of commitments. An executive may need to see which operating tradeoff the supplier understands.

Specificity is more than adding a company name and a number. An irrelevant number can decorate an empty claim. A named customer can lend credibility to a result obtained under conditions the current buyer does not share. Explain the connection and keep the scope beside the evidence.

Point of view matters as well. A supplier can describe the choice it recommends, the constraint that makes it difficult and the alternative it rejected. That shows judgment. “Industry leaders are embracing transformation” gives a reader neither a decision nor a defensible reason for one.

What the authorship evidence actually supports

Kirk and Givi's [The AI-authorship effect](#), published in the Journal of Business Research, reports seven preregistered experiments. When consumers believed emotional marketing communications were AI-written, positive word of mouth and loyalty were reduced. The effect was weaker for factual messages and when AI edited rather than authored the communication.

Those findings make authenticity worth examining, especially when a message asks the reader to trust a person's emotion or experience. They do not establish reliable detection from vocabulary, or a measured general loss of B2B deals. A buyer's belief about authorship, a style pattern and a commercial outcome are different things.

My operating conclusion is to put the real source of an emotional or experiential claim in charge of its substance. A founder story needs the founder's actual experience. A customer story needs the customer's evidence and permission. AI can assist with structure or editing; it cannot provide the experience being claimed.

For factual work, the central check is the claim and its applicability. Heavy AI assistance can still leave a useful, accurate comparison. A human-written comparison can still mislead. The reviewer's job is to establish what the reader will understand and what the company can defend.

Give the model the brief, and give a person the judgment

Start with the reader and the decision. Identify the approved sources, relevant constraints and what would be useful to understand after reading. Give that brief to the model if you use it to draft. A brief the model never sees cannot guide its first pass.

For a factual asset or routine outreach draft, ask AI to work from the approved evidence and keep gaps visible. The responsible writer then chooses the argument, verifies the claims and rewrites what is irrelevant or merely familiar. That review should improve the reader's understanding, not just obscure how the draft was produced.

For a personal or emotional communication, ask the accountable person to supply the original substance and judgment first. AI may help them organise or edit it. Do not ask the model to manufacture a first-person anecdote and then attach someone's name.

Both routes finish with the same responsibility: an identified person approves the meaning, evidence and publication. The byline tells the reader who stands behind the piece; it does not substitute for the review.

Read the result aloud where a spoken tone is appropriate. If a sentence takes two attempts to understand, revise it. Varying sentence length can help clarity, but it is not a test of who wrote the text.

Before and after

Here is an illustrative cold email:

Hi [First Name], I hope this message finds you well. I wanted to reach out because I noticed you're at [Company] and thought you might be interested in how our innovative platform helps forward-thinking teams unlock new efficiencies. We work with leading enterprises to streamline operations and drive measurable results. Would you be open to a quick discovery call to explore potential synergies?

The weakness is not a forbidden word. The recipient has no particular reason to receive it. The message asks for time before showing that the sender understands anything about the account.

A more useful worked example could read:

Maria, I saw three SRE roles advertised for your São Paulo team. Are you adding reliability capacity in the region? We helped [named customer] reduce [relevant problem] under [comparable conditions], with [verified result]. If that is the work behind the hiring, I can send the short implementation note. Would it be useful?

The roles and recipient are part of the example. Replace the placeholders with evidence you may actually share, or leave out the result. The question tests an inference instead of announcing that a job posting proves a board priority.

The second version gives the recipient a business signal, a bounded comparison and a useful next step. It still needs editing for the real account. It is not evidence of a higher reply rate; that would require an actual comparison.

Review usefulness before policing style

I would put a small sample of recent assets in front of the writer and a commercial colleague. Ask what decision each helps, which claim can be checked and what the buyer learns that a competitor's generic page would not supply.

When those answers are weak, revise the substance first. If the reader, decision and evidence are clear but the prose is inflated, cut abstractions, repetitive openings and stock conclusions. Surface habits are good revision prompts. They are poor automatic publication gates.

Keep existing performance in view. A stylistic score alone is no reason to unpublish or noindex a useful page. Correct inaccurate claims promptly; decide on broader revision or removal from the asset's purpose, quality and observed use.

For a larger content programme, look at reply quality, qualified enquiries, buyer feedback and which material appears in real sales conversations. Falling replies or branded search deserve investigation, but neither proves that AI authorship caused the change. Audience, offer, distribution and market conditions can move those measures too.

Where feasible, compare content with similar audiences and purposes, change a clearly identified feature, and inspect the response. Interviews can reveal what the reader found useful or unconvincing even when attribution is difficult. Treat the result as evidence for that context, not a universal rule for every channel.

The standard I would keep is simple: could a named sales leader stand behind this asset in a conversation with an important buyer? They should be able to explain why it matters, where its claims come from and which judgment the company is making.

That gives the writer a better job than hiding the word "delve." Give the buyer something worth reading.

The downloadable guide develops the review method and includes a worksheet for applying it to your own assets.

Reading worksheet

Use these questions with the essay. Record your own evidence and the next decision; the examples in the essay do not supply a result for your organization.

1. What can a buyer learn here that is specific to our offer?

2. Can we show the source or scope of each consequential claim?

3. Which useful detail was lost when we polished the draft?

4. Who will check the final promise against what we can deliver?

Sources and scope

Editorial argument informed by the Kirk-Givi publisher abstract, with a firsthand reading anecdote and explicitly illustrative outreach. The review approach is a proposed practice, not a validated detection instrument or conversion model.

The cited experiments concern consumer responses and beliefs about emotional authorship. They do not establish reliable individual detection, universal B2B conversion effects or channel bans. The worked email has no measured performance result.

Colleen P. Kirk and Julian Givi, Journal of Business Research. [The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications](#). The publisher's abstract and indexed preview report seven preregistered experiments in which belief that emotional marketing communications were AI-authored reduced positive word of mouth and loyalty; the effect weakened for factual messages and AI editing. This does not establish reliable vocabulary-based detection or measured general B2B deal loss. The correct article number is 114984; the abstract and publisher metadata, rather than the full paper, were consulted.

[Read the current article](#)