

The Brutal Page 19—And What It Means for AI



Thomas Wyck: Ein Gelehrter in seinem Arbeitszimmer

Anyone who thinks about writing and publishing a book confronts the Page 19 Dilemma during their research. Following the insight of an author and rhetoric expert, Wladislaw Jachtchenko, e-book streaming platforms report that the average reader of a nonfiction book stops on page 19.

This is not survey data—but mathematically precise, hard usage data, captured because publishers are paid per page read.

Here's the part that makes it worse: that average only counts readers who opened the book at all. Anyone who bought it kept it in the pile for the next reading or just to fulfill Umberto Eco's theorem that you know someone's personality by the books he purchased but never read is not even entering that data.

Now, I can't stop seeing the parallel to how we talk about AI maturity. The market loves proxy metrics: token spend, number of PoCs launched, models deployed, activity on Hugging Face, robustness of the governance structure, etc. That's equal to the book purchase. It tells you someone opened the cover. Nothing more.

The real page-19 question is different: would the workflow visibly break if you silently switched the AI component off for a week?

There can be many answers, from "Actually, nobody noticed it" to "I don't even know how to handle it without AI support." So, it could be consumption without reading capacity, without penetration. And just like with books, the drop-off is rarely about the topic. Need to earn trust, demanding intellectual investment, involvement, and finally, and also therefore, something at risk.

A lot of AI PoCs and projects repeat the mistake that none of the well-written nonfiction books do—they open with definitions ("look what the model can do") instead of having a robust understanding of what problem to solve and the ability to involve the audience beyond the first attack of curiosity or FOMO, which means a real, cost-bearing pain point that tests whether the system can be trusted under consequence.

Three questions I now ask clients before we even talk about "maturity":

- 1. Continuity and Retention without an immediate mandate — does it keep running once the attention of the sponsor and the pilot budget disappear?**
- 2. Deployment with real consequence — does the output sit somewhere a mistake actually costs something, or is it still in the sandbox?**
- 3. The silent-shutoff test — would anyone actually miss it, or does the team just need another Project code to report hours?**

Activity can be bought. Trust can't. That's why "how much is being used" is never the right question — the real one is whether the tool would have survived page 19 if nobody were watching anymore.

Where does your AI deployment actually stand — still at the point of purchase, or already past page 20