


BEYOND THE PRINCIPLES POSTER

Ethics Boards Without Teeth — What the Textbooks Leave Out

A Practitioner's Perspective on the Gap Between AI Ethics Governance and Organisational Reality

Giovanni Leonardi

July 2019



Ethics Boards Without Teeth — What the Textbooks Leave Out

Leadership · No. 65

Originally written July 2019 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“An ethics board that cannot slow a product launch is not a governance mechanism — it is a press release with a meeting schedule.”

1. The Pattern

Over the past eighteen months, I have watched a growing number of organisations establish ethics boards, ethics committees, or responsible AI councils. The pattern is consistent enough to describe with some confidence. A senior leader — usually prompted by media scrutiny, a competitor's misstep, or an internal incident involving algorithmic decision-making — announces that the organisation will take a principled approach to the ethical use of data and artificial intelligence. A set of principles is drafted, typically by a small working group drawn from legal, compliance, data science, and sometimes an external academic. The principles are published, often externally. A governance body is established to oversee their application. And then, with remarkable consistency, nothing of consequence happens.

This is not because the people involved lack sincerity. Many of the practitioners I have observed working on these initiatives are deeply committed to the questions they are grappling with. The problem is structural. The ethics board is established without the one thing it needs to function: the authority to make decisions that cost the organisation something.

2. Why They Lack Authority

The structural deficit is not accidental. It reflects a genuine tension in how organisations think about governance. Governance mechanisms that can delay revenue, block product launches, or impose additional cost on delivery are established only when there is a compelling external force — regulatory penalty, legal liability, or reputational risk severe enough to command board attention. Financial crime compliance has teeth because the penalties are existential. Health and safety governance has teeth because the liabilities are personal.

AI ethics, as of mid-2019, has none of these forcing functions at the level most organisations operate. There is no AI-specific regulation with material penalties in most jurisdictions. The reputational risks, while real, are diffuse and hard to quantify. The legal

liabilities are largely untested. An executive who establishes an ethics board with genuine veto power is, from a purely organisational perspective, creating a constraint on commercial agility without a corresponding external mandate. It is an act of anticipatory governance, and anticipatory governance is rare in organisations under quarterly performance pressure.

The result is that ethics boards are established with advisory mandates. They can review, recommend, and raise concerns. They cannot approve, reject, or require. The distinction is everything.

3. The Advisory Trap

An advisory ethics board creates an illusion of governance that may be worse than no governance at all. It allows the organisation to point to a structure and a process when challenged, while preserving the ability to proceed with any decision the commercial leadership considers important. The board reviews a proposal, raises concerns, and issues recommendations. The delivery team notes the recommendations, addresses the ones that are convenient, and proceeds. If the concerns prove justified and an incident occurs, the organisation can point to the ethics board's involvement as evidence of due diligence. If the concerns prove unfounded, the delivery team congratulates itself on not being slowed down.

The practitioners who serve on these boards understand the dynamic acutely. In my experience, the most capable and principled members are the first to disengage, because they recognise that their participation lends legitimacy to a process that does not take their input seriously. They are replaced by less experienced or less assertive members, and the board's effectiveness degrades further. Within twelve to eighteen months, the board is either dormant or performing a purely ceremonial function — reviewing cases after the fact, producing annual reports that no one reads, and meeting quarterly to discuss principles that have no operational expression.

An ethics board that cannot slow a product launch is not a governance mechanism — it is a press release with a meeting schedule.

4. What the Textbooks Prescribe

The emerging literature on AI ethics governance — and there is a great deal of it — tends to focus on the content of ethical principles and the composition of ethics boards. Which principles should be adopted? How should fairness be defined? Should the board include external members? How should it relate to existing risk and compliance functions?

These are legitimate questions, but they miss the structural issue. The composition of the board and the quality of its principles are irrelevant if the board has no mechanism to influence decisions at the point where they are made. The textbooks describe what an ethics board should think about but are largely silent on what it should be able to do.

The gap is understandable. Prescribing authority structures is context-dependent and politically sensitive in a way that prescribing principles is not. But it is the authority structure, not the principles, that determines whether governance is real.

5. What Would Actually Work

The organisations where I have seen ethics governance function — and they are few — share several characteristics that are worth noting.

First, the ethics review is embedded in an existing decision gate that already has authority. Rather than creating a parallel governance structure, the ethical review is incorporated into the stage-gate process for product development, the risk assessment process for new business initiatives, or the architectural review process for technology changes. This means the ethics review inherits the authority of the host process: a product cannot proceed through the gate without an ethics assessment, just as it cannot proceed without a security assessment or a regulatory impact analysis.

Second, the review criteria are specific and operational, not abstract. Rather than asking whether a proposal is fair or transparent — questions that are important but difficult to adjudicate in a governance meeting — the criteria address concrete, assessable questions. Has a bias assessment been conducted on the training data? Has the decision boundary been tested against protected characteristics? Is there a documented process for a human to review and override an algorithmic decision? Can the logic of the decision be explained to the affected individual in terms they can understand? These are questions that can be answered yes or no, and a no can be a genuine gate.

Third, someone with sufficient seniority has accepted personal accountability for the ethical governance of AI and data use — not in the sense of chairing a committee, but in the sense of owning the risk. When a specific individual knows that an AI ethics failure will be attributed to their governance, their attention to the effectiveness of that governance sharpens considerably.

6. The Leadership Question

The absence of teeth in ethics boards is, ultimately, a leadership failure. It is a failure to accept that ethical governance of data and AI will, at some point, require the organisation to forgo revenue, delay a launch, or spend money it would prefer not to spend. Leaders who establish ethics boards without granting this authority are either hoping the question will never arise — which is naive — or have decided in advance that commercial considerations will always prevail — which makes the board dishonest.

The honest position is either to establish genuine governance with genuine authority and accept the cost, or to acknowledge that the organisation is not yet willing to constrain itself and stop pretending otherwise. Both positions are defensible. The indefensible position — and the one I see most often — is to establish governance that looks real from the outside but is designed, from the inside, never to impede anything that matters.

Regulation will eventually close this gap. The direction of travel in the European Union, and increasingly elsewhere, is toward mandatory requirements for algorithmic transparency, bias assessment, and human oversight. When that regulation arrives with enforceable penalties, ethics governance will acquire the teeth that organisational leadership has been unwilling to provide voluntarily. The organisations that will be best prepared are those that did not wait to be compelled.