


BEYOND THE CHARTER

Ethics Boards Without Teeth — Why the Governance of AI Ethics Became an Exercise in Institutional Performance

A Practitioner's Perspective on the Gap Between Ethical Intent and Organisational Authority in 2018–2019

Giovanni Leonardi

August 2019



Ethics Boards Without Teeth — Why the Governance of AI Ethics Became an Exercise in Institutional Performance

Leadership · No. 79

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

"The ethics board met quarterly, issued recommendations, and changed nothing."

The Board That Could Only Advise

Across financial services, telecommunications, and the technology sector, the past eighteen months have seen a visible proliferation of AI ethics boards, ethical review committees, and responsible AI advisory panels. The announcements are polished. The charters are thoughtful. The membership lists are impressive — a mix of academics, senior leaders, external advisors, and occasionally a philosopher or two.

And yet, in my observation, almost none of these bodies have any meaningful authority over the decisions they are nominally created to govern.

The pattern is consistent enough to be structural rather than incidental. An organisation recognises that its use of data and algorithmic decision-making raises ethical questions it has not previously confronted. It establishes a board or committee to provide oversight. It drafts terms of reference that speak of review, guidance, and recommendation. And there it stops. The board can advise; it cannot direct. It can raise concerns; it cannot block deployment. It can publish principles; it cannot enforce them.

This is not ethics governance. It is ethics theatre.

Why This Keeps Happening

The structural reason is straightforward: ethics boards are typically established by, and report to, the same leadership teams that are driving the commercial adoption of the technologies in question. The incentive to create a body that might slow down or halt a revenue-generating initiative is precisely zero. What leaders want is the appearance of ethical oversight — something to point to when regulators, journalists, or customers ask whether the organisation has considered the ethical implications of its work. What they do not want is a body with the authority to say no.

The result is a governance structure that is advisory by design and toothless by intent. The ethics board reviews cases after the fact, or reviews them in principle without access

to the specific implementation details that would make its review meaningful. It issues recommendations that land on the desks of delivery teams who have neither the mandate nor the incentive to act on them. It meets quarterly in a cadence disconnected from the daily decisions that actually shape how algorithms are built, trained, and deployed.

What Would Be Different If It Were Real

Genuine ethics governance would look quite different from what most organisations have built. It would require, at minimum, three things that are conspicuously absent from the current model.

First, authority to intervene before deployment, not after. An ethics board that reviews a model after it has been deployed is performing an autopsy, not governance. Meaningful oversight requires a gate — a point in the delivery process at which the ethics review is a prerequisite for production, not a retrospective commentary on it.

Second, access to technical specifics, not just strategic summaries. Most ethics boards receive presentations: high-level descriptions of what a model does, couched in reassuring language about fairness and transparency. They do not see the training data, the feature selection rationale, the performance across demographic segments, or the failure modes. Without this detail, ethical review is necessarily superficial.

Third, independence from the commercial function. As long as the ethics board reports to the same leadership that sponsors the commercial deployment, its independence is compromised. The most credible models I have seen — and they are rare — position the ethics function alongside risk or compliance, with a reporting line that does not pass through the business unit whose work it oversees.

The Risk of Getting Comfortable

The danger of the current approach is not that ethics boards do nothing. It is that they create a false sense of assurance. Leaders believe they have addressed the ethics question because they have established a committee. Regulators see a governance structure and assume it is functioning. And the organisation settles into a comfortable arrangement in which ethical review exists in name but exerts no meaningful constraint on behaviour.

When the first serious failure arrives — a biased lending model, a discriminatory hiring algorithm, a surveillance application that violates the spirit if not the letter of the law — the ethics board will be asked what it did. And the honest answer, in most cases, will be: it met, it discussed, and it recommended. Whether anyone was required to listen is a question the charter was carefully drafted not to answer.