


BEYOND THE AUDIT TRAIL

The Three-Lines Model in Programme Delivery — What Defence Taught Us and What We Got Wrong

A Practitioner's Perspective on Adapting Assurance Models for Complex Programme Environments



Giovanni Leonardi

December 2007

The Three-Lines Model in Programme Delivery – What Defence Taught Us and What We Got Wrong

Programme · No. 118

Originally written December 2007 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The three-lines model works brilliantly on paper and fails consistently in practice – not because it is wrong, but because it assumes a clarity of boundary that complex programmes do not possess.”

1. A Model Borrowed, Not Built

The three-lines-of-defence model arrived in programme management the way most governance ideas arrive – borrowed from a neighbouring discipline where it had proved its worth, and applied with more enthusiasm than adaptation. Its origins in financial services regulation are well documented: the first line owns and manages risk, the second line oversees and sets the framework, the third line provides independent assurance. Clean, logical, and in its original context, genuinely effective.

The appeal for programme delivery was obvious. Large programmes – particularly in defence, infrastructure, and government – were struggling with a persistent question: who is actually responsible for ensuring this programme is on track, and who is responsible for checking whether the people who say it is on track are telling the truth? The three-lines model appeared to answer this definitively. The programme team delivers. The programme office oversees. Internal audit or an independent assurance function verifies.

In my experience, the model has been adopted widely and adapted rarely. And the gap between adoption and adaptation is where most of the problems live.

2. Where the Model Meets Reality

The difficulty is not with the principle of separation. Separation of delivery from oversight from independent assurance is sound governance design. The difficulty is with the assumption that these separations can be maintained cleanly in the operating reality of a complex programme.

Consider what happens in practice. The programme team – the first line – produces a status report. It shows the programme as amber, with mitigations in place. The programme management office – the second line – reviews this report. But the PMO is typically embedded within the programme structure, staffed by people who report to the programme director, and dependent on the programme team for their information. The "oversight" they provide is, in reality, a review of information produced by the people they are supposed to be overseeing, conducted by people who share the same reporting line and the same incentives.

The third line — independent assurance — faces a different problem. It is genuinely independent, but it is also genuinely distant. An assurance review conducted quarterly or bi-annually can identify systemic issues, but it cannot provide the real-time challenge that complex programmes need. By the time the third line identifies a problem, the programme has typically been living with it for months.

The model assumes that each line can do its job without depending on the others for its information, its access, or its authority. In programme delivery, all three dependencies are the norm.

3. The Boundary Problem

The deepest structural issue is what I think of as the boundary problem. The three-lines model requires clear boundaries between each line — between delivery and oversight, between oversight and assurance. In financial services, where the model originated, these boundaries are reinforced by regulation, by separate reporting lines, and by professional standards that make the independence of each line a matter of institutional design.

In programme delivery, no such reinforcement exists. The boundaries are organisational choices, and they are subject to the same pressures as every other organisational choice in a programme under pressure. When a programme is running late, the PMO is pulled into delivery support. When an assurance review raises uncomfortable findings, the programme director's first instinct is to challenge the methodology rather than address the findings. When the senior responsible owner needs a single coherent story for the board, the three lines are quietly collapsed into one narrative.

This is not cynicism. It is the natural consequence of applying a model that requires institutional separation to an environment where institutional integration is the dominant force. Programmes, by their nature, pull people together around a shared objective. The three-lines model requires them to remain apart.

4. What Actually Works

The programmes I have seen navigate this most effectively have not abandoned the three-lines principle, but they have been honest about its limitations and pragmatic about its application.

The most important adaptation is making the second line genuinely independent of the programme director. This typically means the PMO or programme assurance function reports to a different senior leader — the portfolio director, the chief operating officer, or a dedicated assurance director. The reporting line matters enormously, because it determines whose questions the second line is trying to answer. If the second line reports to the programme director, it answers the programme director's questions. If it reports to someone with a broader mandate, it asks different questions entirely.

The second adaptation is investing in the quality of the second line. In too many programmes, the PMO is a scheduling and reporting function — competent administrators who can produce a Gantt chart and consolidate a status report, but who lack the seniority and the technical depth to genuinely challenge the programme team. An effective second

line requires people who have run programmes themselves, who can read between the lines of a status report, and who have the standing to ask difficult questions without being dismissed.

The third adaptation is increasing the cadence and reducing the formality of third-line assurance. The traditional model of a formal assurance review every six months produces a document that is comprehensive, carefully worded, and almost always too late to be useful. The programmes that get value from independent assurance tend to use a lighter, more frequent model – shorter reviews, more often, focused on specific risk areas rather than attempting a comprehensive assessment each time.

5. The Lesson Beneath the Lesson

The deeper point is not about the three-lines model specifically. It is about the habit of importing governance models from one context to another without examining whether the conditions that made them effective in the original context exist in the new one.

The three-lines model works in financial services because the regulatory environment enforces the separations the model requires. It works in defence procurement – to the extent that it does – because the scale and formality of defence organisations provide the institutional mass to maintain separate lines even under pressure. In programme delivery more broadly, these conditions rarely hold. The result is a model that provides the appearance of robust governance while the reality is something considerably less structured.

This is not an argument for abandoning the model. The principle of separating delivery from oversight from assurance remains sound. But it is an argument for being honest about the gap between the model as drawn and the model as lived – and for investing the organisational effort required to close that gap, rather than assuming that adopting the model is the same as implementing it.

“The three-lines model works brilliantly on paper and fails consistently in practice – not because it is wrong, but because it assumes a clarity of boundary that complex programmes do not possess.”

The organisations that have made this work have done so by treating the three-lines model as a starting point for governance design, not as a finished product. They have asked: given our programme structure, our reporting lines, our culture, and our risk appetite, how do we create genuine separation where it matters most? The answer is always specific to the context. And it is always harder than the model suggests.