Thrive Or Die - UK Telecoms

UK Telecoms Edition — Two Operators, Two Governance Philosophies, One Live Test

A Charterhouse Consultants white paper by Bridget Convey, July 2025

Executive Summary

BT and Vodafone are currently running the closest thing to a live, head-to-head experiment in this entire paper series. BT has attributed up to 55,000 planned job cuts partly to AI, with chief executive Allison Kirkby stating the current plan 'did not reflect the full potential of AI.' Vodafone's chief technology officer has explicitly ruled out following BT's lead, describing the company's approach instead as a skills rebalancing. Both are UK-headquartered, both operate in the same market, and both are making a public, falsifiable bet on how much AI capability should be allowed to restructure the organisation, and how fast.

The Premise: The Gap Is The Danger, Not The AI

In October 2023, a group including Turing Award winners Yoshua Bengio and Geoffrey Hinton, alongside Stuart Russell, Daniel Kahneman and more than twenty other senior AI researchers, published a consensus paper in the journal Science on the risks of advanced AI. Their argument was precise. AI capability and autonomy are increasing rapidly. Organisations are shifting toward systems that can act and pursue goals with reduced human oversight. Governance is not keeping pace with that shift. The danger they describe does not come from AI itself. It comes from the widening gap between what a system can now do on its own and how closely anyone is still watching it.

That paper is about safety and governance at a societal scale. It does not comment on organisational strategy, and I want to be precise about that rather than overstate its relevance. What it gives this paper is not an endorsement, it gives us the mechanism. Capability grows. Autonomy grows. Oversight does not grow at the same rate. In UK telecoms specifically, the gap is visible in a warning from inside BT's own wholesale division: agentic AI systems deployed across horizontal operations can each develop a different interpretation of basic business terms, such as what counts as an 'unlimited plan' or 'roaming', and that interpretation silently drifts further apart every time the underlying product catalogue changes, unless something is actively governing consistency across agents.

This is the basis for the claim in the title. Thrive or die is not a slogan layered on top of the argument. It is the plain consequence of the mechanism. An organisation that builds governance at the same pace it builds AI capability closes the gap and survives it. One that lets capability outrun governance does not get a soft landing.

Thrive Or Die: The Pattern In Organisations That Did Not Close The Gap

Kodak invented the first digital camera in 1975 and its own internal research predicted by 1981 that digital would eventually replace film. What it did not close was the governance gap between that capability and a business model still built entirely around film margins. The company filed for bankruptcy in 2012, thirty seven years later. Blockbuster was offered the chance to buy Netflix outright for fifty million dollars in 2000, declined, and filed for bankruptcy in 2010, having eliminated close to eighty four thousand jobs in the collapse. Both prove the same point. The danger was never a lack of capability. It was a governance gap left open until it was unmanageable.

Kodak and Blockbuster failed by protecting an old model for too long. UK telecoms shows a live split between two operators making opposite bets on how fast to let AI capability change the organisation. BT has committed to cutting up to 55,000 roles, including contractors, by 2030, and CEO Allison Kirkby told the Financial Times that this figure did not reflect AI's full potential, meaning the company could become smaller still. Vodafone's chief technology officer Scott Petty has taken the opposite public position, stating the company does not plan to follow BT's lead in AI-driven job replacement, framing the impact instead as a need to rebalance skills toward software roles. Independent analysis of layoffs across the wider telecom sector, including BT, Vodafone, Telefónica and Telecom Italia, found no reliable link between headcount reduction and margin improvement, suggesting AI is, in some cases, being used as the public explanation for cuts that are structurally driven by debt reduction, network consolidation and fibre rollout completion rather than by AI capability itself. If that is true in some cases, it is itself a governance gap, a mismatch between what is actually driving the decision and what is being said about it publicly.

The Evidence In UK Telecoms: A Public Divergence, And A Warning From Inside The Machine

Two pieces of evidence show this sector's version of the gap especially clearly, one at the level of two companies' public strategy, the other inside the technical detail of how agentic AI actually behaves once deployed.

BT And Vodafone: Two Public Bets On The Same Technology

BT's EE division already runs Aimee, a virtual assistant handling up to 60,000 customer conversations a week, and BT has stated that AI is being used across sales and support operations at both BT and EE. Kirkby's comment that the 55,000 figure 'did not reflect the full potential of AI' signals a company willing to let capability drive organisational size further and faster than its own current public commitments. Vodafone, facing the same underlying technology, has drawn a different public line, ruling out following BT specifically while still acknowledging AI will require a 'rebalancing' of skills, notably toward the roughly 7,000 software roles it has separately committed to hiring. Two operators, comparable scale, same national market, opposite public governance stance on how much authority AI capability should have over headcount decisions.

The Agentic Drift Warning: The Gap Mechanism Inside The Technology Itself

A director at BT Wholesale has warned publicly that 2026 will bring an 'agentic AI reckoning' for telecoms operators running horizontal AI platforms across their operations. The specific risk described is technical and precise: each AI agent deployed across a large, interdependent network of systems develops its own interpretation of core business terms through prompting alone, and that interpretation can work acceptably in a pilot, then drift further from the others every time the underlying product catalogue changes, unless something is actively monitoring and correcting for consistency across agents. This is the capability-outrunning-governance mechanism described in the Science consensus paper, translated into a specific, named technical failure mode inside live UK telecoms infrastructure, not a hypothetical one.

UK Telecoms In Three Horizons: The Gap Closes Or It Does Not?

If the gap between capability and governance is the real variable, then each horizon below has two outcomes, not one. These are deliberately stated as firm positions, not hedged possibilities, because a board or leadership conversation needs something to agree or disagree with.

Two Years: 2026 To 2028

What happens everywhere: agentic AI moves from customer service pilots, like EE's Aimee, into core network and business operations across major UK operators.

- Thrive path. Operators build active, ongoing monitoring for agent consistency across their AI deployments, treating the drift risk BT Wholesale identified as a standing operational requirement, not a one-off configuration task.
- Die path. Operators scale agentic AI across operations without a mechanism for catching drift between agents, and the first visible failure, an incorrectly applied roaming charge or an inconsistent plan interpretation at scale, becomes a regulatory and customer trust event.

Five Years: 2028 To 2031

What happens everywhere: the BT and Vodafone divergence resolves one way or the other, as results from each company's approach become visible in margins, customer satisfaction and market share.

- Thrive path. Whichever strategy proves correct, aggressive AI-led restructuring or measured skills rebalancing, becomes the template the rest of the sector converges on, based on demonstrated results rather than public positioning.
- Die path. Neither company's public narrative fully explains its actual results, since the sector-wide data already shows a weak link between headcount cuts and margin improvement, and telecoms boards continue attributing structural decisions to AI without the underlying evidence to support the claim either way.

Ten Years: 2031 To 2036

What happens everywhere: 'fully autonomous networks,' the destination Nokia's European CTO has predicted the sector is moving toward from 2026 onward, become operational reality rather than a stated ambition.

- Thrive path. UK operators that built governance for agent consistency and organisational change together, rather than treating headcount reduction and AI deployment as the same decision, run genuinely autonomous networks with a defensible audit trail behind every automated action.
- Die path. Operators that used AI primarily as a public justification for structurally-driven cuts, without building matching technical governance, run autonomous networks nobody can fully explain when something goes wrong, the network engineering equivalent of the agentic drift problem, at national infrastructure scale.

What I Am Recommending To UK Telecoms Boards

- Separate, explicitly and publicly, the portion of any restructuring driven by AI capability from the portion driven by debt reduction, network consolidation or fibre rollout completion, since sector-wide data suggests these are currently being blurred together in public messaging.
- Build active monitoring for interpretation drift across every agentic AI system operating on live business rules, following the specific warning from BT's own wholesale division, before scaling agentic deployment further.
- Decide deliberately, at board level, whether your organisation's public position is closer to BT's or Vodafone's, and ensure the internal evidence supports whichever public claim is being made, since both companies are now accountable to that claim in a way that will be visible in results within a few years.

Conclusion

UK telecoms is currently running the clearest live test in this entire series of what happens when two comparable organisations make opposite public bets on the same technology. BT is betting that AI capability should be allowed to shrink the organisation further than currently planned. Vodafone is betting the opposite, that the change is about skills, not headcount. Underneath both bets sits a specific, named technical warning from inside BT's own operations, that agentic AI systems drift apart in their interpretation of business rules unless something is actively governing that consistency. Kodak and Blockbuster took decades to prove what an unmanaged gap costs. BT and Vodafone will prove it, one way or the other, within the next few years, in full public view.

Sources: RSInc, Data Center Dynamics, Light Reading, Mobile World Live, Mobile Europe, Telecoms.com, TelcoTitans, Business Insider via AOL. All figures current as of mid-2026, verified at time of writing. This is a first draft prepared for internal discussion.