Thrive Or Die – Insurance

UK Insurance Edition — Why The Gap Between AI Capability And AI Governance Decides Which Insurers Survive

A Charterhouse Consultants white paper by Bridget Convey, March 2026

Executive Summary

This paper makes one claim, applied to UK insurance specifically. The gap between how fast AI capability grows inside an insurer and how fast that insurer's governance of it grows is not a compliance issue. It is the variable that decides which UK insurers thrive over the next decade and which do not survive it. That is not a metaphor, it is already operational in the UK market, visible in Aviva's fraud detection figures and in Lemonade's claims model, and it follows the same mechanism senior AI safety researchers have identified at a societal scale.

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. Companies 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 corporate strategy, and I want to be precise about that rather than overstate its relevance. What it gives this paper is not a business endorsement, it gives us the mechanism. Capability grows. Autonomy grows. Oversight does not grow at the same rate. The gap between those lines is where damage happens, whichever scale you are looking at, and in UK insurance specifically it is now visible on both sides of the same claim, in the fraud detection tools insurers deploy and in the AI-generated evidence fraudsters increasingly use against them.

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. A firm that builds governance at the same pace it builds AI capability closes the gap and survives it. A firm that lets capability outrun governance does not get a soft landing. History already shows what that failure looks like when the underlying capability was digital and connected rather than AI and autonomous. It looks like Kodak. It looks like Blockbuster.

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

Kodak and Blockbuster are usually told as stories about missing a trend. Looked at through the gap mechanism, they are something more specific, cases where the company had the capability internally, had time to act on it, and still let the gap between what it could do and what it was willing to change grow too wide before closing it.

Kodak invented the first digital camera in 1975 and its own internal research predicted by 1981 that digital would eventually replace film. Kodak had the capability decades before its collapse. What it did not close was the governance gap between that capability and a business model still built entirely around film margins. Every digital product Kodak launched was positioned as a companion to film rather than its replacement, because closing that gap meant threatening the profitable core. The company filed for bankruptcy in 2012, thirty seven years after it had already built the technology that ended it.

Blockbuster was offered the chance to buy Netflix outright for fifty million dollars in the year 2000, when Blockbuster held over six billion dollars in annual revenue and nine thousand stores. Its leadership saw a small, unprofitable company and declined. What they failed to see was the same gap, a digital delivery capability accumulating quietly while governance and decision making stayed anchored to the physical store model. Blockbuster filed for bankruptcy in 2010, having eliminated close to eighty four thousand jobs in the collapse.

Both companies prove the same point this paper makes about AI. The danger was never that they lacked the capability. The danger was a governance gap that let a known, capable threat sit unmanaged for years until it was no longer manageable. That is the exact mechanism now playing out inside UK claims handling, at a faster tempo, described in the evidence below.

The Evidence In UK Insurance: Firms Already Living This Pattern

Two UK-facing examples show both sides of the gap mechanism in insurance, one showing what happens when capability accelerates on both sides of a transaction at once, the other showing what a governed AI core looks like when it is built in from day one rather than bolted on.

Aviva And Direct Line: The Fraud Arms Race As Proof Of The Gap

In 2025, following its acquisition of Direct Line, Aviva detected more than 18,400 suspect claims worth £233 million across its combined brands, equivalent to £638,000 of fraud stopped every day. The same data reveals the mechanism this paper describes, in miniature and in real time. Fraudsters are now using AI-generated images and manipulated documents to fabricate accident scenes and inflate claims, and insurers are responding with their own AI-driven detection, overseen by human review, to keep pace. This is the capability-versus-governance gap playing out on both sides of a single transaction. Aviva has also worked with McKinsey to rebuild its claims journey around AI end to end, treating the claim itself, not the policy document, as the primary customer relationship.

Lemonade: What A Governed AI Core Looks Like From Day One

Lemonade, which operates in the UK alongside the US, Germany, France and the Netherlands, was built AI-native from its first day rather than having AI added to a legacy claims process. Its AI system, AI Jim, holds the world record for the fastest insurance claim ever settled, at two seconds, and the company now resolves close to half of all claims through AI without human involvement, while more complex claims are escalated to human staff. That combination, instant resolution for simple claims and mandatory human escalation for complex ones, is the gap mechanism managed correctly, not a slogan. Lemonade did not automate the customer-facing shell first and figure out governance later. It built the escalation rule into the AI core from day one.

UK Insurance 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 conversation needs something to agree or disagree with.

Two Years: 2026 To 2028

What happens everywhere: AI-driven fraud detection and claims triage become standard across UK motor and home insurance.

- Thrive path. Insurers embed hard escalation rules alongside fraud AI, following the Aviva model, catching more fraud without wrongly rejecting genuine claims.
- Die path. Insurers deploy fraud AI or claims automation without matching investment in escalation governance, and either lose the fraud arms race to AI-generated evidence, or damage customer trust through wrongful auto-rejections.

Five Years: 2028 To 2031

What happens everywhere: claims settlement speed becomes a genuine, provable point of competitive differentiation rather than a marketing line.

- Thrive path. Insurers with a governed AI claims core, in the Lemonade pattern, use verified settlement speed as a real acquisition advantage, since it can be demonstrated rather than merely claimed.
- Die path. Insurers still running claims on largely manual review with AI bolted onto the front end fall behind on cost ratio and customer experience simultaneously, and become acquisition or reinsurance capacity targets for firms with a governed core already built.

Ten Years: 2031 To 2036

What happens everywhere: the underwriting licence becomes secondary to the data and decisioning infrastructure sitting underneath it, following the pattern already visible at Ping An.

- Thrive path. The insurer's product becomes largely indistinguishable from its AI decisioning engine, and the trust built into that engine, not the policy wording, is what customers are actually buying.
- Die path. UK insurers that never built a proprietary, governed AI core become underwriting capacity behind broker, MGA or Big Tech distribution brands, a role many smaller UK insurers already occupy today for reasons unrelated to AI, and now find themselves unable to exit.

What I Am Recommending To UK Insurance Boards

- Map every point where AI already touches a claims decision today, and identify which of those points has a genuine human escalation rule, versus which has become fully automated without anyone deliberately deciding that.
- Treat the fraud arms race as a permanent condition rather than a one-off project, since Aviva's own figures show fraud methods evolving as fast as detection does.
- Before marketing claims speed as a customer benefit, be able to evidence the governance behind it, not just the speed, since speed without governance is a Klarna-shaped risk waiting to happen inside a claims department rather than a customer service queue.

Conclusion

Insurance runs on trust in a promise that is only tested once, at the moment of loss. The capability-outrunning-oversight mechanism identified by Bengio and Hinton is now visible on both sides of every UK motor and home claim, insurers are automating faster, and so are the fraudsters targeting them. Kodak and Blockbuster show what happens when a knowable gap is left open for years. Aviva and Lemonade show, in different ways, that closing it is not hypothetical, it is already operational UK insurance practice in 2026. The only question left for an insurance board is which side of that gap their own claims process currently sits on.

Sources: Aviva plc newsroom, Insurance Business UK, Insurance Times, McKinsey & Company, Lemonade Inc. press materials, Reinsurance News, InsurTech Insights. All figures current as of mid-2026, verified at time of writing. This is a first draft prepared for internal board discussion.