

FINANCIAL RISK · SCENARIO ANALYSIS

The 2028 Intelligence Scenario: A Business Translation

A prominent financial research firm published a scenario where AI capability outpaces institutional adaptation — and it maps the exact transmission mechanism from "AI disrupts white-collar work" to "systemic financial risk." Here's why it matters, and what to do about it before the scenario stops being hypothetical.

By **Pedro Bandeira** · Analog AI · February 2026 · 11 min read

In February 2026, Citrini Research — a financial research firm followed closely by institutional investors — published a scenario analysis that I haven't been able to stop thinking about. Not because it's doom-mongering. Because it's engineering.

The piece is written as a fictional memo dated June 2028. It traces, step by step, how AI capability could transmit from "software disruption" through to "the US mortgage market faces a structural crisis." It names specific companies as examples. It models the feedback loops. It identifies exactly where the system has no natural brake.

I want to be precise about what this is: it is a thought experiment, not a forecast. The authors say so explicitly. But thought experiments from people whose job is to be right about systemic risk are worth reading differently than opinion pieces. They're asking: *if the capability continues accelerating at its current rate, and if institutions don't adapt fast enough, what breaks, and in what order?*

The answer, laid out across five phases, is instructive — not because it will happen exactly this way, but because the vulnerabilities it identifies are real, and they exist today.

▣ BEFORE WE CONTINUE

This article references the Citrini Research scenario analysis published February 22, 2026. I'm not a financial analyst and this is not investment advice. What I am is someone who implements AI systems inside regulated businesses and has spent twenty years watching companies get caught by structural changes they could see coming but chose not to prepare for. That's the only lens I'm applying here.

Five Phases of a System Under Stress

The scenario describes five phases of transmission — each one building on the last. I've compressed them here into the questions they raise for business leaders, not the financial mechanics of each phase.

PHASE	WHAT THE SCENARIO DESCRIBES	THE BUSINESS QUESTION IT RAISES
PHASE 1 Software	Agentic tools allow companies to replicate SaaS functionality in weeks. Enterprise software vendors face margin collapse.	Are your vendor contracts built on assumptions about what software will be worth in two

PHASE	WHAT THE SCENARIO DESCRIBES	THE BUSINESS QUESTION IT RAISES
Disruption 2026	Seat-based pricing breaks as headcount falls.	years? Are you locking in multi-year commitments now?
PHASE 2 Friction Elimination 2026–2027	AI consumer agents automate purchasing. Intermediation value collapses across travel, insurance, real estate. Platform loyalty disappears when machines optimise every transaction.	If your business model depends on friction — information asymmetry, switching costs, habitual repeat engagement — what replaces it?
PHASE 3 Labour Displacement Spiral 2027	White-collar displacement accelerates. Displaced professionals flood service and gig roles, compressing wages sector-wide. Top earners face structural income impairment.	Does your workforce resilience strategy assume the 2024 job market will still exist in 2027? Have you modelled what workforce anxiety does to productivity and retention?
PHASE 4 Private Credit Stress Q3 2027	PE-backed software companies that built on human-labour assumptions face revenue collapse. Private credit defaults begin. Life insurers holding this debt face capital pressure.	Are your banking or financing arrangements tied to lenders with hidden exposure to this? Are your own credit facilities stress-tested against a revenue environment where your SaaS vendors disappear?
PHASE 5 Mortgage Market Stress Late 2027– 2028	Prime mortgage borrowers — historically bulletproof — face structural income collapse. Home prices fall in tech-heavy metros. A \$13T market built on income assumptions begins to crack.	The scenario's conclusion: economic systems built around scarce human intelligence become unstable when that scarcity disappears. Policy institutions move too slowly to manage it.

Unlike cyclical recessions, AI capability *improves* during economic weakness. When companies face pressure, they cut payroll — and reinvest those savings into AI tools. This creates what the scenario authors call a negative feedback loop with no natural brake. Waiting for conditions to improve before you start adapting is not a strategy. It is the opposite of one.

Why the Underlying Logic Is Sound

I don't know if this scenario will materialize. Nobody does. But strip away the timeline and the financial contagion mechanics, and what you're left with is a set of structural vulnerabilities that are demonstrably real right now, in February 2026.

The SaaS assumption is already breaking. The scenario's Phase 1 — where AI enables companies to replicate enterprise software functionality in weeks — isn't hypothetical. I'm doing this for clients today. The tools exist. The economics are compelling. Enterprise SaaS pricing built on per-seat, per-month recurring revenue faces a genuine structural challenge when the intelligence required to build competing functionality becomes abundant and cheap.

White-collar work is being restructured faster than anyone admits publicly. The labour displacement isn't coming — it's happening in slow motion right now, concentrated in specific roles: data analysis, document review, first-draft generation, customer service tier-one, code review. The question isn't whether this is occurring; it's whether companies are managing it deliberately or just watching it happen and hoping for the best.

Income assumptions in financial models are lagging reality. The scenario's mortgage crisis is extreme. But the observation underneath it — that financial models assume income stability that may not materialise for white-collar professionals — is worth taking seriously. If you're a law

firm, a consultancy, or a professional services firm, your clients are running the same models. The disruption to their revenue affects their ability to pay yours.

Policy will be late. The scenario describes bipartisan political gridlock preventing effective response. That's not a 2028 prediction — it's a 2026 description. The EU AI Act debate, the US AI regulatory vacuum, the G7 disagreements on AI governance: these aren't early-stage problems being solved efficiently. They are late-stage problems being managed badly. Counting on regulation to set your strategy is, as always, the slowest possible approach.

"The system's vulnerability stems from decades of assumptions that human intelligence remains scarce. Multiple leverage points were built on that scarcity simultaneously."

— Citrini Research scenario analysis, February 2026

The Three Types of Company in This Scenario

Scenarios like this aren't useful for predicting the future. They're useful for identifying which side of a structural divide you're currently on. In this scenario, there are broadly three types of company:

×
EXPOSED

• AI
pilots

POSITIONED
TO
ADAPT

- AI is already embedded in workflows under human governance
 - Workforce has been upskilled deliberately, not just exposed to tools
 - Vendor stack is rationalised — not locked into per-seat SaaS at scale
 - Business model doesn't depend on information asymmetry or
- running but nothing in production governance
 - Staff using ChatGPT, Copilot, and Gemini ad hoc with no policy
 - Workforce strategy assumes 2023 job descriptions will still make sense in 2027
 - Multi-year SaaS contracts recently signed at peak pricing
 - Leadership saying

switching	"we'll
costs	address
• Has	AI
an	governance
explicit	when
AI	it
policy:	becomes
what	a
data	regulatory
leaves	requirement"
the	
perimeter,	
who	
approves	
AI	
outputs	

→
**COULD
GO
EITHER
WAY:
THE
MIDDLE
GROUND
(MOST
COMPANIES)**

- Has done some AI experimentation but lacks a coherent architecture
- Governance conversations

are
happening
but
not
completed

- Workforce
has
mixed
levels
of
AI
literacy

—
some
enthusiastic
early
adopters,
significant
resistance
from
others

- No
formal
AI
risk
classification:
nobody
has
actually
mapped
which
tools
touch
sensitive
data,
hiring
decisions,
or
client-
facing
outputs

Most of my clients start in the middle category. That's not a criticism — it reflects how fast the landscape has moved. The ones who will be in the "positioned to adapt" column by the time any version of this scenario plays out are the ones making deliberate decisions now, rather than waiting for a compliance trigger to force the issue.

The Human Side Nobody Is Pricing In

The Citrini scenario's Phase 3 includes something that most AI strategy conversations skip entirely: the social consequences of workforce displacement at scale. The scenario describes demonstrations blocking AI labs, public anger targeting compute owners who captured most of the productivity gains, and a "fraying social fabric as displacement accelerates faster than institutions adapt."

I'm not going to argue whether this specific social backlash materialises. What I will argue is that the human cost of poorly managed AI transitions is already visible inside individual companies — it just doesn't show up on the balance sheet yet.

When companies deploy AI without workforce preparation, without clarity about what it means for roles, without wellbeing support for people navigating genuine professional anxiety — they don't just create individual distress. They destroy the trust and psychological safety that allows people to actually use new tools well. They create resistance that slows adoption. They lose the people who choose to leave rather than work in an uncertain environment, which is often the best people, because they have options.

The companies that navigate structural AI transitions well — at sector level and at company level — are not the ones who automate the most aggressively. They are the ones who manage the human transition deliberately: investing in psychological safety, communicating honestly

about what is changing and what is not, and building the workforce capability that makes AI actually valuable rather than just deployed.

This is the other side of the coin that any serious AI strategy has to address. Technical governance and human governance are not separate problems.

What To Do Right Now

The scenario's authors conclude with three words I find hard to argue with: *time is the real villain*. Not the technology. Not the economy. The gap between how fast capability advances and how fast institutions — companies, regulators, labour markets, financial systems — adapt to it.

You cannot close that gap by waiting for clarity that isn't coming. You close it by starting with what you can control:

- **Map your AI exposure — now, not when the regulation requires it**

List every AI system your company uses, including the ones embedded in SaaS tools you already pay for. Copilot, Salesforce Einstein, any AI-assisted recruiting or contract review tool. For each one: what decisions does it influence? Who approves its outputs? What data does it see?

- **Review your workforce strategy against 2027 job descriptions, not 2023 ones**

Which roles in your organisation are most exposed to structural change in the next 18–24 months? This is not about replacing those people — it's about investing in their transition before you're forced to manage it in crisis mode. The companies in the "positioned to adapt" column have already started this conversation.

- **Audit your SaaS contracts for Phase 1 exposure**

Where are you locked into per-seat, per-month pricing for tools that could be replicated — or where the value proposition depends on functions AI is actively commoditising? This isn't a reason to cancel everything tomorrow. It's a reason to stop renewing blindly on auto-pilot.

- **Establish an AI governance policy before you need one**

What data is permitted to leave your perimeter? Which AI outputs require human review before acting on them? Who is accountable for AI decisions in client-facing contexts? These are not hard questions. They take a day to answer. Not having answered them is a choice.

□ **Build human resilience alongside technical resilience**

AI literacy training, honest communication about how AI is being used in your organisation, psychological safety for people raising concerns about AI tools — these are not soft HR programmes. They are the difference between an organisation that adopts AI well and one that fights itself while trying to.

□ **THE EUROPEAN CONTEXT**

The Citrini scenario is written primarily through a US lens — US labour markets, US mortgage market, US private credit architecture. The European context differs: stronger labour protections slow displacement velocity, different financial system architecture changes the contagion mechanics, and the EU AI Act creates governance obligations that force some of this conversation whether companies want it or not. But the underlying dynamic — AI capability outpacing institutional adaptation — is not geographically bounded. The question for European SMBs and regulated industries is not whether to prepare. It's how to prepare efficiently, proportionately, and without losing the operational momentum that makes adaptation possible.

The Only Response That Makes Sense

I've been doing this long enough to know that the right response to a scenario analysis is not to build a bunker or to dismiss it as fearmongering. It's to take the vulnerabilities seriously, make a

proportionate assessment of your own exposure, and act on the things you can control before the timeline makes them urgent.

The companies that navigate structural transitions well are not the ones with the most sophisticated AI. They are the ones that prepared deliberately — technically, organisationally, and humanly — before the scenario stopped being hypothetical.

The scenario says time is the villain. I'd say it differently: the villain is the gap between knowing you should prepare and deciding you can wait another quarter. That gap is the only thing in this whole analysis that you actually control.

Not Sure Where Your Company Sits?

The gap between "exposed" and "positioned to adapt" is almost always smaller than it looks from the outside. Start with a conversation about what you actually have in place — and what the realistic gaps are.

[Request a Reality Test →](#)

Pedro Bandeira is a Fractional CAIO and AI Governance Advisor working with regulated SMBs across Europe. He is based in Warsaw and operates across EU/UK time zones. This article is analysis, not financial or legal advice. The Citrini Research scenario referenced is a thought exercise published February 22, 2026 — it is not a prediction. [Read more on Analog AI Pulse →](#)