



When more becomes too much

Marcel Prins
1 October 2026

CFA Society Netherlands - AI in Risk Management

Technology makes things cheaper

Technology wave	What became cheap
Electricity (1800-1930)	Mechanical power
PC (1980-2000)	Computing
Internet (1995-...)	Communication & distribution
Social Media (2010-...)	Audience building
AI (2020-...)	Intelligence & execution

\$\$\$

\$\$\$

\$\$\$

... and creates value in risk management



Credit risk	Market & liquidity risk	Operational risk	Compliance & reg risk	Model risk	Operational efficiencies
• AI-driven credit scoring and underwriting using bureau, transaction, and alternative data • GenAI drafting credit memos, summarising customer information, and flagging policy gaps • Early-warning systems for portfolio monitoring using financials, news, and market data • Collections optimisation and restructuring scenario analysis • ...	• Real-time market risk monitoring (VaR, stress metrics) with ML • Liquidity and cashflow forecasting using behavioural and macro data • Trading strategy support, backtesting, and capital optimisation • ...	• Real-time fraud and anomaly detection in payments and cards • AML alert triage, false-positive reduction, and SAR narrative drafting • Conduct surveillance in communications and trading behaviour • Cyber threat detection via logs, network traffic, and user behaviour analytics • ...	• GenAI summarising new ECB/DNB/AFM rules and mapping to internal policies • Automated regulatory reporting and transaction monitoring narratives • SupplierTech analytics for internal self assessment and third-party inspection prioritisation • ...	• AI-assisted back-testing, benchmarking, and drift detection for credit, market, and fraud models • GenAI maintaining model inventories and drafting model documentation • AI supporting model validation workflows and challenge processes • ...	• GenAI drafting routine risk reports, board papers, and RCSA documentation • Automated data extraction and reconciliation across risk systems • AI-supported workflow orchestration • Knowledge search and Q&A over internal policies, procedures, and past decisions • ...

AI (2020-...)

Intelligence & execution

\$\$\$

\$\$\$


\$\$\$

... with new bottlenecks emerging

Technology wave	What became cheap	New bottleneck
Electricity (1800-1930)	Mechanical power	Distribution
PC (1980-2000)	Computing	Software
Internet (1995-...)	Communication & distribution	Relevance & attention
Social Media (2010-...)	Audience building	Trust & differentiation
AI (2020-...)	Intelligence & execution	Trust, originality & judgement

The illusion of certainty

- You read a perfectly formatted, assertive paragraph and assume it represents deterministic, reliable truth, forgetting that it is still statistical guesswork.
- This is the illusion of certainty: one of the core risks in enterprise AI adoption.




'Too much' would be stupid, not?

Defining stupidity:

1. Cognitive offloading
2. Wilful ignorance
3. Conformity



Shades of AI usage



AI overreliance examples

- Risk analysts accept AI-generated model documentation, stress-test scenarios, or parameter estimates without challenging assumptions or re-running key calculations.
- A compliance officer uses an LLM to map AFM/DNB/ECB requirements to internal policies, then implements controls based on a plausible but incorrect interpretation.
- Security or operational risk teams trust AI summaries of logs or incidents, close cases without verifying root causes, and leaving control gaps as a result.



Wrong output

Reduced oversight

Bias amplification

AI overdependence examples

- Staff stop developing or exercising core skills because AI drafts everything.
- Critical risk processes become structurally dependent on a small set of AI tools.
- Boards and senior management treat AI-supported risk reports as 'good enough', reducing deep-dive challenge and independent validation.




Skill atrophy

Operational fragility

*Loss of
independent judgement*


Who is most at risk?

- AI-overdependence poses a significant threat to cognitive skills. It also undermines creativity, memory, and executive functions.
- Those most vulnerable are:
 - younger learners
 - less formal educated
 - professionals in high-stakes fields



I understand
you're the
human in
the loop.

**AI APPROVAL
REQUESTS**



'Loop' differences

	In-the-loop	On-the-loop	Next-to-the-loop
When humans act	Before decision / action	During / after	Nominally...
Typical context	High-stakes, regulatory, irreversible	High-volume, time-sensitive	Poorly designed oversight, excessive automation pressure
AI autonomy	Low	Medium / high	High, human oversight largely a formality
Human authority	Final decision-maker	Supervisor with override power	Formal accountability, limited actual influence
Risk profile	Low overreliance risk	Medium, needs strong monitoring	High overreliance and accountability risk

The TIME litmus test



T - Time

Is there sufficient time to review the output?

I - Information

Are the source data and assumptions available?

M - Mandate

Is the reviewer actually authorised to intervene?

E - Expertise

Can that person recognise the error?

No TIME? Then the human is probably next-to-the-loop!

Practical mitigation strategies

To avoid overreliance

- Targeted AI literacy, incl limitations
- Create friction with user prompts
- Show model confidence in context
- AI recommend-or-decide thresholds
- Independent verification / 4-eyes
- Log AI usage, sample it, discuss it (RSCAs, audit, etc)

To avoid overdependence

- Invest in skill retention
- Define AI knowledge boundaries
- Define no-AI zones
- Deliberate AI-off exercises
- Treat AI tooling as critical infrastructure incl. citizen AI
- Add AI risk & dependency metrics in ERM framework



AI as a risk driver in your ERM framework

Risk category	Underlying risk
Model risk	Outputs may be inaccurate, unstable, biased or insufficiently explainable
Conduct risk	Employees or customers may rely on AI outputs inappropriately
Operational risk	Critical processes may fail when AI systems, data or providers are unavailable
Human-capital risk	Overreliance may weaken professional expertise, challenge and independent judgement
Third-party risk	Dependence on external models, platforms and data providers may create concentration risk

Culture eats strategy for breakfast

Enablers

- Growth mindset
- Professional pride
- Constructive scepticism
- Leadership example
- Psychological safety

Blockers

- AI-knows-best culture
- Skill complacency
- Efficiency-at-all-cost mindset
- Siloed mindset
- Displacement anxiety



**Use AI to scale
human judgement,
not to outsource it**



marcel@covalens.nl
+31 6 1002 1677

