cognizant®

Cognitive debt

and the erosion of insurance expertise

Introduction

The insurance industry has spent the past few years exploring AI's potential. Early wins came in IT and other processes with minimal regulatory exposure. **Now many firms are turning to the core of insurance: underwriting, claims, and service.**

This shift is not without risks, and requires insurers to navigate carefully. In particular, insurers must determine how and where to apply non-deterministic AI tools in an industry where regulators demand predictability and auditability.

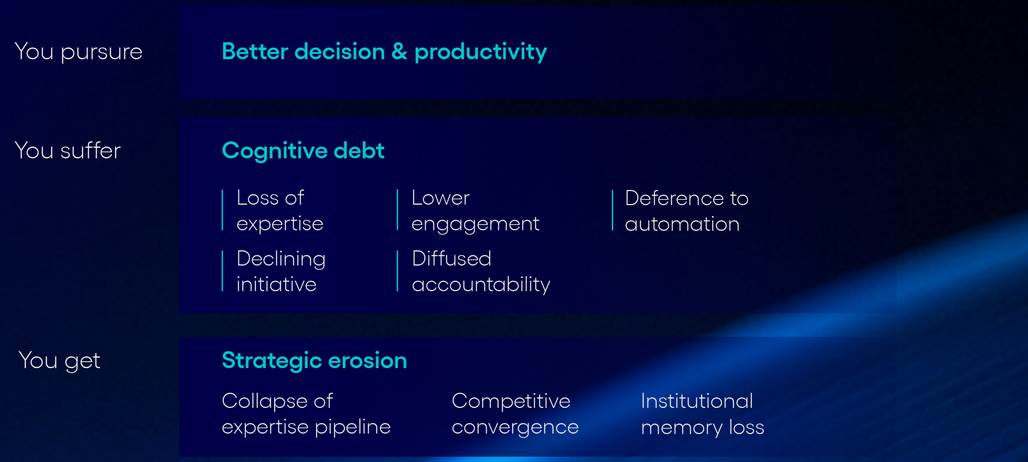
While the industry's ability to implement AI has matured, we have mostly ignored another risk: what AI use does to the knowledge workers who increasingly depend on it. We call this **cognitive debt**, leveraging some foundational thinking from academia.

Cognitive debt: the erosion of human capability that occurs when people rely on AI as a substitute for their own expertise.

Like financial debt, cognitive debt is easy to accumulate and hard to discharge. Both are also predictable in that humans are more inclined to think about short-term benefits and convenience than they are about long-term costs.

Consider how the cognitive debt cycle in insurance works. Insurer staff pursue productivity improvements and decision support, and along the way their behavior changes in foreseeable ways. These changes may be subtle, but over time they lessen the amount of human expertise in the firm, diminish institutional memory, and impact the firm's ability to differentiate itself from competitors.

To our knowledge, this dynamic has never been studied within insurance. Given AI's explosive growth and the long-term risks associated with cognitive debt, we decided to pursue this data-driven analysis.



A note on the scenarios in this report: several findings open with a short scene. These are illustrative composites — not real firms or real people — written to show what the dynamics of AI feel like to insurer leaders. The analysis that follows are findings from our survey of 542 leaders across property & casualty, life and annuities, specialty, reinsurance, and retirement. The survey was fielded in August 2026.

Cognitive debt and the erosion of insurance expertiese, a survey of 542 Director-and-above insurance leaders, fielded August 2026.



What the industry is chasing

Let's back up for a moment to think about the industry's goals in implementing AI.

You might think that cost improvements are the drivers, but you'd be at least partly wrong. According to our survey, the industry ambition is anchored around improvements in decision making, experience, and speed.

Asked what is driving AI adoption, insurance leaders point first to decision quality (47%), customer and distribution experience (46%), and speed in underwriting and claims (45%). Cost is further down the list. This is not an industry chasing headcount reduction via AI, though it often does that in other ways. It is an industry chasing better, faster decisions.

The industry is moving toward the next phase: deploying AI where the judgment lives. Our respondents said their firms' AI work outside of IT centers on claims (53%), customer service (51%), and underwriting (48%). Internal functions like fraud prevention (24%), marketing (22%), distribution and agent support (21%), and actuarial and pricing (19%) trail well behind. Insurers are aiming AI squarely at the work where seasoned human judgment has always been the product.

Which raises a central question: If you use AI to replace human expertise, what happens to that expertise? How will humans respond to this rewiring of the enterprise?

Cognitive debt defined

The concept of cognitive debt was established decades ago. Various research efforts have described what happens when people offload thinking due to the presence of an external aid. For more, read about transactive memory (Wegner), cognitive offloading and the "Google effect" (Sparrow, Risko & Gilbert), and automation complacency in high-stakes settings (Parasuraman & Manzey, Bainbridge). More recently, a study by the MIT Media Lab discussed an AI-specific version of cognitive debt and identified reduced engagement and ownership in the human brain when a language model does the first draft of the thinking.

The concept takes on increased importance in the age of AI, but it has not been examined in insurance specifically. That's why we created the Cognitive Debt Index (CDI): a measure of perceived drift toward over-reliance on AI, and administered it via survey to 542 insurance leaders.

CDI is composed of four dimensions

Skill atrophy

People lose proficiency in the tasks AI now handles.

Judgment deference

People defer to AI instead of exercising their own judgment.

Agency and initiative

People stop driving problems they assume AI will handle.

Automation bias

People accept AI output without challenge, even on high-stakes calls.

The CDI is the average of the four dimension scores. We asked respondents to indicate their level of agreement with two statements about potential negative impact of AI in each dimension, on a 5-point scale where 1 means Disagree Strongly and 5 means Agree Strongly. Thus, higher agreement scores indicates more perceived debt.

Now let's move on to our findings.

Note. To test the validity of the CDI, we looked at the statistical correlation of index components. We also compared the CDI results to results of a separate question about the respondents' concern about AI's impact on their workforce and found the two to be statistically, positively correlated.

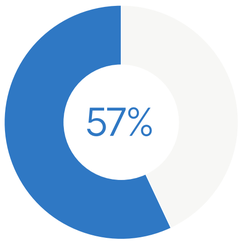
Finding 1

Insurance expertise is at risk

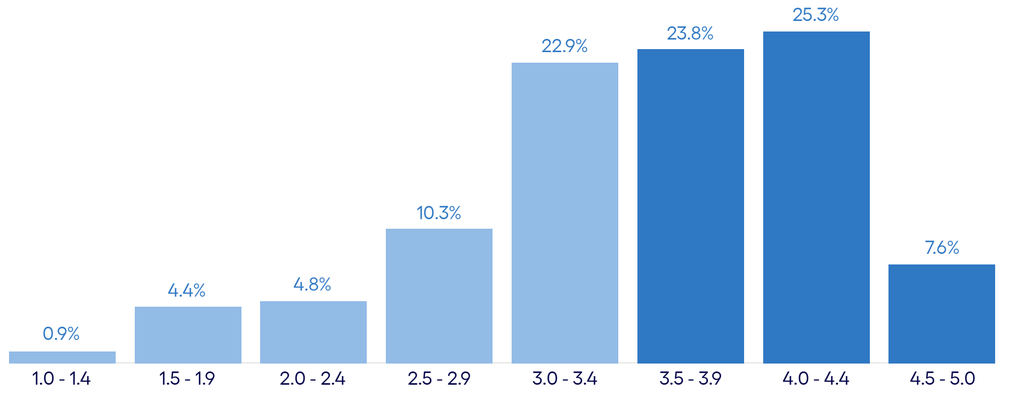
Marcus has run underwriting at a mid-size commercial carrier for eleven years, and he can remember how he learned the craft: badly priced accounts, corrected by someone senior, the reasoning drilled in until it became instinct. His newest underwriters are not learning that way. The model prices the account; they check that it looks reasonable and move on. They are faster than he was at their experience level and, on the hard accounts, thinner — they have never sat with a risk long enough to develop a feel for it, because the model never asks them to. Marcus is not worried about this year's numbers. He is worried about who is running his desk in five years, and whether the expertise that took him eleven years to build is being cultivated by their development process.

For most of the industry this is not hypothetical. **The overall CDI stands at 3.48 on a 5-point scale, above the midpoint on all four dimensions.** Scores are also tightly centered on the mean: 57% of the insurance leaders we surveyed fell in the high-debt bands, a score of 3.5 or higher. At that level, high cognitive debt appears to be an industry baseline.

Skill Atrophy is among the highest-scoring dimensions, and the item that drives it is the one Marcus is living: leaders agree that newer staff are not building the expertise earlier cohorts did. The tools are teaching people to get the answer without teaching them the complex storylines behind the answer.



of leaders in the high-debt bands (CDI ≥ 3.5)



CDI score (1 = low debt → 5 = high debt)

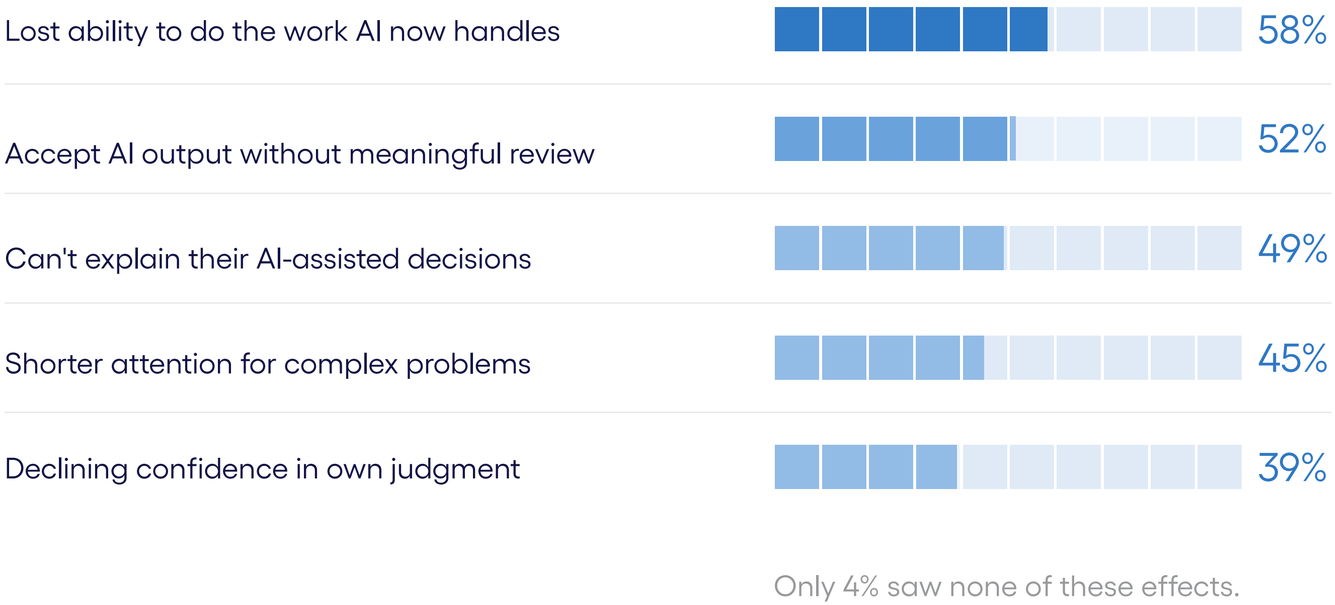
Finding 2

This is observed behavior, not a forecast

Fast forward a year or two. Priya runs claims at a personal-lines insurer, and the settlement figures her adjusters produce are, by every audit she runs, accurate. That accuracy is…unsettling. In a file review she found three approvals in a row where the adjuster had accepted the model's recommended figure without a note, without a flag, without the small hesitation she used to see on a borderline call. The LLM adjusters use is spot on often enough that questioning it has begun to feel like wasted effort. Priya knows a large claim will eventually come along that the model gets badly wrong. What she cannot promise her division head is that anyone on her team will be in the habit of catching it.

The erosion from cognitive debt is not something leaders fear might happen. It is something they report watching happen today. Since scaling AI, they say they have personally observed staff who have lost the ability to do the work AI now handles (58%), staff accepting AI output without meaningfully reviewing it (52%), and staff who cannot explain the reasoning behind their AI-assisted decisions (49%). Shorter attention for complex problems (45%) and declining confidence in people's own judgment (39%) follow close behind.

The most telling number is the one that almost nobody chose: only 4% said “none of the above / too early to tell.” Ninety-six percent of insurance leaders can already name the erosion that they see occurring around them.



Only 4% saw none of these effects.

Finding 3

Scaling outran planning

Sixty-two percent of insurers are scaling or maturing AI in production today. Fifty-eight percent say the long-term impact on human capability did not materially shape that rollout.

Put those two numbers together and the problem comes into focus: the industry moved AI into production faster than it worked out what that meant for how people think and act. That is not carelessness so much as the ordinary physics of a fast rollout, an industry trying to move “at AI speed.” The human-capability question is slower and quieter than a throughput target, and it loses every time the two compete for attention. But deferring consideration of negative impacts does not make them free. It is a recipe for cognitive debt accumulation.

62%

are scaling or maturing AI in production today



58%

say human impact didn’t materially shape that rollout

Finding 4

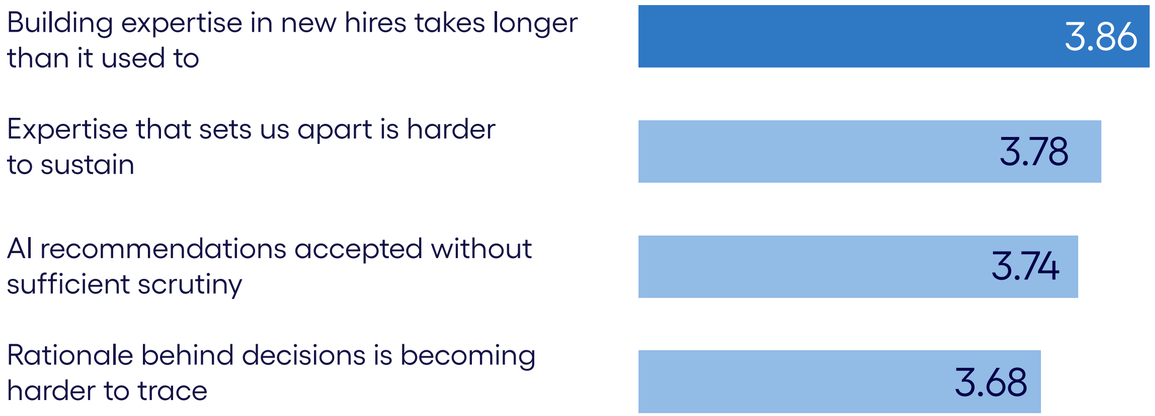
Leaders can see what it will cost

Tom, a chief actuary at a life and annuity carrier, faces a critical compliance risk: he cannot fully reconstruct how an AI-assisted reserving assumption was derived. Although the assumption shapes product pricing and appears sound, the weeks-long, iterative workflow—including the specific AI prompts and intermediate reasoning—was never documented. In an industry built on auditability and regulatory transparency, this missing paper trail leaves Tom with an unverified conclusion and a growing concern that other undocumented models exist across the business.

Insurance leaders are not naive about where this leads. They rate as likely a set of concrete business consequences: mispriced or poorly understood risks, claims and reserving errors, increased regulatory and market-conduct scrutiny, and client or distribution dissatisfaction. And they see the mechanism underneath — agreeing, at means of 3.68 to 3.86, that building critical expertise in new hires now takes longer, that the expertise which differentiates the firm is harder to sustain, that AI recommendations are accepted without sufficient scrutiny, and that the rationale behind decisions is becoming harder to trace.

The danger here is not one dramatic AI failure. It is subtler and worse: an organization that slowly loses the ability to explain its own decisions, or to catch them when they go wrong.

Q: How is AI affecting underwriting and claims in your firm?



mean agreement, 1-5 scale (n=542)

Finding 5

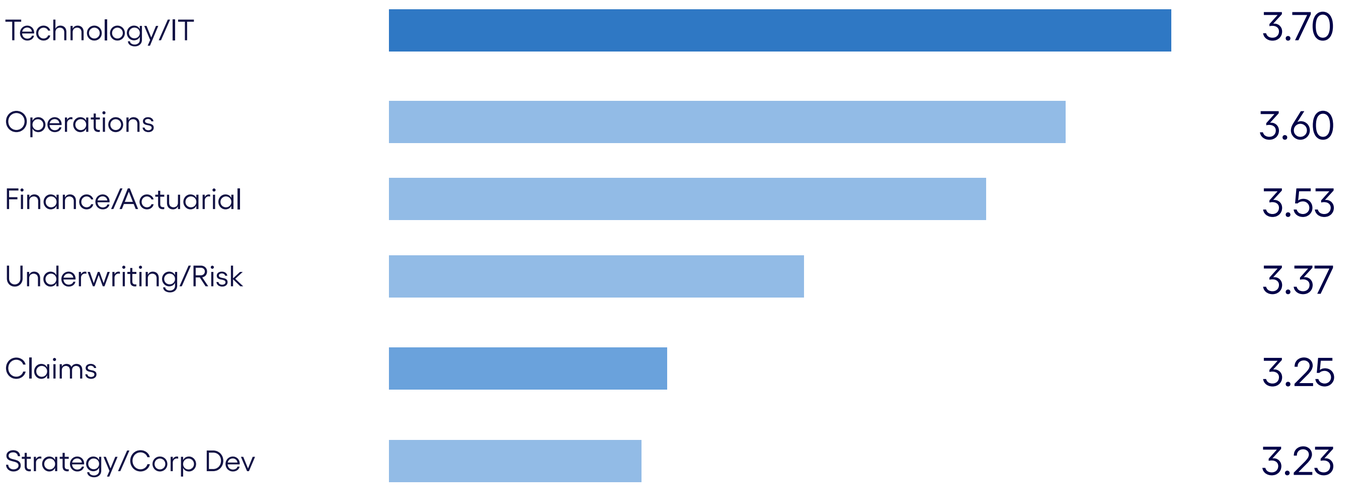
The experts are the most worried

Ana and her team built the platform. As head of data science, the LLM now powering underwriting and claims is her output, and this makes Ana uneasy. She knows precisely where the models are strong and where they are fuzzy, and she watches her colleagues extend trust into places she would not. To everyone else she is the builder, so she must be the believer; in fact her fluency is what makes her cautious. Adoption keeps climbing and the organization reads the curve as success. Ana reads the same curve and sees how much is being trusted by people who cannot see what she sees.

Ana is not an outlier, and the numbers say so. Sort the CDI by each respondent's own functional background and the people closest to the tools report the most cognitive debt. Technology and IT leaders score 3.70, the highest of any function; front-line claims sits lowest at 3.25, with underwriting in between.

Call this a vantage-point effect, and read every CDI number through it. IT departments are not more debt-ridden than the rest. Every respondent was rating the whole organization. What changes with fluency is what a person can see. The people who understand the machine best tend to have the sharpest feel for its limits. So the reassuring number from your front line and the alarming number from your technologists are describing one and the same firm — and it matters a great deal whose reading you act on.

Cognitive debt index by respondent function



Finding 6

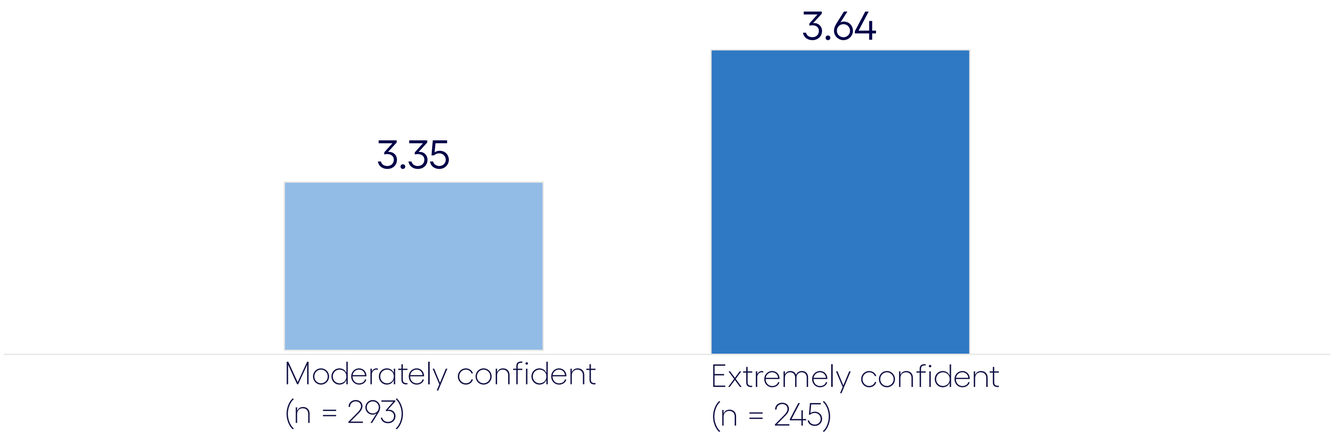
The confidence paradox

David is COO of a specialty carrier, and he would tell you without hesitation that his AI program is a success. The scorecard is green, the leadership team is aligned, no one brings him more than an occasional hiccup. What David has not registered is that certainty this complete is a warning sign. The people who tend to be most sure an AI deployment is working are the ones no longer looking hard for the ways it isn't — conviction has quietly replaced scrutiny.

This is a finding that should trouble the most confident insurer leadership team. Leaders extremely confident their firm is capturing value from AI score a CDI of 3.64. Those only moderately confident score 3.35 — a gap of nearly three-tenths of a point on a five-point scale.

In other words, the firms most sure they are winning with AI may be carrying the most cognitive debt. That is not a coincidence to wave away. Deference is comfortable, and automation bias is confident by definition. A process stops generating friction when people stop pushing back, and that reads as success right up to the moment someone needs to catch a mistake and no one can. Confidence tells you how a team feels. It does not tell you whether the judgment underneath it is still there.

Cognitive debt index by confidence of respondent



Q: How confident are you that your organization is capturing value from its AI investments?

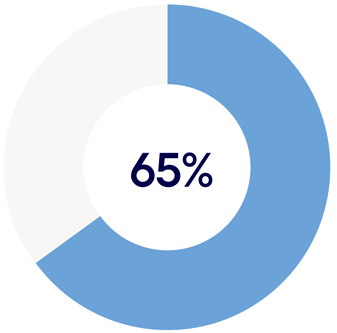
Finding 7

Success has a cost

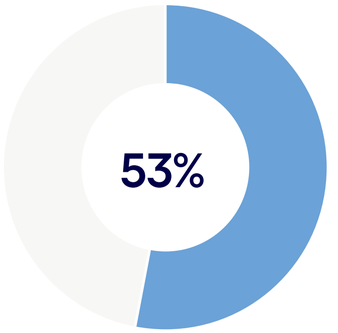
Rachel is chief transformation officer at a large multiline, and by the metrics she was hired to move, she is winning — the productivity gains are not incremental but transformative. There is a second pattern in the same teams that no dashboard tracks. They have stopped bringing her problems. They used to surface odd cases, argue for exceptions, propose things the process did not anticipate; now the work flows through the model and comes out clean, and the initiative that used to sit on top of it has gone quiet. Her fastest teams are also her most passive. Rachel was promoted for the first half of that sentence, and she is beginning to grasp that it may be inseparable from the second.

The data suggests she should worry. Firms reporting transformative AI productivity gains score higher on cognitive debt than everyone else — a CDI of 3.58 versus 3.43 ($p < 0.05$). Put in plain terms: 65% of the transformative-gains firms sit in the high-debt band, against 53% of everyone else — a twelve-point difference.

The most successful AI adopters appear to be accruing debt the fastest. The gains are not illusory — they are real, and they are the whole reason anyone is pursing AI. But the gains and the debt come out of the same motion, and a scorecard that tallies only the first is telling half the story. Speed has a cost.



of transformative-gains firms sit in the high-debt band



of everyone else — a twelve-point spread

Q: Over the past 12 months, describe the level of AI-driven productivity gains your firm has achieved.



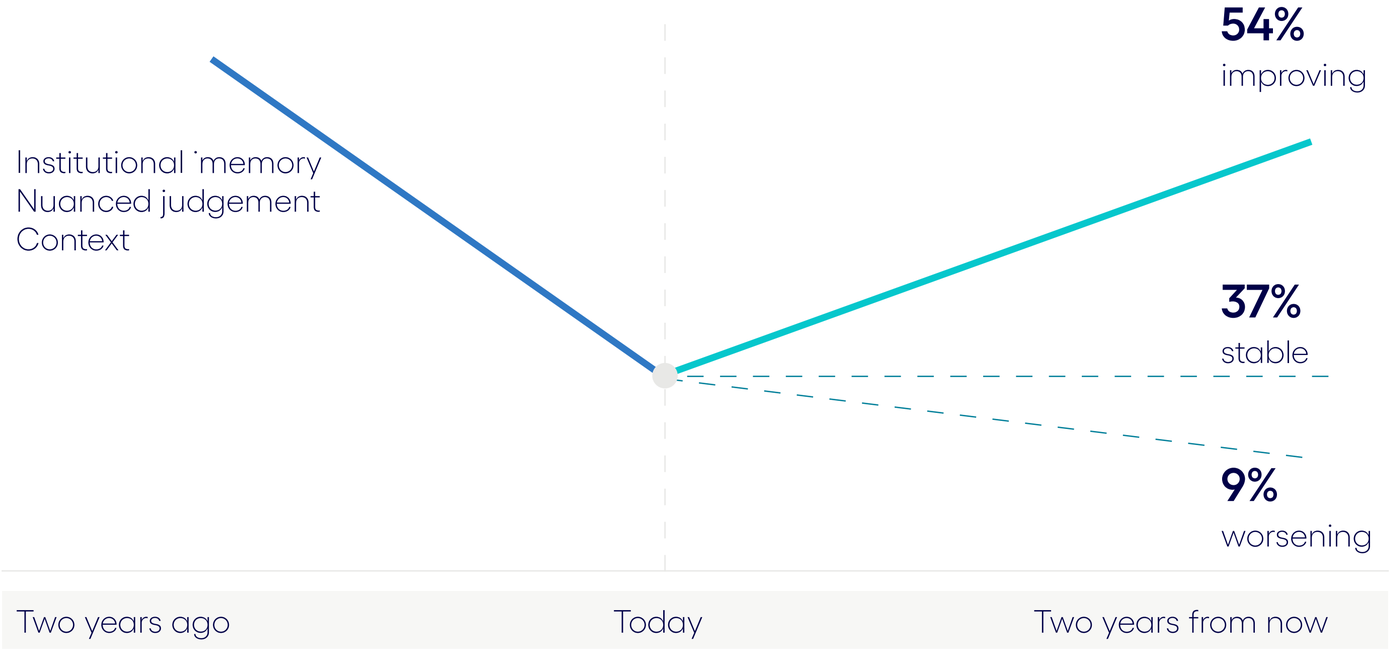
Finding 8

Reason for optimism

Taken together, we believe our survey findings suggest the industry faces a real issue of cognitive debt. Even so, we found a ray of hope in our data. When asked how AI-related loss of human capability is trending in their organization, 54% of leaders say it is improving, 37% say it is stable, and only 9% say it is worsening.

Put another way, nine in ten leaders believe the decay has leveled off or turned around. The debt piled up fast across two years of nearly unmanaged adoption; the question now is how quickly that debt might come back down.

That optimism is worth something only if it is acted on. Recovery does not happen because leaders feel hopeful about it. It happens when someone puts a number on the debt, hands it to an owner, and works it down on purpose. The firms that treat “improving” as a plan rather than a mood are the ones for whom it will become a reality.



Your next steps

The thread that runs through our findings is that the industry has managed the model risk of AI far better than the human risk. Governance frameworks watch for model bias, drift, and compliance gaps. Very few watch whether the people around the model still have the judgment to catch it when it is wrong. That is a gap to close.

We propose you address your cognitive debt in two ways. First, assess the human side. Cognitive debt remains hidden until someone takes specific responsibility for it, measures it, and reports out the results. The fix isn't creating more overhead. It's building judgment checks into the moments where people already work alongside AI, defining how human capability sits next to AI in key processes, and turning those observations into metrics that are tracked over time.

Second, you should deliberately manage down any existing cognitive debt. Firms that recognize the strategic implications will create workflows that preserve judgment, and build governance that delivers consistency across the enterprise. The compounding effect of those results over time will be powerful.

If you need assistance getting started, we're happy to help. Cognizant's expertise as an AI builder and our learnings from this study can be delivered via:

- **An executive briefing.**
  We will run a 60-minute session for your leadership team on our Cognitive Debt findings and what they mean for your organization. This is the fastest way to get the conversation on track.

- **A custom advisory project.**
  These are highly tailored to your specific needs but would likely include components aimed at an organizational assessment, experience design for participants in your key underwriting, servicing, and claim processes, and governance.

To take the next step, contact us here, or reach out to the study authors.

About Cognizant Moment

Cognizant Moment is Cognizant's dedicated practice area that provides brands with a unique perspective to foster growth, attract new customers, and strengthen loyalty with intuitive and engaging content and products. Our approach combines big data insights with the imaginative application of AI and human-centered design to drive innovation and deliver hyper-personalized experiences. As AI takes on more of the judgment behind how organizations sell and service risk, Cognizant Moment designs for a simple principle: the technology should sharpen human expertise, not replace it.

Study authors



Craig Weber is head of insurance strategy for Cognizant. He leads a team responsible for Cognizant's insurance offerings, and frequently speaks and publishes as a technology evangelist. He has also been a passionate advocate for the insurance industry for more than 30 years. Prior to Cognizant, Craig was CEO of Celent, a well-known analyst firm focused in financial services technology. He began his career as a management consultant working for a Boston-based insurer.



Stephanie Wan is Head of Experience Strategy at Cognizant Moment for the Americas, driving growth for Banking, Financial Services and Insurance clients through experience transformation, design, and digital product strategy. A leading voice on experience-led transformation and the impact of AI on the future of work, she has led digital transformation for Fortune 500 financial institutions and previously built Idea Couture's business transformation practice as a Partner prior to its acquisition by Cognizant. She holds a BASc in Engineering from the University of Toronto and is a Lean Six Sigma Black Belt.



Mohan Krishnaraj is Global Head of Strategy, Design, Studios & Growth at Cognizant Moment. With over 20 years across design, technology, and business — including CXO and Experience Design leadership roles at Infosys WongDoody, HARMAN (Samsung), and Wipro — he helps organizations drive growth and transformation through design, AI, and human-centered experience. He has been named among India's Most Influential Business Leaders and has received global honors including the iF Design Award, Red Dot Award, Webby Awards, and CODiE Awards.

About the research

Cognitive debt and the erosion of insurance expertise is based on a survey of 542 Director-and-above insurance leaders fielded in August 2026. It operationalizes an established body of research, addressing topics such as cognitive offloading, transactive memory, and automation complacency specifically within the insurance industry. Findings reflect leaders' perceptions of their own organizations.



Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

© Copyright 2025–2027, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.



© 2025–2027 Cognizant