

PwC and Salesforce

The AI Trust Gap: The need for speed vs. the psychology of risk

Authors:

Scott Likens, Principal and Chief AI Engineering Officer,
PwC Global and US

Susan Emerson, Senior Vice President – Product: AI,
Analytics & Data, Salesforce



Table of contents



CEO sentiments and the Davos dichotomy



AI and the pace of change



The risk vs. trust mindset



Behavioral economics applied to AI



Transparency leads to trust



What now: human prompts and the new circle of trust



CEO Sentiments and the Davos Dichotomy

As CEOs and C-suite leaders convened earlier this year at prestigious events like the 54th World Economic Forum in Davos, Switzerland, the buzz around AI was palpable. Onstage, speakers championed AI's potential to revolutionize industries, drive growth, and foster innovation. It was all about efficiency, optimization, and staying ahead of the curve.

But behind closed doors, some of the conversations took on a more measured tone.

PwC and Salesforce brought a group of C-suite leaders together for a discussion led by special guest and global trust expert Rachel Botsman on 'The AI Trust Gap: Perception vs. Reality vs. Action', and we were intrigued by what we heard. While leaders acknowledged and even embraced the transformative power and potential of AI, many also expressed concerns about the inherent challenges posed by such quick and widespread adoption. They knew that not embracing AI was no longer an option, but they also understood the imperative to balance the needs of people with the promise made possible by technology.

"Rebuilding trust in the future" was the official theme at Davos – but the idea of "rebuilding trust" implies that, at some point, our trust in technology was compromised. Trust, after all, can't be created with code and algorithms – it should be earned through a deep understanding of the many things that make us human.

“Trust is a human emotion, but AI isn't emotional. AI may be able to diagnose diseases, but it can't comfort a patient at their bedside.”

Scott Likens

Behavioral economics applied to AI

The presence of AI in contemporary media is prolific; it's the equivalent of celebrity status.

AI is often both villainized and revered: villainized because some fear it will replace human creativity and even jobs but revered because it can expand human potential and improve efficiency and innovation when used responsibly.

Neither of these perspectives are absolute. As with most technologies, it's not about the solution, it's about how the solution is used.

Integrating AI isn't just a technical challenge — it's a human one. It requires a nuanced understanding of human nature and the delicate balance between optimism and caution.



Behavioral economics, which examines the psychological factors that influence human decision-making, can provide some valuable insights into how executives might make decisions about AI:

Cognitive biases:

Many of us are susceptible to cognitive biases that can affect our decision-making. For example, confirmation bias may lead executives to seek out information that confirms their preconceived notions about AI's benefits or risks, while availability bias (leaning on readily available information) may cause them to overestimate the likelihood of high-profile AI failures or successes.

Loss aversion:

Executives may be more sensitive to potential losses, such as financial setbacks or damage to their organization's reputation, than to an equivalent level of potential gains. This could lead them to be overly cautious or conservative in their approach to AI adoption.

Herding behavior:

Executives may base their decisions on the actions of others rather than on independent analysis. If they see competitors or industry peers rapidly adopting AI, they may feel pressured to follow suit to avoid falling behind, even if they haven't thoroughly evaluated the suitability or risks of AI for their organization.

Overconfidence:

Executives may underestimate the complexity of AI projects, overestimate their organization's readiness for AI adoption, or downplay the potential challenges and uncertainties involved. This overconfidence can lead to poor decision-making and inadequate preparation for the realities of AI implementation.

Temporal discounting:

Executives may focus on immediate cost savings or efficiency gains from AI adoption while overlooking potential long-term implications such as workforce disruption, ethical dilemmas, or societal impact. This shortsightedness can result in decisions that neglect important considerations for the future. The first step is recognizing these behavioral biases and tendencies – only then can executives take the necessary additional steps to mitigate their impacts.

This may involve a strong responsible AI governance process, implementing decision-making frameworks that encourage careful consideration of both short-term and long-term implications, seeking diverse perspectives from stakeholders to counteract groupthink, and fostering a culture of transparency.

Most of all, it's important to remember that success with AI isn't just about technology, it's about sociology. It's about the human intelligence behind the artificial intelligence, the decisions those humans make, and the trust they build along the way.



Here's what senior leaders should consider, from the CEO to the CISO, for GenAI:

What are we concerned about?

As the momentum to adopt generative AI surges, CEOs recognize that this technology could introduce potentially significant, unintended consequences.

Explore the CEO perspective and implications in "[The AI Challenge](#)" findings from [PwC's 27th Annual Global CEO Survey](#), "Thriving in an age of continuous reinvention".

- CEOs are most concerned about GenAI increasing cybersecurity risks — and over half agree that it is likely to increase the spread of misinformation in their company.
- One-third of CEOs expect GenAI to increase bias towards specific groups of employees or customers in the next 12 months. Almost as many disagree, suggesting bias is likely to be an area of growing attention as the scope and complexity of generative AI's role in business expands.

“As unanticipated consequences become apparent, it's up to entrepreneurs to implement, upgrade, or completely rethink the business models and structural mechanisms they have in place to reduce the negative impacts.”

Rachel Botsman, Author and Trust Fellow, Oxford University

When it comes to generative AI risks, CEOs are most concerned about cybersecurity


Cybersecurity risk




Spread of misinformation



Legal liabilities and reputational risks



Bias towards specific groups of customers or employees



■ Disagree ■ Neither agree nor disagree ■ Agree

Question: To what extent do you agree or disagree that generative AI is likely to increase the following in your company in the next 12 months?

Note: Disagree is the sum of 'slightly disagree,' 'moderately disagree' and 'strongly disagree' responses; Agree is the sum of 'slightly agree,' 'moderately agree' and 'strongly agree' responses. Percentages shown may not total 100 due to rounding.

Source: PwC's 27th Annual Global CEO Survey

AI and the pace of change

In 1948, Alan Turing, the “father” of modern computer science, wrote code for a machine to play chess against a human. This was one of the first examples of AI – but the past 12 months have arguably seen more groundbreaking AI advancements than Turing could have ever imagined.

Technology innovation today is accelerating – closing the gap between what’s “now” and what’s “next” at an increasing rate. This is thanks largely to generative AI, which removed barriers to entry, allowing anyone to use AI solutions via natural language creation of text, images, or video.

The pace of change is almost as monumental as the change itself. Moore’s Law declared that computing speed would double every 24 months, an unquestioned reality that held for decades. Generative AI has reduced that timespan to less than four months, and over the next 10 years generative AI could grow by more than 100,000,000x faster than Moore’s Law.

Speed – how to react to it and manage it – is an ever-present topic among executives who are feeling under pressure to mobilize and demonstrate AI competence. Whereas we had decades to prepare for previous technological advances from the industrial revolution to the Internet, with generative AI we’ve had mere months, adding to the impulse to step on the reinvention accelerator.

But rushing into AI without a clear understanding of your organization’s needs or risks can spell disaster. First, you should articulate your plans for AI adoption and how they align with your overall business objectives. And you should consider the impact on employees, how you can upskill your teams and set a strong foundation for success.

It’s essential to not let the fear of falling behind dictate your actions. Take the time to do your due diligence, to assess the opportunities and challenges, and to chart a course that makes sense for your organization.

AI isn’t going anywhere – it’s only going to become more commonplace. But while we can’t put the brakes on the future, we have a responsibility to get the future right. We should balance the needs of humans with the potential of technology as we continue to innovate and create.



Transparency leads to trust

As a business leader, you should engage with your stakeholders, listen to their concerns and address them head-on. Be upfront about the risks and uncertainties of AI, but also emphasize the potential benefits and opportunities AI presents.

Talk is only half the battle, however. As the survey illustrates, you should demonstrate that you're committed to doing things the right way, even if that means sacrificing speed in the short term.

Invest in strong governance and responsible AI structure. Take the time to thoroughly vet AI solutions before deploying them. Use AI to inform decision-making but not make decisions (keep humans "in the loop.") Most importantly, bring your employees with you and equip them with the new skills they will need to succeed.

Success is less about speed and more about direction – not the time it takes to adopt AI but confirming that you're moving down the right path.

In the rush to get ahead of the AI curve, how can leaders confirm "trust" doesn't get left in the dust?

It starts with transparency. According to PwC's 2024 Trust Survey, employees and consumers would like to see more transparency from businesses around their AI and data privacy policies:

- Nearly nine in 10 of both employees (89%) and consumers (88%) consider data privacy policies an important business practice for companies to publicly disclose, but only 32% of executives say their companies disclose these policies.
- With AI, 69% of employees and 66% of consumers say it's important for businesses to be transparent about their AI governance framework, but just 33% of business executives say they disclose this information.

Read [AI is Coming – Here’s How to Get Ready](#) to learn about getting employee buy-in.

Trust by design integrates privacy, security, transparency, and responsible practices from the very beginning of the design process, rather than as afterthoughts or additions. The goal is to create technologies and systems that users can trust inherently, not just because they’re secure, but because they’ve been thoughtfully designed and deployed.

PwC’s responsible AI Toolkit delivers trust by design through the overall AI life cycle with frameworks, templates and code-based assets that are relevant to executives at every level. The CEO and board set the strategy, with special attention to public policy developments and to corporate purpose and values. Chief risk and compliance officers are in charge of control, including governance, compliance and risk management. Chief information and information security officers take the lead on responsible practices, such as cybersecurity, privacy and performance. Data scientists and business domain specialists apply responsible core practices as they develop use cases, formulate problems and prompts and validate and monitor outputs.

Salesforce has a holistic AI strategy built into everything they do. Their **Einstein 1 platform** prioritizes data masking, inputting clean data, and secure data retrieval. They built an internal ethical AI policy team to confirm their people can stay abreast of the latest regulations.

Susan Emerson, Senior Vice President – Product: AI, Analytics & Data, Salesforce, notes that embedding responsible AI in organizational values lays the foundation. “It very much is part of our values. Our culture is driven all the way from the top and throughout the organization. Trusted AI is really one of the core elements that Salesforce has been focusing on.”

Discover how companies can unlock the value in AI while mitigating the downsides: **“From principles to practice: Responsible AI in action”**

“Trust is not just a layer; it’s a road map. Build that muscle of trusted, ethical AI now, because the road map is going to become more and more complex as we go toward autonomous AI.”

Marc Mathieu, Head of AI Transformation, Salesforce

Source: strategy + business, **From principles to practice: Responsible AI in action**, February 20, 2024



Rachel Botsman, a leading author on trust and technology, summarized the insights from the dialogue at Davos well when she said the following:

Throughout the week, I jotted down the tensions I kept hearing leaders talk about:

- *The confidence to embrace innovation but also hold back to ensure safety.*
- *The need to augment people's capabilities but not replace humans.*
- *The need to regulate and set guard rails but not box people in.*
- *The need to maximize benefits in the hands of 'good actors' whilst minimizing the harm of 'bad actors.'*

From "Growth and Jobs" to "Cooperation and Security" to "Climate Nature Energy", two key themes kept emerging: trust and uncertainty. Indeed, the official theme was "rebuilding trust." But (and it's a big but) where will the deep sources of trust come from? We have a long way to go.

"My trust takeaways from the World Economic Forum"

“For me, the week always feels like a pulse-check of leading decision-makers. The overall feeling is as revealing as what people are saying...”

Rachel Botsman

The risk vs. trust mindset

The risk mindset is about moving fast, seizing opportunities, and mitigating potential pitfalls along the way. It's fueled by the fear of missing out, no matter the cost.

In our conversation, Rachel shared a framework called The Trust Mindset vs The Risk Mindset. The trust mindset is about building relationships, being uncomfortable with uncertainty, constantly thinking about the trust state of all stakeholders and how it is possible for both to coexist.

The World Economic Forum dialogue presented an opportunity to challenge executives to identify which way they leaned, stressing that it is possible for both mindsets to co-exist in the context of an organization's unique environment and priorities.

Risk mindset

In a risk mindset, the primary focus is on identifying and mitigating potential negative outcomes associated with AI adoption.

Executives operating from a risk mindset are concerned with factors such as data security, privacy breaches, algorithmic bias, regulatory compliance, and the potential for job displacement.

Decision-making tends to be cautious, with a strong emphasis on risk assessment and risk management strategies.

There may be a tendency to prioritize short-term gains over long-term considerations, leading to a focus on immediate benefits while downplaying potential risks and ethical concerns.

Trust mindset

In a trust mindset, the emphasis shifts towards building and maintaining trust in AI systems and the organizations deploying them.

Executives operating from a trust mindset prioritize factors such as transparency, accountability, fairness, and ethical considerations.

Decision-making is guided by a commitment to fostering trust among stakeholders, including customers, employees, regulators, and the broader community.

There is a recognition that trust is a critical asset that should be earned over time through consistent and responsible behavior.

While still cognizant of potential risks, executives in a trust mindset place greater emphasis on building trust through actions such as implementing strong ethical guidelines, engaging stakeholders in meaningful dialogue, and prioritizing the long-term impact of AI deployments on society.

Designing responsible AI strategies with trust from the beginning can help balance risk while enabling your organization to move at speed.

Open-minded, agile leadership also is critical: risk-minded leaders deliver better, faster guidance as they internalize the momentous significance of the generative AI revolution. Opportunity-seekers are well-served by spending time immersing themselves in what can go wrong to avoid costly mistakes. And both groups need a healthy dose of appreciation for the priorities and concerns of the other.⁵

“We should become the architects of a balanced human-machine world. This requires harnessing the power of technology to solve complex problems, to create new opportunities, and to improve the quality of life, while always keeping the human element at the forefront.”

Scott Likens



What now: human prompts and the new circle of trust

We encourage leaders to assess risk and trust as they endeavor to implement AI and generative AI. We encourage you to do the same by leading an exercise with your leadership team. Create a space to incubate ideas and have transparent, authentic conversations about the impact of AI within their respective lines of business.

Treat this closed-door activity as a collaborative ‘open mic’ style exchange to explore ideas and be vulnerable as a collective. Remember to involve individuals from diverse backgrounds who hold different roles.

Here are a few prompts to get the conversation started:

Perspective

- Do you trust technology?
- What do you think the AI trust gap is? What are some ways to bridge it through education?
- Does increased use of AI correlate to an increased understanding of AI?
- Do you think about AI from a place of risk or trust?
- Identify pros and cons of adoption from each of those mindsets.
- Identify which of those risks are real threats vs. perceived threats.
- What would good and bad actors do with AI?
- How can we embrace innovation while keeping our data safe?

Process

- What are the desired impacts of AI adoption within your daily responsibilities? For your team? For the company?
- How are you already integrating AI into your work?
- What does your ideal AI adoption look like? What risks come with that?
- How can we innovate within the confines of regulations?

People

- How do you bring your organization along with you?
- How can you use interpersonal company relationships to build trust?
- How can we upskill our employees without replacing them?
- How can each person on your team better utilize their unique skill set?
- How do you create a safe space for them?
- How do you identify how high you’re asking them to jump?
- Who should verify the results of AI and when?

Ready to embrace AI and transform your business?

Explore the [AI Strategy Guide](#) with resources on how to plan your approach, ready your technology and enable your people.

PwC's AI solutions, powered by Salesforce, keep humans at the center of innovation. We bring companies and customers closer together in trusted, meaningful ways. Our perspective on AI centers provides human oversight, enhanced security measures, and trusted customer data to power better customer experiences.

Contacts



Scott Likens
Principal and Chief AI Engineering
PwC Global and US



Susan Emerson
Senior Vice President – Product: AI,
Analytics & Data, Salesforce

