

APRIL 2026

# THE WHARTON BLUEPRINT FOR AI AGENT ADOPTION

The Psychological Frictions That Slow Adoption  
and The Scientific Evidence on How to Overcome Them



Human-AI  
Research



# Authors



## Thomas McKinlay

Founder of Science Says

Thomas is the founder of [Science Says](#), a consultancy and media company that gives businesses insights from the latest scientific research. An ex-Googler, Thomas was frustrated by how far removed business practices were from the latest evidence. Today, Science Says helps thousands of companies—from small businesses to Fortune 500 firms—apply the latest research, and successfully adopt AI in their organizations.



## Stefano Puntoni

Professor of Marketing and Faculty Co-Director of Wharton Human-AI Research, The Wharton School, University of Pennsylvania

Stefano is the Sebastian S. Kresge Professor of Marketing at The Wharton School and Faculty Co-Director of Wharton Human-AI Research. Prior to joining the University of Pennsylvania, Stefano was a professor of marketing and head of department at the Rotterdam School of Management, Erasmus University, in the Netherlands. He holds a PhD in marketing from London Business School and a degree in Statistics and Economics from the University of Padova, in his native Italy.

## Research & Insights



## Serkan Saka

Assistant Professor of Marketing, San José State University  
Senior Research Analyst, Science Says

# CONTENTS

## 01 Authors

## 03 Contributions

## 04 Introduction

## 05 The Problem with AI Agent Adoption

## 11 Psychological Friction 1: Perceived Competence

People Want Competent, not Warm AI Agents

People Focus on Four Elements to Determine whether the Agent is Adding Value

Explanations Make AI Agents Feel More Reliable

Presenting Agents as Supporting Human Experts Makes Them Feel More Competent

## 19 Psychological Friction 2: Trust

People Have Higher Trust in AI Agents When They Know Their Limitations

Showing Proof of Successful Outcomes Builds Trust in AI Agents

People Trust AI Agents that Reduce Uncertainty Before, During, and After Use

Feeling Understood Makes People Accept AI's Recommendations More

A "Learning" Label Makes People More Likely to Trust the Agent

When AI Makes the Process Too Easy, People Trust Outcomes Less

People Have Higher Trust in AI's Work When it Uses Precise Numbers

AI Loses Trust When It Feels Generic Instead of Tailored

## 30 Psychological Friction 3: Delegation of Control

People Delegate Best When Agents Have a Moderate Level of Autonomy, not Minimal or Total

If People Don't Feel in Control, They Won't Adopt Your AI Agent

AI Adoption Increases When People Feel Agents Are Their Own

People Rely More on Agents When Under Pressure to Move Fast

AI Can Reduce Concerns When It Feels Less Powerful Than a Human

The Higher the Delegation, the Less the Oversight

## 42 Summary of Research Findings and Recommendations

Psychological Friction 1: Perceived Competence

Psychological Friction 2: Trust

Psychological Friction 3: Delegation of Control

## 46 Stay Updated

# With Contributions From Industry Leaders and Wharton Faculty



**Kartik Hosanagar**

John C. Hower Professor of Technology and Digital Business, Professor of Marketing, and Faculty Co-Director of Wharton Human-AI Research



**Prasanna (Sonny) Tambe**

Professor of Operations, Information and Decisions, and Faculty Co-Director of Wharton Human-AI Research



**Hamsa Bastani**

Associate Professor of Operations, Information and Decisions, and Associate Professor of Statistics and Data Science



**Gérard Cachon**

Fred R. Sullivan Professor, Professor of Operations, Information and Decisions, Professor of Marketing, and Department Chair



**Lyle Ungar**

Professor of Computer and Information Science, and Professor of Operations, Information and Decisions



**Katherine Milkman**

James G. Dinan Endowed Professor, and Professor of Operations, Information and Decisions



**Shiri Melumad**

Associate Professor of Marketing



**Ethan Mollick**

Ralph J. Roberts Distinguished Faculty Scholar, Associate Professor of Management, and Faculty Director of Wharton Generative AI Labs



**Chris Caldwell**

CEO, Concentrix

Fortune 500 company deploying AI agents to enhance operations for 2,000+ clients



**Wade Foster**

CEO, Zapier

Platform allowing millions of its users to deploy AI agents across 8,000+ web apps



**Neil Hoyne**

Chief Strategist, Google

Deploying AI agents across products like Google Workspace, Search, and Chrome



**Brian Solis**

VP and Head of Global Innovation, ServiceNow

Running 80+ billion workflows per year on a platform unifying legacy systems, cloud applications, and AI agents



**Maria Montenegro**

CEO, Wolters Kluwer Corporate Performance & ESG

Introducing AI agents for finance and ESG reporting across their customers, including 88% of the Fortune 500



**Adam Seligman**

CTO, Workato

Enterprise MCP leader connecting AI agents to business data and apps for half the Fortune 500

# Brought to you by Science Says and Wharton Human-AI Research



## Science Says

Wharton has partnered with Science Says to bring you this Blueprint.

Science Says specializes in turning the latest academic research into practical, digestible insights.

For weekly insights in AI and marketing, you can subscribe for free to the [Science Says newsletter](#).

For research reports in your industry, as well as resources and support with AI adoption in your organization, you can find out more [here](#).

Human-AI  
Research

## Wharton Human-AI Research

Wharton Human-AI Research advances the exploration of how artificial intelligence is transforming how we live, learn, and work.

By bringing together top researchers and industry leaders, Wharton Human-AI Research provides insights that help organizations leverage AI for competitive advantage.

This Blueprint is a product of that mission, combining rigorous research with practical applications.

For more human-AI research insights, visit [ai.wharton.upenn.edu](https://ai.wharton.upenn.edu) or [subscribe to the newsletter](#).



# THE PROBLEM WITH AI AGENT ADOPTION

“The number one place I see folks struggle is that they are not sure where to start. What are the right use cases? How can they get impact? As a result, they keep doing what they have always done.”



**Wade Foster**  
CEO of Zapier

“When AI first got better, people had their first ‘wow moment’ - maybe they wrote an email or summarized a document - but then they gave it real work and it failed or made stuff up because it couldn’t access any real data. So now people are discouraged, there is skepticism baked into their perception of AI agents.”



**Adam Seligman**  
CTO of Workato

“Mostly, people do not trust AI to do the right thing. Often, they shouldn't trust it. But they should be trying AI agents more to see where they do work.”



**Lyle Ungar**

Professor of Computer and Information Science, and Professor of Operations, Information and Decisions

# The Problem with AI Agent Adoption

AI agents are here, and they work. It is already possible to use agentic AI not just to research a trip, but actually book it. Yet outside of early adopters and technology enthusiasts, most people are not ready to let an AI agent email their boss, manage their clients, access their personal files, or use their credit card.

Adoption is moving from a technological problem to a psychological one. Giving an AI agent control over fundamental parts of your life is a completely different proposition than just asking a question to a chatbot\*. It requires believing the agent can do the job, trusting it will do it right, and being willing to hand over control.

This Blueprint helps AI agent developers and deployers close this psychological gap. It draws on the latest behavioral science in human-AI interaction, recommendations from Wharton Human-AI Research faculty, and what's happening in organizations that are already deploying AI agents at scale.

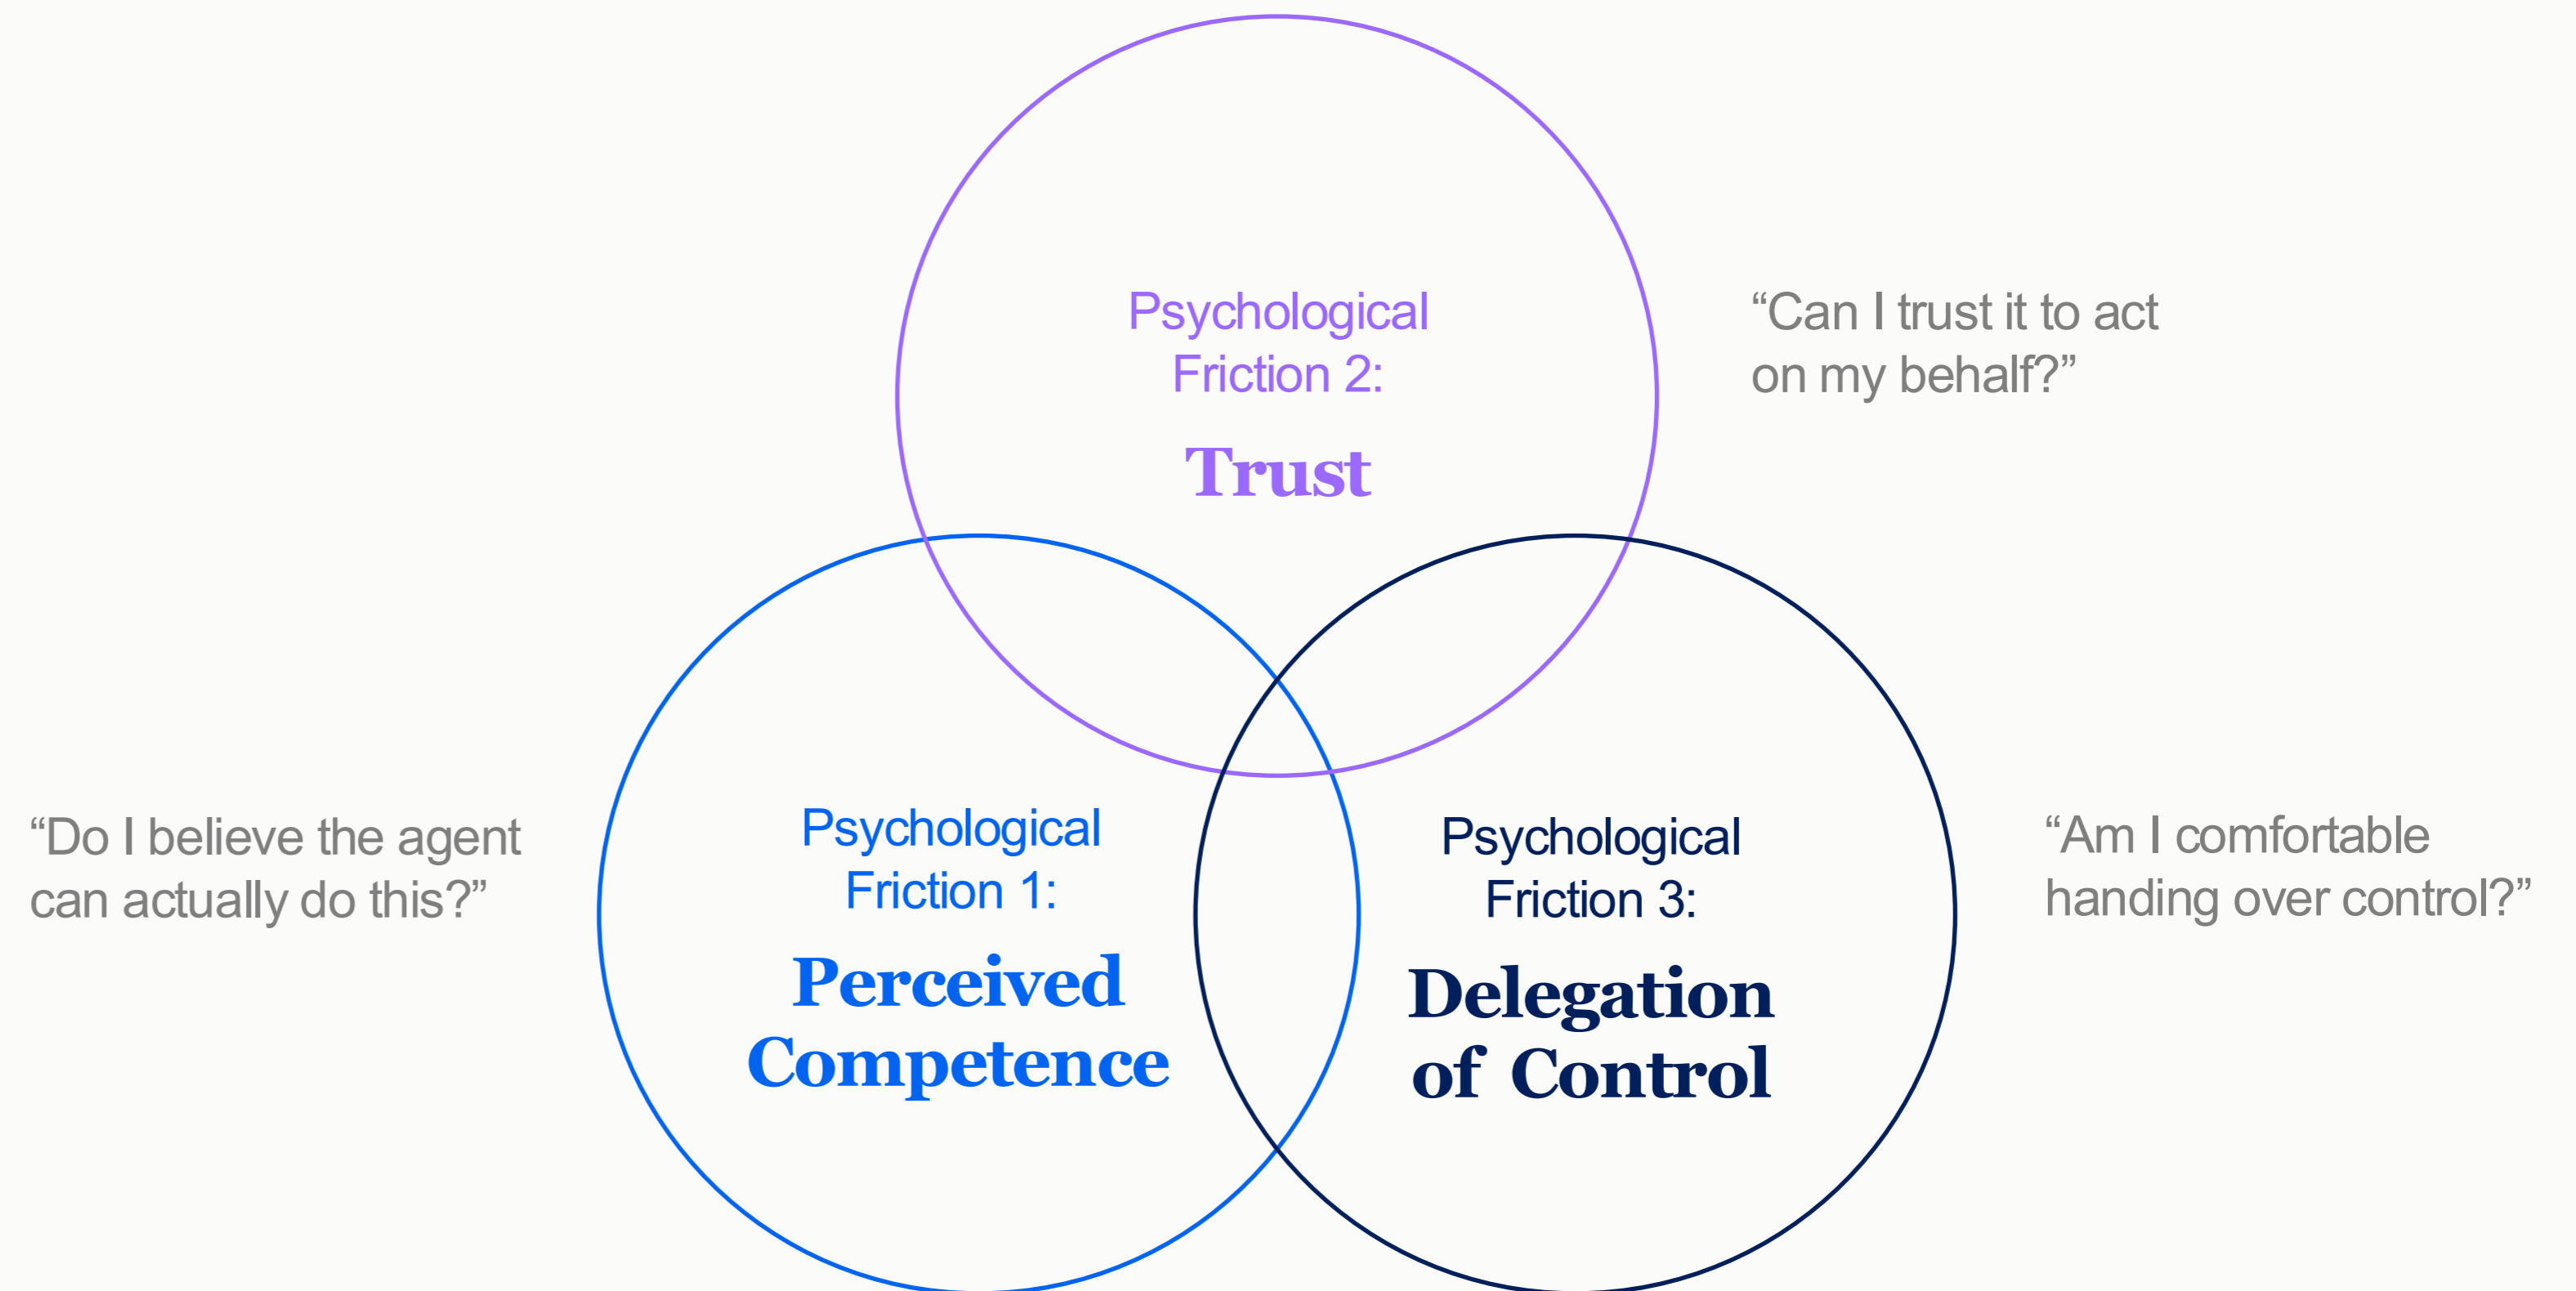
\*To optimize your AI chatbot, see our [Wharton Blueprint for Effective AI Chatbots](#)

# How to Overcome the Problem

We turned to companies leading the deployment of AI agents and asked their executives to share the challenges they're seeing on the ground.

We then answered those challenges with insights from the latest scientific evidence and recommendations from Wharton's leading researchers in the topic of human-AI research.

The findings are organized into three frictions you need to overcome to drive AI agent adoption:



## ■ Note on the Research


This Blueprint is based on the latest scientific evidence in human-AI interaction. The research we cover is our current best knowledge based on rigorous scientific research from leading universities and scientists around the world. Given the novelty of AI agent technology, we have extrapolated insights from AI and organizational behavior research that we believe also have implications for agents. While the insights presented here are solidly grounded, remember that this field is rapidly evolving, and future research may change our understanding of these insights.



# THE BLUEPRINT FOR AI AGENT ADOPTION

# Psychological Friction 1: Perceived Competence

Do I believe the agent can accomplish what I want it to?



Perceived competence is the user's subjective belief in an AI agent's ability to perform desired actions, rather than a measure of its technical capability. Elements of this friction include explaining AI agent decisions, communication style, and handling of mistakes.

“People are not as forgiving of AI errors as they are of human errors. They systematically misjudge AI partners by focusing on where the system fails on ‘easy’ cases. Those failures loom larger psychologically than the harder cases where the AI adds real value.”



**Kartik Hosanagar**

John C. Hower Professor of Technology and Digital Business, Professor of Marketing, and Faculty Co-Director of Wharton Human-AI Research

## PSYCHOLOGICAL FRICTION 1: PERCEIVED COMPETENCE

# People Want Competent, not Warm AI Agents

## Research Findings

- Across multiple experiments, people were less willing to use AI that was friendly and warm (e.g., sounds cheerful), compared to competent (e.g., cites criteria used or explains the reasoning).
- AI that seemed competent led to higher trust, which was a strong predictor of willingness to use the AI and share personal information. The higher the trust, the higher people's willingness.

## Practical Implications

- Ensure your agent signals competence through consistent recommendations, as inconsistency is seen as a sign of lower competence.
- When the agent makes a recommendation or takes an action, state the criteria it followed (e.g., "I prioritized location, cancellation flexibility, and reviews."). This reinforces the sense of competence and reasoning.
- In most situations (e.g., involving health, money, legal, and professional), avoid designing agents with an overly-friendly tone or warm personalities (e.g., avoid "Great choice! You're going to love this option!").

## What Drives This

- Users judge AI agents like people, asking "Can it do the job?" and "Can I rely on it to produce reliable outputs?"
- Feeling the AI can do the job increases perceptions of competence and fairness, which leads to trust when using it.

## Limitations

- The research focused on AI chatbots as advisors. We expect the effect to work similarly for AI agents taking real actions.

“Most people do not care if the AI feels empathetic. They care if it solves their problem.”



**Lyle Ungar**

Professor of Computer and Information Science, and  
Professor of Operations, Information and Decisions

“People can resist agents that don't match the tone of the consumer – customers tend to get frustrated with overly polite and obedient technology that isn't accomplishing things at speed.”



**Chris Caldwell**

CEO of Concentrix

## PSYCHOLOGICAL FRICTION 1: PERCEIVED COMPETENCE

# People Focus on Four Elements to Determine whether the Agent is Adding Value

## Research Findings

- Based on an online survey using scenario-based evaluations of AI-enabled travel agents, people had a higher intention to use an AI travel agent if they felt the agent was working alongside them and co-creating value together.
- People judged the agent's added value based on four main characteristics of convenience, personalization, ubiquity (always available), and superior functionality.
- However, privacy concerns, technology anxiety, usage barriers, and the desire for personal human interaction decreased intentions to use the AI.

## Practical Implications

- Strengthen reasons for use while actively reducing reasons against use.
- To strengthen reasons to use, clearly show how the agent is helping the user add value. For example, have the agent:
  - Narrow down options and explain why they'd be a good fit, but let the user add an opinion.
  - Show tradeoffs between different courses of action and let people adjust assumptions.
  - Remember preferences and adapt over time.
- Address privacy and anxiety concerns directly, before people even use the agent (e.g., clearly state your privacy policy in simple terms).

## What Drives This

- When deciding whether to use agents, people consider benefits (is it adding value?) and concerns (do I risk my privacy?) simultaneously.
- So even when people see clear benefits, unresolved concerns can still limit perceived value.

## Limitations

- The study focused on travel planning, which is an information-heavy but relatively low-risk context. The effect might differ for high-risk agents.

“Giving up tasks is not the issue; employees do not want to lose their identity or give up agency. If anything, they seek empowerment.”



**Brian Solis**

VP and Head of Global Innovation of ServiceNow

“Talk about use cases. How do you actually get value out of these? There’s a lot of jargon, hype, and tech lingo. What most folks really want to know is how’s this going to impact my day-to-day?”



**Wade Foster**

CEO of Zapier

To learn more about this see: [How AI Affects Our Sense of Self](#), *Harvard Business Review*, Stefano Puntoni & Gizem Yalcin, 2023

## PSYCHOLOGICAL FRICTION 1: PERCEIVED COMPETENCE

# Explanations Make AI Agents Feel More Reliable

## Research Findings

- In several experiments in high-stakes service contexts (e.g., healthcare and insurance scenarios), people felt the AI was more reliable and safe, and were more likely to accept its response, when it provided detailed explanations (e.g., explained the process to get the information asked or complete a task).
- When AI explained its process (e.g., steps taken, data considered, or how a task was completed), people perceived the system as more reliable and safer to use.

## Practical Implications

- Explain the process as well as the outcome. Show the calculation steps or data sources used to reach a recommendation.
- Automate detailed explanations. Display the process breakdown automatically, without requiring the user to click on a button to view.
- Frame explanations to provide reassurance and clarity, not just technical transparency (e.g., use “I used X methodology because it provides accurate data”, not “I used X methodology”).

## What Drives This

- People see detailed explanations as a signal of competence and seriousness, and interpret them as a trust signal of the AI’s work quality and safety.
- This is especially important in high-stakes decisions where users feel vulnerable (e.g., financial or medical decisions).

## Limitations

- The study looked at high-stakes services (e.g., insurance, healthcare); the effect might be smaller for low-stakes decisions (e.g., tourism), where competence is less important.
- Trust in AI declines among experts in its domain of application, limiting applicability in settings with highly experienced users (e.g., a legal AI targeted to lawyers).

“People do not want AI making silent decisions. They want transparency. What does the agent recommend? What does it execute? Where does human approval sit? That transparency builds trust quickly.”



**Maria Montenegro**

CEO of Wolters Kluwer Corporate Performance & ESG

“AI agents are too often built in the image of yesterday’s workflows. If we build agents that expose their decision logic in business terms people use, we gain agentic empathy, trust built through auditable reasoning. This way, people can understand them better.”



**Brian Solis**

VP and Head of Global Innovation of ServiceNow

## PSYCHOLOGICAL FRICTION 1: PERCEIVED COMPETENCE

# Presenting Agents as Supporting Human Experts Makes Them Feel More Competent

Associate AI agents with credible human figures (e.g., let people know an expert is working alongside the agent) when competence and credibility matter (e.g., medicine, finance) and position AI as a subordinate or support system (e.g., use “assists,” “supports,” “prepares,” not “co-decides”).

Adding a human doctor to an AI consultation reduced resistance to AI by 16% (vs. AI alone). The trust people have in experts is transferred onto other people working with them; the same happens with AI. However, if the AI is positioned as an equal, people feel like it is competing with the expert’s authority.

“Programmers are the only group that has already broken through the skepticism stage, because they’ve seen agents that deliver at a high bar. But for everyone else, they have not seen that level of impact yet.”




**Adam Seligman**

CTO of Workato

# Psychological Friction 2: Trust

Should I trust the agent with this task?



Trust is the user's willingness to rely on the AI agent despite uncertainty. It reflects a belief that the system will behave reliably, consistently, and in line with the user's interests. Elements such as showing successful outcomes, being transparent about limitations, and using communication styles that reinforce reliability all influence trust.

“People struggle with AI agents because they do not know when to trust them. The hesitation is not just fear of errors—it is also uncertainty about scope, reliability, and how much oversight they will need.”



**Prasanna (Sonny) Tambe**

Professor of Operations, Information and Decisions,  
and Faculty Co-Director of Wharton Human-AI Research

## PSYCHOLOGICAL FRICTION 2: TRUST

# People Have Higher Trust in AI Agents When They Know Their Limitations

## Research Findings

- Across two experiments using a dynamic task in which the AI had a known weakness, people who were clearly told where an AI system was likely to fail showed higher trust in using it and worked more effectively with it than people who were not told, even though the AI itself was identical.
- Being upfront about limitations did not reduce trust after people experienced using the AI. Once people interacted with the system, their confidence and reliance were shaped more by how well they understood its boundaries than by abstract concerns about imperfection.

## Practical Implications

- Do not hide your AI agent's limitations. Clearly explain where your agent is less reliable. For example, "This agent finds hotels in your area, but is not very good at checking amenities."
- Do not frame explanations as technical details (e.g., "This agent is not directly connected to the Google Maps MCP"); keep it simple and clear.

## What Drives This

- When people are told the limitations of the AI, they feel they understand how to use it better.
- This helps them select when to pay attention to the output (e.g., if they know that a specific task might be inaccurate), rather than feeling they have to pay attention all the time, which makes people feel more confident when using it.

## Limitations

- The AI weakness in this research was known in advance (e.g., people were told the system was less reliable for one particular car color because it had limited training data for that color), whereas real-life AI systems may have shifting or context-dependent limitations.

## PSYCHOLOGICAL FRICTION 2: TRUST

# Showing Proof of Successful Outcomes Builds Trust in AI Agents

## Research Findings

- Based on controlled experiments where people evaluated AI-generated recommendations in a speed-dating context, people trusted AI more when they saw evidence of its success in real-world tasks (e.g., “Booked 531 flights today!”) rather than detailed technical explanations about its decision-making process (e.g., “I’ve used an API to scrape websites”).
- Seeing evidence of the agent’s work (e.g., how many flights it successfully booked) increased people’s belief in the AI’s future accuracy by up to 22.1%.

## Practical Implications

- Highlight the accuracy of AI predictions or task performance through clear, measurable outcomes (e.g., “94% of users loved the products recommended by our AI”).
- Include mentions of outcomes and performance in your messaging and AI-user interactions when relevant (e.g., “I suggest this option, 89% of users say my recommendations are accurate”).
- Provide relatable examples of its performance (e.g., show the AI resolving recurring customer issues).
- Avoid overwhelming people with overly technical or complex explanations of how AI works unless absolutely necessary (e.g., for highly specialized products).

## What Drives This

- Many people are not interested in how an AI works. Rather, they want to know whether it works.
- Outcome proof reduces uncertainty by showing that the system has already performed well in similar situations. This makes the AI feel tested, reliable, and safe to rely on, even if people cannot explain its logic.

## Limitations

- The study looked at speed-dating AI recommendations; the effect might not be the same for all AI applications and when the user is an expert in a domain (e.g., a VP in marketing may ask, “What model assumptions drove this forecast?” rather than being persuaded by “Used by 2,000 brands”).

“You need hands-on time to develop a feel for what works, what breaks, and where the edges are. You have to tinker with this stuff to get a tactile feel for how it works. It’s like learning to fly, you just can’t learn it from reading a manual.”



**Adam Seligman**

CTO of Workato

## PSYCHOLOGICAL FRICTION 2: TRUST

# People Trust AI Agents that Reduce Uncertainty Before, During, and After Use

## Research Findings

- People experience uncertainty when using AI agents at different phases: before, during and after interactions.
- Transparency and autonomy reduce this uncertainty and increase trust in the agent (e.g., ensuring users have the final decision).

## Practical Implications

- Build trust by reducing uncertainty across the interaction:
  - Before action: Make goals and criteria explicit (e.g., “You said you wanted cancellation flexibility, so I’m excluding all non-refundable options.”). Clarify why additional information is needed (e.g., “I’m asking this to refine trade-offs.”).
  - During action: Make steps visible (e.g., “Booking flight...selecting refundable option...confirming cancellation policy.”). Add approval checkpoints (e.g., “Do you approve that a 4-hour layover is acceptable?”).
  - After action: Show how feedback changes future behavior (e.g., “I’ve noted you prefer earlier departures and will prioritize them going forward.”). Summarize trade-offs clearly.

## What Drives This

- People feel vulnerable when using AI agents because they can’t fully see or predict how agents make decisions, or whether outcomes will match their expectations.
- Trust depends on how willing users are to accept this vulnerability.
- Transparency (e.g., making the agent’s reasoning clear) and autonomy (e.g., granting users checkpoints and override options) reduce this uncertainty.

## Limitations

- This paper outlines design principles rather than testing specific product features in real-world deployments.

“Early on, an agent can feel like managing a promising, but still unreliable junior hire. You delegate something, but then spend twenty minutes checking whether it was done the right way. Then you fix it. Then you wonder why you didn’t just do it yourself in the first place. Trust is not built and you’re returning to your old patterns.”



**Neil Hoyne**

Chief Strategist of Google

## PSYCHOLOGICAL FRICTION 2: TRUST

# Feeling Understood Makes People Accept AI's Recommendations More

Make AI sound like it understands people for decisions that relate to long-term or virtue goals (e.g., “You mentioned you’re trying to eat healthier, so I prioritized lower-fat options.” Or, “I know you’re trying to improve your sleep, so I prioritized hotels with quieter locations and strong sleep quality reviews.”).

People accepted AI recommendations 54.2% more when the AI showed a strong understanding of people’s intentions and emotions, but only for long-term or virtue choices. Avoid using it for pleasure-driven or vice decisions (e.g., romantic restaurant choice).

When AI seems to understand people's intentions, they trust its advice for long-term or practical goals. However, for choices made for pleasure, the same advice feels intrusive and reduces trust.

**Research:** Liu, T., Xu, Y., Yang, J., & Li, K. (2025). [Mindful machines: understanding how AI's theory of mind capabilities influence consumer response to product recommendations](#). *Psychology & Marketing*.

**Universities:** Durham University, Southampton University

---

## A “Learning” Label Makes People More Likely to Trust the Agent

Add a "learning" or "constantly improving" label to your AI agent (e.g., "This AI adapts and improves with use"). In experiments, people chose AI labeled as "learning" 55% of the time, compared to 43% with no label. The label signals that the AI is dynamic and gains experience over time, and not static or outdated.

This makes people more confident in using it, increases trust, acceptance, and reliance, especially in high-stakes contexts like financial or healthcare advice.

Pair the label with a brief explanation of how it improves (e.g., "This agent improves based on your feedback"). Avoid presenting AI as fixed or outdated (e.g., "Last updated in 2025"). This hurts confidence in reliability and adaptability.

**Research:** Chacon, A., Kausel, E. E., Reyes, T., & Trautmann, S. (2025). [Preventing Algorithm Aversion: People Are Willing to Use Algorithms with a Learning Label](#). *Journal of Business Research*.

**Universities:** Universidad Técnica Federico Santa María, Pontificia Universidad Católica de Chile, University of Heidelberg

## PSYCHOLOGICAL FRICTION 2: TRUST

# When AI Makes the Process Too Easy, People Trust Outcomes Less

When AI performs all the work upfront, with no input from users, people spend 21.3% less time engaging and feel 5.4% less psychological ownership over the outcome (e.g., not feeling like the solution is theirs or that it does not reflect their own thinking). That produces advice that people are less likely to adopt.

To support adoption, design agents so they share the work before acting (e.g., narrow options, surface tradeoffs, involve users in small decisions). Increased participation in the process helps users feel invested and makes AI-supported outcomes more defensible and easier to rely on.

“While learning through web search requires us to navigate web links, interpret different sources and synthesize them ourselves, LLMs conduct this process on our behalf. This ease is an obvious driver of AI adoption, but it can come at the cost of developing deeper knowledge on topics by transforming learning from an active to a passive process.”

**Shiri Melumad**

Associate Professor of Marketing

## PSYCHOLOGICAL FRICTION 2: TRUST

# People Have Higher Trust in AI's Work When it Uses Precise Numbers

## Research Findings

- Across multiple controlled experiments, people trusted AI recommendations 12% more when numbers were precise (e.g., 8.2 out of 10) rather than rounded (e.g., 8 out of 10).
- This was especially effective when AI made practical, functional recommendations. People responded 32.7% more positively to precise numbers and were more likely to purchase the recommendation.
- When AI showed precise numbers, people were also 15.7% more willing to act on the recommendations compared to rounded numbers.
- The effect was strongest under high uncertainty (e.g., people were unsure how accurate the AI was). Precise numbers drove higher trust when AI accuracy was unclear or when product quality was difficult to judge (e.g., the product is totally new, and users were unfamiliar with the product).

## Practical Implications

- Display precise numbers when your AI makes analytical or performance-based claims (e.g., “8.2/10 match”, not “8/10 match”).
- Use precision in early interactions when people are deciding whether to trust the agent at all, and when trying to increase delegation to AI.
- Do not use precision to mask poor recommendations; it will break trust.

## What Drives This

- People expect AI to be analytical. Precise numbers signal rigor and confidence, matching these expectations.
- Rounded numbers feel human and approximate, clashing with how AI is thought to operate.
- When AI communicates with numerical precision, people think that the system is grounded in analysis rather than guesswork. That perception of competence increases trust and makes people more willing to rely on the agent.

## Limitations

- Most scenarios were hypothetical rather than based on real purchase behavior, so effects may differ when decisions involve real money or long-term consequences.

## PSYCHOLOGICAL FRICTION 2: TRUST

# AI Loses Trust When It Feels Generic Instead of Tailored

## Research Findings

- Across multiple experiments, people were 8% less willing to accept AI recommendations and outcomes when they believed the advice was generic (e.g., “most users chose this flight” or “this is our most popular savings strategy”), rather than tailored to them (e.g., “based on your preference for direct flights, I booked this option”).
- When people believed preferences varied widely across consumers (e.g., people thought that different people like different things), they were 6.2% less willing to accept AI recommendations.

## Practical Implications

- Make the agent reference the user’s criteria when carrying out tasks (e.g., “You selected low-risk investments, so I excluded higher-volatility options.”).
- Be careful when using social proof and avoid sentences that refer to other users’ preferences and actions (e.g., “Most users choose...”).

## What Drives This

- When people think preferences differ substantially from person to person, they doubt that a generic AI (e.g., based on general consensus, and not the user’s criteria) can understand what matters to them. That doubt targets the AI’s judgment and decision-making ability.
- If an AI agent appears to be optimizing for the average user rather than the specific individual, users hesitate to let it act on their behalf.

## Limitations

- Most experiments were conducted in e-commerce contexts.

“Users disengage when outputs feel generic or ungrounded. To drive real productivity and decision quality, agents must be context-aware and informed by deep domain expertise and authoritative content.”




**Maria Montenegro**

CEO of Wolters Kluwer Corporate Performance & ESG

# Psychological Friction 3: Delegation of Control

How and how much should I delegate to the agent?



Delegation is the user's willingness to grant an AI agent the autonomy and control required to act on their behalf. Elements include the AI agent's level of control, the user's feeling of ownership, and the perceived efficiency of the agent's actions.

“The hard part of adopting AI agents is deciding what to delegate. When people do not have a clear sense of what the AI should do, what they should do, and when to step in, adoption stalls. Agentic AI works best when delegation is explicit, not implicit.”



**Ethan Mollick**

Ralph J. Roberts Distinguished Faculty Scholar, Associate Professor of Management, and Faculty Director of Wharton Generative AI Labs

## PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# People Delegate Best When Agents Have a Moderate Level of Autonomy, not Minimal or Total

## Research Findings

- Across multiple controlled experiments, people were more likely to adopt AI agents that had a moderate level of decision autonomy (e.g., they can suggest a shortlist of best hotel options independently, but do not book autonomously until told to do so).
- Moderate decision autonomy increased willingness to buy, whereas too little or too much of it decreased it.

## Practical Implications

- Design AI agents as co-deciders, not fully autonomous nor passive tools.
- Avoid full automation (e.g., selecting and booking hotels without human input) unless people explicitly opt in.
- Build features that make people feel like they're in control (e.g., adjustable settings, easy overrides).

## What Drives This

- When AI agents fully take over decisions (high autonomy), people feel their freedom and control are being taken away, which leads us to respond negatively to the agent.
- Similarly, when agents have too little autonomy (e.g., they need very specific guidance for each task), people feel the agent is unhelpful and requires too much work.

## Limitations

- The experiments focused on purchase decisions only; it is unclear how the effect impacts long-term use or repeated delegation to AI agents.
- Scenarios were experimental (news platforms, shopping guides, home AI robots), people might respond differently in real-life scenarios.

“Our research suggests adoption improves when agents mirror organizational design principles: specialization, hierarchy when needed, and explicit handoffs. Paradoxically, adding structure—rather than intelligence—can do more to unlock usage.”



**Prasanna (Sonny) Tambe**

Professor of Operations, Information and Decisions,  
and Faculty Co-Director of Wharton Human-AI Research

“Agents require delegation. You need to break down what you do into steps clear enough for someone else to follow. While people are pretty good at their jobs (most of the time), they can be terrible at explaining how they do them.”



**Neil Hoyne**

Chief Strategist of Google

## PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# If People Don't Feel in Control, They Won't Adopt Your AI Agent

Let people know that they are always in control of the agent. In your messaging, clearly state that all agent recommendations and actions are easy to review and track.

Based on three experiments with over 3,000 people, researchers found that control concerns (e.g., people worried they couldn't audit the agent's process) accounted for 26% of the weight of the decision whether to adopt AI. 31% of it was linked to privacy concerns (e.g., people didn't trust the agent with personal information), a similar weight to price.

**Chris Caldwell**

CEO of Concentrix

“AI agents should know when to escalate to humans and provide contextual handoffs that are structured to reinforce human decision-making rather than competing with it.”

## PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# AI Adoption Increases When People Feel Agents Are Their Own

## Research Findings

- Across experiments and a survey, people were more likely to adopt AI assistants when they felt a sense of ownership over them. People who actually owned an AI assistant (e.g., they purchased it and named it) showed 15% higher intention to use it than those who did not.
- Creating a psychological sense of ownership (e.g., by letting people name the AI) increased adoption by up to 20%.
- Ownership mattered only for functional AI (e.g., a scheduling agent) and task-focused AI (e.g., drafting a report). For companion AI (e.g., AI friends), people feeling like they owned the agent did not change their willingness to use it.

## Practical Implications

- Create a sense of ownership for functional AI agents (e.g., carrying out technical and admin tasks like accounting agents) using language like “your assistant,” “your setup,” or “your preferences.”
- Encourage light personalization early: naming, preferences, routines, or setup choices.

## What Drives This

- People are more committed if they feel they own the agent.
- This makes people more willing to invest time, effort, and attention into using the AI.
- However, this does not increase willingness to adopt the AI if the reason to use the agent is an emotional connection, as in the case of companion AI.

## Limitations

- Effects were tested primarily with smart-speaker style AI assistants.

“Design agents like employees you’d like to hire and manage. Give them a job description, permissions, goals, escalation paths, and measurable outcomes. Onboard them like you would a high-performance candidate, and nurture them to thrive.”



**Brian Solis**

VP and Head of Global Innovation of ServiceNow

“Companies need to be able to say, ‘you can draft this email, but you can’t send it,’ or ‘you can recommend inventory moves, but you can’t execute them.’ Without clear controls, agents stay stuck on trivial tasks because nobody trusts them with anything important.”



**Adam Seligman**

CTO of Workato

## PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# People Rely More on Agents When Under Pressure to Move Fast

## Research Findings

- In a lab experiment involving a task where people made decisions under different workload conditions, when decision pressure increased (e.g., the stakes are high, the volume of tasks is large, or the timing is short), people relied on AI more often.
- People followed an AI algorithm for 48% of decisions under high pressure, compared to 39% under lower pressure (e.g., fewer customers, fewer items to process).
- The difference wasn't about liking AI more. It was about speed. Under pressure, people stopped re-evaluating every recommendation and defaulted to the algorithm more quickly.

## Practical Implications

- Make the AI the path of least resistance by saving prior preferences, applying them, and allowing one-click approvals.
- Do not manufacture urgency or add unnecessary time pressure, but use real workflow pressures (e.g., busy inbox labels, high ticket volume trackers, deadline reminders).

## What Drives This

- People face two decisions when working with AI: whether to trust it and how much effort to spend deciding whether to trust it each time.
- Under low pressure, people tend to deliberate repeatedly, which adds cognitive friction and slows adoption. Under high pressure, people are more willing to default to the algorithm, reducing mental effort and allowing the AI's advantages to show.

## Limitations

- The study examined reliance behavior in a controlled task environment rather than long-term consumer adoption of autonomous agents.

“Some people underuse AI not because they doubt its value, but because remembering to use it at the right moment is difficult when they aren’t yet in the habit. One solution to this is engaging in more deliberate planning. For example, you might form the plan: Whenever I draft a document longer than five sentences, I will run it by an LLM for feedback.”



**Katherine Milkman**

James G. Dinan Endowed Professor, and Professor  
of Operations, Information and Decisions

## PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# AI Can Reduce Concerns When It Feels Less Powerful Than a Human

## Research Findings

- Across multiple experiments, people felt 13.8% lower risks (e.g., privacy concerns) when interacting with an AI agent compared to a human agent.
- When AI seemed less powerful (e.g., limited authority on decisions, ability to influence outcomes), people shared both highly sensitive and less sensitive information more willingly and showed a higher intention to reuse.
- When people were explicitly reminded of AI agents' power (e.g., told the agent would make decisions on their behalf without input), the effect reversed, and people showed greater concerns about giving up control to the agent (vs. humans).

## Practical Implications

- Frame AI agents as limited in power and authority using language like "assists," "supports," or "helps," not "decides," "evaluates," or "controls."
- Avoid emphasizing AI capabilities or decision-making authority (e.g., do not use "Our AI can take this over for you") when collecting sensitive information.

## What Drives This

- People worry most about sharing information with actors they perceive as judgmental, authoritative, capable of influencing outcomes, and able to misuse information (e.g., evaluating your financial profile and determining less-than-ideal options for you).
- Humans naturally trigger these concerns because they are seen as intentional, evaluative, and socially powerful. By contrast, AI agents often feel more neutral and less judgmental as long as they are framed as tools rather than authorities.
- When an AI is positioned as "just helping," people feel safer disclosing information. When it is positioned as powerful or controlling, it triggers the same (or even stronger) privacy anxiety as a human would.

## Limitations

- The study focuses on retail and service contexts (e.g., florist, home décor, food, and pharmacy). The effect might differ for high-risk domains such as healthcare or law.

”Let the agent assist, explain, and adapt before it acts independently. Autonomy should be earned rather than assumed.”



**Kartik Hosanagar**

John C. Hower Professor of Technology and Digital Business, Professor of Marketing, and Faculty Co-Director of Wharton Human-AI Research

“Agents need human managers. They will not behave like Agent Smith in the Matrix and take over your entire enterprise. Human oversight, governance, and training are essential in managing and collaborating with agents.”



**Brian Solis**

VP and Head of Global Innovation of ServiceNow

PSYCHOLOGICAL FRICTION 3: DELEGATION OF CONTROL

# The Higher the Delegation, the Less the Oversight

Design your agent so that people know when it is likely to make mistakes and in a way that keeps people attentive (e.g., include pop-up alerts where outputs have high stakes, ask for confirmations for the output, and force humans to review the process).

As agents improve and commit fewer mistakes, people see them as more reliable and competent. However, this can also make people pay less attention to potential mistakes, which leads to scaling costs. At some point, that cost outweighs the benefit of collaboration.

“If the AI is right 99% of the time, the human-in-the-loop has almost zero incentive to inspect the outputs. To force them to stay vigilant against that rare mistake, the employer has to pay a massive wage premium—one that scales inversely with the error probability.”



## Hamsa Bastani

Associate Professor of Operations, Information and Decisions, and Associate Professor of Statistics and Data Science



## Gérard Cachon

Fred R. Sullivan Professor, Professor of Operations, Information and Decisions, Professor of Marketing, and Department Chair

“When systems feel too automated, people lose a sense of agency. Collaboration models in which humans merely confirm AI output reduce engagement because users are no longer thinking, just rubber-stamping.”



**Kartik Hosanagar**

John C. Hower Professor of Technology and Digital Business, Professor of Marketing, and Faculty Co-Director of Wharton Human-AI Research



# **SUMMARY OF RESEARCH FINDINGS AND RECOMMENDATIONS**

# Psychological Friction 1: Perceived Competence

## Recommendation

Avoid giving your AI agent an overly-friendly or warm personality.

Clearly show how the agent is helping the user co-create value together.

Show detailed explanations for the agent's process.

Pair AI agents with a credible human professional and position the AI as supporting the human expert.

## Effect

→ Users judge agents based on their competence in task completion, not their niceness.

→ People are more likely to use an agent when they feel they are working alongside it. Highlighting benefits like convenience, personalization, superior functionality, and 24/7 availability helps overcome privacy or technology anxiety.

→ AI providing detailed steps signals competence and seriousness, making people perceive the AI system as more reliable and safer to use, especially in high-stakes decisions.

→ People's trust in an expert is transferred to the AI, which can reduce resistance to AI by 16%. Positioning the AI as an equal can compete with expert authority.

# Psychological Friction 2: Trust

Recommendation	Effect
Clearly define the agent's limitations, and highlight where it fails.	→ People trust agents more when they are transparent about capabilities, which helps them work more confidently and know when to step in.
Show proof of successful outcomes carried out by the agent.	→ Evidence of past success (e.g., "500 flights booked today") can increase belief in the AI's accuracy by 22.1%. This is more effective than technical explanations for most users.
Build trust by reducing uncertainty before, during, and after action.	→ People feel vulnerable when using agents at different stages, before, during, and after use. Decreasing uncertainty makes people feel less vulnerable and increases trust.
When the agent helps people with tasks related to long-term goals (e.g., eating healthier for better aging), make the AI sound like it understands their intentions and goals.	→ People accepted recommendations 54.2% more when AI agents showed a strong understanding of their intentions and emotions.
Use the label "learning" or "improving" to describe your agent.	→ People chose an AI labeled as "learning" 55% of the time vs. 43% without the label. It signals that the system is dynamic rather than static or outdated.
Involve people in tasks and decisions, and avoid full automation with no human input.	→ When the AI makes things too easy, people trust outcomes less because they incorrectly attribute the ease to an oversimplified process rather than competence.
Use precise numbers and metrics in all of the agent's outputs, even if rounded figures seem easier.	→ Precise numbers signal rigor and analytical depth. This can increase trust in recommendations by 12%, especially when the user is uncertain.
Ensure the agent references specific user criteria, not general advice.	→ Generic advice ("most users like this") reduces acceptance by 8%. Framing actions as tailored to the individual prevents the "it wouldn't work for me" bias.

# Psychological Friction 3: Delegation of Control

Recommendation	Effect
Design the agent for a moderate, 'human-in-the-loop' mode that requires some quick checks or final approval.	→ People delegate best with moderate control because they want to remain in a monitoring role to feel a sense of control and validate the outcome.
Clearly communicate the control the user has, showing where they can edit, pause, stop, or reverse the agent's actions.	→ Control concerns account for 26% of the decision to adopt AI. Users must see that they can easily review, track, and reverse actions.
Make people personalize the agent's name, settings, and functions to make it feel like "their" tool.	→ Letting users name the agent or personalize settings can increase adoption by 20%. This applies specifically to functional and task-focused agents.
Leverage workflow pressure and speed. For example, highlight full inboxes and upcoming deadlines to encourage agent use.	→ Under high pressure (high volume/short time), people follow AI recommendations 48% of the time vs. 39% under low pressure. Speed becomes the primary driver of reliance.
Frame agents as "less powerful" (e.g., influential) than humans.	→ People feel 13.8% lower privacy risk with AI compared to humans because they perceive the AI as less judgmental. However, if the AI is framed as powerful/controlling, this benefit disappears.
Design your agent to keep people attentive, especially in high-stakes outputs.	→ As agents become more reliable, people pay less attention to potential mistakes, which leads to scaling costs that can outweigh the benefits of collaboration.

# STAY UPDATED

Thank you for reading the Blueprint for AI Agent Adoption by Wharton Human-AI Research and Science Says.

We hope you find this Blueprint as useful and actionable as we did while creating it.

If this was helpful, you may also like to:

- [Sign up](#) to receive the latest insights from Wharton Human-AI Research
- [Get regular research updates](#) from Knowledge at Wharton
- For weekly insights from the latest research, subscribe to the [Science Says newsletter](#)
- Create a research blueprint for your own industry. [Reach out to Science Says here](#)
- [Reach out here to Thomas](#) for support with adoption of AI in your organization
- Read our [Wharton Blueprint for Effective AI Chatbots](#) to optimize your chatbot