

Artificial Intelligence and Business Model Uncertainty

Building the Digital Twin of Your Organization

Narrative script and presentation companion. The live demonstration and audience responses vary by session.

Good morning. I am Rico Kerstan, Managing Partner of KR Krisensicher Risikoberatung. We work with organizations that need to make consequential decisions under uncertainty. Today, the uncertainty is artificial intelligence - and what it does to the way your organization creates and defends value.

Before we begin, I think we need a different way of looking at this. The familiar template is falling apart. It is a visual metaphor for habits that no longer fit the environment. The question is whether our organizations can change those habits as quickly as their environment changes.

Your organization. Under pressure.

When we talk about AI, we tend to ask tool questions. Which model should we buy? Which process should we automate? What productivity gain can we expect? Those are useful questions, but they start too far downstream. I want to start with the business model: how do customers find us, what do they pay us for, and which parts of that journey can someone else now deliver differently?

The central question is whether two capable operators, armed with frontier AI, could take one valuable part of your customer journey. They do not need to manufacture every physical component. They might own the first conversation, the proposal, the recommendation, or the expert routine. The organization may keep its machines while losing influence over the work that connects those machines to customer value.

Evidence and further reading

Rico Kerstan & Miriam Büxenstein - AI Management: Organization in Agentic Times. Owner-provided manuscript; chapters 3, 11 and 20.

How much of your day?

How much of your working day involves AI? Scan the QR code and use the slider, from zero to eight-plus hours. No account, no exam and no leaderboard. We mean active use during your working day, not how long a browser tab is open.

The distribution is appearing live. More hours do not necessarily mean better results. A capable workflow may reduce the time we personally spend interacting with AI. This question simply establishes where this room is starting.

And without you?

If you closed your laptop now, would any work continue? What share of your work gets done by AI without you guiding every step? At one end: nothing moves without me. At the other: it would ask why I am interrupting.

This is a self-report, not a measurement of organizational automation. But it reveals a different operating logic: an assistant that waits for every instruction versus a bounded task that can genuinely be delegated. More autonomy needs a clear objective, suitable tools, limits and an accountable owner.

Have you let an agent act?

Have you given an agent a task, tools and permission to act? Not yet, tried it, or use it regularly? Asking a chatbot to draft a paragraph is different from allowing a system to take several steps towards an outcome.

We are finding a starting point, not ranking people. Systems differ and the label agent is used loosely. Let us look at the consequences of putting these capabilities into a workflow.

Now close your laptops.

These laptops represent the room. The moving subset illustrates the proportion of respondents reporting some autonomous work. It is not an exact measure of hours, jobs or productivity. Some work keeps moving; some waits for the next human input.

We should not infer that autonomy is universal or reliable. Useful progress depends on what happens when the system meets an exception. The management task includes specifying where it may act, when it should stop, and how outcomes are checked.

A different clock.

These curves illustrate linear planning meeting accelerated technical change. They are not a measured forecast of intelligence, quality, total task cost or a universal exponential law. The gap is the point: assumptions can become obsolete during the life of a perfectly coherent plan.

A firm can make a perfectly coherent annual plan and still misjudge what becomes possible during its execution. Capability, inference price, integration cost and adoption do not move at the same rate. A lower model price does not automatically create a lower end-to-end cost. Organizations have to keep testing the assumptions behind their roadmap.

The problem is not that planning is obsolete. It is that a plan needs feedback and revision. If it takes us a year to understand and approve an experiment, a smaller competitor may run many learning cycles while we are still discussing requirements.

Not on your competitor list.

Imagine two capable operators with frontier AI. They can combine coding, research, analysis and orchestration in ways that previously required a much larger team. This does not mean any two-person team can replace a regulated manufacturer. It means your competitor list may leave out the people who can now challenge a valuable part of your work.

Ask: which fragment of your customer journey is visible, repeatable, knowledge-intensive or profitable? A proposal, a design conversation, market intelligence, service triage, compliance preparation. The physical and regulated parts remain important. But someone else may get between you and the customer before a production order reaches you.

Give us a website.

Give us a public organization or brand website. Each device can enter one URL and replace it. No private portals or confidential links. We will choose one eligible entry at random, and the slot machine will display the server-selected result.

If no entries arrive, we can openly choose a public site. The point is to test a real workflow on a source we did not prepare around. A site can also fail to load or offer no compatible logo; that is a limitation we should make visible.

From URL to possibility.

Here is Project Zero Factory. From the chosen website, we retrieve public identity assets and select the genuine logo. A model proposes titles and rationales for three directions. The existing composition engine then combines that identity with card geometry and visual styles. We are assembling capabilities into a workflow; we are not claiming that a model drew every line.

Retrieving public information, identifying reusable assets, structuring design decisions, rendering alternatives and creating review material used to involve several steps and handoffs. Some of those steps can now run together. The exact website, asset format and provider conditions still matter.

These are nonfunctional concepts and proof separations, not bank-approved or printer-approved production files. If the live run fails, a prepared example can be shown with a clear label. A prepared result is not evidence that the randomly selected website worked.

They only need a valuable part.

Consider the customer journey: discover, understand, design, quote, approve, produce and deliver. The demonstration touched discovery and design preparation. It did not complete a commercial quote or production approval. Yet someone can challenge a valuable early part without reproducing the entire company.

Ask where customer trust, margin and proprietary knowledge sit. A small service can capture the first conversation and leave an established firm with fulfilment. Conversely, the established firm can combine its physical capability, relationships and domain knowledge with a faster front end. The threat and the opportunity are organizational choices.

Bridge to AI Management: the tool demonstration establishes a bounded technical capability. It does not establish an operating model. The real questions begin at handoff, review, exception, intervention and learning.

Evidence and further reading

Rico Kerstan & Miriam Büxenstein - AI Management: Organization in Agentic Times. Owner-provided manuscript; chapters 3, 11 and 20.

What if it comes from elsewhere?

OpenAI published a large collection of model-generated mathematical manuscripts on 6 October 2026. The repository distinguishes manuscripts from result families and describes varying verification status. Hundreds of manuscripts do not mean hundreds of independently verified cryptographic breaks.

Quantum computing has long been a prominent concern in cryptography. AI-assisted mathematical research creates another question worth monitoring. This release does not establish an efficient attack on deployed encryption. It also does not verify that OpenAI deliberately withheld cryptographic breakthroughs. The question remains open.

For management, the implication is to monitor assumptions and preserve options. We cannot wait until every technical question has been settled before deciding how to prepare.

Evidence and further reading

OpenAI - Sharing AI progress in mathematics, 6 October 2026 · <https://openai.com/index/sharing-ai-progress-in-mathematics/>

OpenAI mathematical manuscript collection and verification caveats · <https://github.com/openai/math>

Can you see how value moves?

The network on screen represents customer promises, expertise, decisions, delivery, quality, margin and learning. Follow the work: where does it wait? Who makes a judgement? Which assumptions remain unwritten? What happens when the person who knows the exception is unavailable?

This is where transformation starts. Make the actual workflow legible: inputs, decisions, handoffs, outputs, exceptions and evidence. A digital organizational twin needs a representation of how value moves, not just an inventory of systems or an impressive organizational dashboard.

Evidence and further reading

Rico Kerstan & Miriam Büxenstein - AI Management: Organization in Agentic Times. Owner-provided manuscript; chapters 3, 11 and 20.

Build your own challenger.

The Organizational Singularity by Salim Ismail and contributors describes the Edge Twin and an AI-native transformation argument: practical execution at the edge, rather than bolting tools onto every existing handoff. These are proposals for organizational design, not universally established empirical laws.

My transfer is to treat the twin as a management capability under uncertainty. It models real workflows, runs bounded experiments, exposes assumptions and challenges how value is created. It can act as a challenger to the business model while also learning how to improve execution.

Start with one valuable workflow. Define the customer outcome and the conditions under which it is acceptable. Extract the knowledge required, make decision rights explicit, run the alternative alongside the current way, compare real results and decide what to change. A twin that only watches is useful, but it has not yet demonstrated the ability to perform the work.

Evidence and further reading

Salim Ismail with contributors - The Organizational Singularity, living book · <https://openexo.com/organizational-singularity>

Rico Kerstan & Miriam Büxenstein - AI Management: Organization in Agentic Times. Owner-provided manuscript; chapters 3, 11 and 20.

Attack. Test. Learn. Rebuild.

The loop is model, challenge, decide, execute, observe and rebuild. At its centre sits human accountability. This does not imply a person must approve every minor action. It means people decide the purpose, constraints, authority and acceptance criteria - and remain answerable for the system they deploy.

Transformation therefore needs technical understanding at management level. Leaders do not all have to become software engineers, but they need enough understanding to distinguish a demonstration from reliable operation and to commission something more useful than disconnected tools.

Your vision should describe how the organization will create value in the new environment. Your operating approach should allow the route to change as you learn. We help organizations connect those two things: a direction worth pursuing and practical execution that can be evaluated and improved.

AI Management also brings us to intervention points: who can stop a run, return work, escalate an exception and revise the rule? Human accountability has practical meaning only if intervention is possible before the final outcome. The costs of quiet human repair must remain visible.

Evidence and further reading

Rico Kerstan & Miriam Büxenstein - AI Management: Organization in Agentic Times. Owner-provided manuscript; chapters 3, 11 and 20.

100x.

Consider a thought experiment, not a measured gain or a promise. Assume a hundredfold productivity gain on a defined activity at constant quality. One possibility is to reduce the people doing it from one hundred to one for the same theoretical capacity. Another is to keep one hundred people and produce the equivalent capacity of ten thousand baseline workers.

Capacity is not demand, revenue, profit or social benefit. Bottlenecks may move to production, review, customer acquisition or attention. The task may only be a small part of a job. Employment effects depend on what the organization decides and what the market actually buys.

Some organizations will reduce certain roles, some will grow, and many will do both in different areas. What is irresponsible is pretending the distribution of those gains is simply a technical parameter.

The financial figures deliberately illustrate a conditional hundredfold gain per employee. Start at €100,000 revenue and €10,000 profit per person. At 100x sold output per person, with unchanged prices and a constant 10% profit margin, these become €10 million revenue and €1 million profit per person. In the one-person case total sales stay €10 million; in the 100-person case total sales reach €1 billion. These are arithmetic scenarios. Additional AI and transition costs are excluded; constant margin and sufficient demand are strong assumptions, not predictions.

Capacity can change the job.

Ingka Group reported in June 2023 that Billie resolved approximately 47% of received customer enquiries from 2021 to 2023. In the same account, it described 8,500 call-centre co-workers reskilled in remote interior design, digital retail sales, relationships and complex enquiries.

This is a concrete example of automation connected to a new use of human capability. It is not evidence that AI never causes job losses, that all 8,500 received a single identical job, or that IKEA guaranteed employment forever. The important question is what the organization chose to build around the capacity it gained.

In your organization, could experts spend more time on unusual customer requirements, materials, quality and design judgement if the surrounding coordination work changed? The answer needs evidence from your workflows, not a generic workforce forecast.

Evidence and further reading

Ingka Group - AI and Remote Selling bring IKEA design expertise to the many, 29 June 2023 ·
<https://www.ingka.com/newsroom/ai-and-remote-selling-bring-ikea-design-expertise-to-the-many/>

What would you choose?

The experiment is now on your phone. Change people, productivity and the share of capacity actually sold. Your choices stay on your device and do not change the stage. The first figure is hypothetical capacity relative to a baseline worker.

Choose a concrete activity. What limits the benefit first: demand, quality, trust, physical throughput, review time or budget? Those are the first assumptions the organizational twin should test.

The third slider controls the share of theoretical capacity actually sold. Revenue equals capacity times €100,000 times that share; profit is an illustrative constant 10% of sales. Total and per-person figures move together. Set capacity sold to 10% to show how demand realization changes the apparent economic upside. This model does not estimate fixed costs or layoffs; use real unit economics in a decision.

Redesign it. Before they do.

Start with one valuable workflow that is slow, coordination-heavy or exposed to a small AI-enabled competitor. Make it visible. Give a bounded team the mandate to challenge it. Judge the alternative on real outcomes and total effort, including human supervision and correction.

AI does not remove the need for management. It increases the need for clear decisions, technical understanding, controlled experimentation and accountable execution. If you want a partner to connect those things in your organization, that is the work we do.

Scan the same QR code for the presentation and this narrative script. Keep exploring the simulation. If you want to build the management capability behind it, the upcoming AI Capabilities training on the KR eCampus will continue that work.