P H I - L E A N P R E S S
Occasional Papers in Operational Excellence

Occasional Paper No. 1 October 2026

THE SYSTEMIC ATTRITION SPIRAL

When Scale Learns Proximity: How Artificial Intelligence Threatens the Last Advantage of SMEs, and the Lean AI Response

MANUEL GOURMAND
Phi-Lean
Liverdun, Grand Est, France · Royal Tunbridge Wells, United Kingdom

A B S T R A C T

SME's (that is, in the context of this research paper, small and medium enterprises that are less than 150 employees) have always lost to big corporations on scale. They survive on intimacy—on being closer, faster, and more human. They know when and how to bend the rules when the situation requires.

This paper studies our past technological revolutions – personal computers, robotisation and automatization of the shop floor, ERPs, the Internet – and predicts the impact Agentic AI will have on SME on a two to five years horizon. Each of the past revolution got multinational corporations stronger, each caused further extinctions of smaller businesses. The author studied what behaviour helped successful SMEs to survive, and what behaviour drove others to fail. The findings – that a Lean, or related operational excellence approach, is the underlying recipe to success - can be summarise in those Deming words:

"It is not necessary to change. Survival is not mandatory."

Attributed to **W. Edwards Deming**[1]

Most analysts get the threat of artificial intelligence dead wrong. They think it is a cost story. It isn't. Small firms were never going to win a price war anyway. The real danger is that AI lets massive multinational companies, fake intimacy at scale.

Consider IKEA's primary retailer. They offloaded routine customer service to an automated agent, then redeployed 8,500 call-centre employees as remote interior designers. In a single shift, they got cheaper and grew closer to their customers. This is what the author decided to call "the double squeeze" on SME.

Anything purely informational can be duplicated instantly. True craft, physical presence, and real local accountability still protect the small businesses—for now. But basic customer expectations move fast. Today's instant quote is tomorrow's baseline.

[1] As Deming famously noted in his seminars, organizational change is optional because "survival is not mandatory" (as cited in Latzko & Saunders, 1995, p. 19).

SMEs can choose to wait and slide into a Systemic Attrition Spiral of widening cost gaps, sadly at the same time generating exploding Failure Demand, all this while losing margins and talent. Two temptations arise: use AI to cut heads, or use AI to automate what shouldn't be automated, manufacturing Failure Demand. Accounts from Chinese business leaders, published in 2026, point the same way: fix the process before adding AI, keep experienced people close to the work, and protect the junior tasks where judgement is learnt. Our response is a five-step Lean AI Roadmap.

1. Identify what the customer actually cares about.
2. Map the work required to deliver it.
3. Stabilise that workflow until it stops breaking.
4. Automate the precise bottleneck holding it back.
5. Shift the staff you free up directly toward customer relationships.

Handled this way, cheap tools allow a twenty-person shop to scale across international borders and clear regulatory hurdles that used to require an army of lawyers.

The software costs near to nothing. The discipline to implement it rightly, is what costs everything.

Keywords: agentic artificial intelligence; SMEs; customer intimacy; Kano model; customer value; economies of scale; lean management; Failure Demand; reskilling; China.

1. THE THREAT IS PROXIMITY, NOT COST

Ask a room full of SME owners how AI threatens them and most will talk about price. The big players automate, their unit costs fall, and a gap that small firms already live with gets wider.

That is right. It is also the lesser danger.

Small firms have never won on cost. They survive because they know customers by name, pick up the phone, and bend the rules when a job needs it — and AI is now giving large firms a credible version of that closeness, at scale. The serious threat is not that AI makes big firms *cheaper*. It makes them *closer*.

First, whom this paper is about. An SME here is an independent company with fewer than 150 employees that is not part of a larger corporate entity: no parent group, no franchisor, no investor platform supplying its systems.[2] The European Union draws the line higher, at 250 staff, with ceilings on turnover and balance sheet (European Commission, 2003). We draw it lower on purpose. This paper's argument turns on proximity, and in a small firm proximity is a personal habit before it is a process. Somewhere around 150 people the habit stops scaling: the owner no

[2]Statistics cited in this paper use official size classes; Eurostat, for instance, counts enterprises with 10–49 employees as small and 50–249 as medium-sized. Those classes bracket our population rather than match it, so the figures are read as indications of direction.

longer knows every customer and every employee by name. The figure is close to Dunbar's (1993) estimate of the number of stable relationships a person can keep track of, which makes it a rough anchor rather than a law. Independence matters for a different reason. A subsidiary can borrow its group's data, IT department and capital, which puts it on the other side of the squeeze described below.

Waiting used to be a sensible answer to new technology. In the microcomputer and early internet years a late adopter could let prices fall and standards settle, then buy the mature product. Not any more. The payoff from information technology has always come less from the kit than from the reorganisation built around it (Bresnahan, Brynjolfsson and Hitt, 2002; Brynjolfsson and Hitt, 2000), and that reorganisation takes years to build and shows up in the figures late (Brynjolfsson, Rock and Syverson, 2021). Software can be bought in a week. An organisation that knows how to use it cannot.

SMEs start behind. They already lag larger firms on digital adoption, held back by management time, skills and finance (OECD, 2021), and early AI use has been concentrated in the largest firms (McElheran *et al.*, 2024). An SME that defers adaptation now is not postponing a gain; it is losing ground every month to competitors whose gains compound. Its customers, meanwhile, are learning to expect what those competitors offer (Section 4).

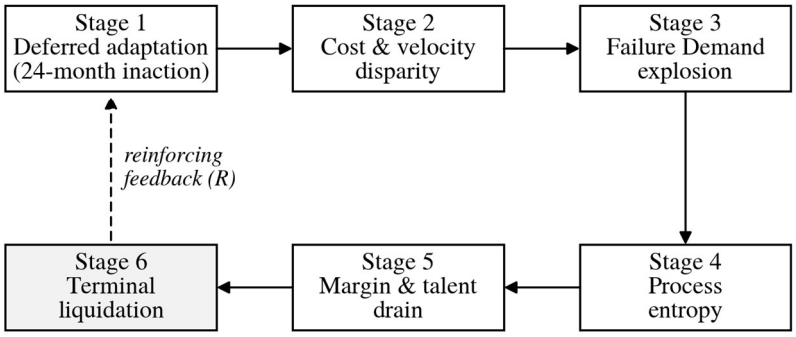


FIGURE 1. *The Systemic Attrition Spiral (SAS): a reinforcing loop of operational decay.*

Figure 1 shows the loop. It begins with a decision that does not feel like one: waiting. Competitors who move first open a cost and speed gap (Stage 2). Customers used to instant answers elsewhere start chasing — status calls, complaints, corrections — and Failure Demand explodes (Stage 3). The team spends its days firefighting, workarounds multiply, and the process degrades (Stage 4). Margins thin and the best people leave for firms where their output, and their pay, is higher (Stage 5). Liquidation or a distressed sale follows (Stage 6). Each stage makes the next harder to escape. Hence a spiral, not a decline.

One warning before going further. Inaction is not the only way in. A firm that uses AI as a blunt tool for removing people, rather than for redesigning work, enters the same loop at Stage 5 (Section 10). A firm that puts a rigid chatbot between itself and its customers enters at Stage 3, manufacturing its own Failure Demand (Section 6). Speed matters. So does the kind of action taken.

2. THE DOUBLE SQUEEZE

2.1 Two ways to win

Large firms win on cost across broad markets; small ones win through focus in narrow ones (Porter, 1980). Small firms carry *material* handicaps, such as less capital and less specialist knowledge, but enjoy *behavioural* advantages: flexibility, motivation, closeness to the customer (Nooteboom, 1994). Operational excellence against customer intimacy, in Treacy and Wiersema's (1993) terms.

What kept the two apart was a trade-off. Evans and Wurster (1997) described it as *reach* against *richness*. Information could be shared with many people, or it could be detailed, tailored and two-way, but not both, because rich exchange needed people in the room. The big firm reached millions, thinly. The small firm served fewer customers, richly. That trade-off — not any virtue of being small — is what protected the SME.

2.2 Proximity is not a postcode

The word *proximity* carries most of this paper's weight, so it needs pinning down. It is not the distance between the shop and the customer's front door. It has four faces.

Physical proximity is being there: the van that turns up the same afternoon, the fitter who sees that the wall is not square. *Relational* proximity, or customer intimacy, is knowing the customer — history, house, habits — well enough to anticipate the next request (Treacy and Wiersema, 1993). *Responsiveness* is the speed and flexibility of the answer: a decision in the time it takes to ask, from someone with the authority to bend a rule. *Emotional resonance* is the hardest to name and the easiest to feel. The customer believes they matter to this firm, shares something with it (a town, a craft, a set of values), and would feel a small betrayal in going elsewhere.

AI does not touch the four faces equally. Table 1 shows where it adds value and where people stay indispensable, using the pairings every SME owner knows: the local shop against the big footprint, the artisan against the factory.

TABLE 1 — *Four faces of proximity: where AI adds value, and where people stay indispensable*

Face	Where AI adds value	Where people stay indispensable	Illustration
Physical (being there)	Scheduling and routing; live arrival times; forecasting which parts the van needs.	The skilled visit: measuring, fitting, repairing, noticing what the customer did not mention.	A hardware shop texts "ready at 14:30" and holds the right stock; the counter assistant who knows which plug holds in a 1930s plaster wall stays human.
Relational (customer intimacy)	Memory: the full customer history at hand; service reminders; drafted follow-ups.	Reading the person: the customer who "just wants a quote" but is renovating after a bereavement.	An independent wine merchant's system recalls every bottle bought; the recommendation across the counter, and its credibility, stay with the merchant.

Face	Where AI adds value	Where people stay indispensable	Illustration
Responsiveness (speed and flexibility)	Instant answers to known questions at any hour; same-day quotes; order status without a call.	Exceptions, the authority to bend the rule, an apology that means something.	A kitchen studio's agent answers "where is my worktop?" at 22:00; the owner decides on the spot to replace a damaged door before fitting day.
Emotional resonance (mattering)	Little directly. It frees time and removes the irritants that wear goodwill down.	Shared identity, care, craft and the accountability of a named person.	An artisan baker's pre-orders and baking quantities are forecast by software; the baker's hands, the sourdough's story and the morning greeting are the product.

One pattern runs down the table. AI is strong wherever proximity is information: knowing, remembering, answering. It is weak wherever proximity is presence, judgement or care. Section 3 turns that distinction into a test.

2.3 Blade one: AI deepens scale

AI performs best on processes that are high-volume, standardised and well documented, which is exactly what large firms have spent decades building. Adoption so far has clustered among them (McElheran *et al.*, 2024), and the intangible investment AI demands favours whoever can pay for it (Brynjolfsson, Rock and Syverson, 2021). The first effect of AI is to widen a cost gap that already exists.

Serious. Not new.

SMEs have never competed on unit cost. A wider gap hurts in price-sensitive segments, yet it leaves the logic of focus intact; if this were the only effect, the old SME playbook would still work.

2.4 Blade two: AI attacks proximity

The second effect is different in kind. The internet broke the reach–richness trade-off for information (Evans and Wurster, 1997); generative and agentic AI are breaking it for conversation and judgement, delivering personalised service at scale (Huang and Rust, 2021). Volume becomes an input too, since every interaction a large firm handles is material for serving the next customer better (Varian, 2019).

Ingka Group, the largest IKEA retailer, shows what this looks like on the ground (Table 2). From 2021 an AI agent called "Billie" handled around 47% of customer enquiries, saving roughly €13 million over two years (Ingka Group, 2023). Ingka did not bank the saving by cutting staff. It retrained about 8,500 call-centre co-workers as remote interior-design advisers; asked whether AI would lead to job cuts, its global people and culture manager indicated that this was not what the company was seeing (Retail Gazette, 2023). The remote channel brought in €1.3 billion in the 2022 financial year (Ingka Group, 2023; Retail Gazette, 2023).[3]

[3]In March 2026 Ingka announced the removal of around 800 roles, reported as concentrated in its corporate group functions as part of an organisational simplification; the announcement did not refer to AI or to customer-facing roles (Retail Gazette, 2026).

TABLE 2 — *Exhibit: Ingka Group (IKEA), AI deployment and human redeployment*

Indicator	Reported value	Source
Share of enquiries handled by AI agent "Billie" (2021–2023)	≈47% (≈3.2 million interactions)	Ingka Group (2023)
Operating savings attributed to Billie	≈€13 million	Ingka Group (2023)
Call-centre co-workers reskilled as remote design advisers	≈8,500	Ingka Group (2023)
Remote design and selling sales, FY2022	€1.3 billion (3.3% of total)	Retail Gazette (2023)
Target share of sales from remote channel	10% by 2028	Retail Gazette (2023)

IKEA used AI to get closer. The agent took the routine questions; the people it released now sit with customers planning kitchens and living rooms, giving the kind of trusted one-to-one advice that used to belong to the independent kitchen studio down the road. Nothing suggests IKEA was aiming at small competitors. It did not need to. In the terms of Table 1, it took the relational and responsive faces of proximity without sending a single van. The capability now exists, and any large firm willing to redeploy its people rather than remove them can use it.

2.5 Stuck in the middle without moving

Porter (1980) argued that profitability follows a U-shape. Cost leaders and focused differentiators do well; firms "stuck in the middle" do not. The reach–richness trade-off kept the small, focused firm safely on one arm of that U, and the double squeeze takes the protection away (Figure 2). Blade one pushes the large firm further along the cost axis. Blade two lifts it up the proximity axis. An SME whose closeness no longer stands out, and whose costs are still higher, lands in the middle without having changed anything at all.

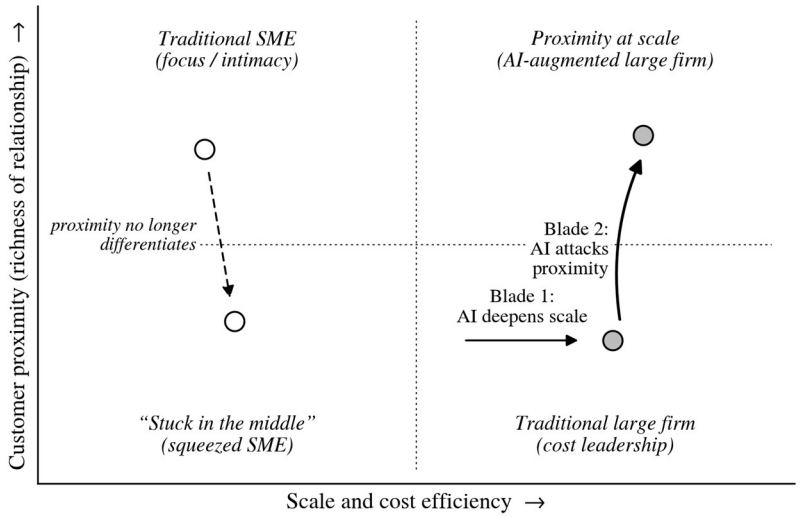


FIGURE 2. *The double squeeze: AI deepens the scale of large firms and lets them attack the SME's proximity advantage.*

Three claims follow from this, and the rest of the paper tests them.

1. AI widens the cost gap between large firms and SMEs, because it rewards the standardised, high-volume processes that large firms already have.

2. AI erodes the proximity advantage of SMEs, because it lets large firms deliver rich, personalised, human-assisted service at scale.
3. The only durable SME response is to use AI the same way (automate the routine, move people closer to the customer) and to do it faster than large firms can.

The first claim is the weakest of the three. Section 11 takes on the best argument against it.

3. WHAT AI CAN COPY, AND WHAT IT CANNOT

The attack has limits, and an owner needs to know exactly where they lie, because the answer decides what to automate and what to defend.

3.1 A test for protection

Resource-based theory supplies the test. An advantage lasts when rivals cannot copy it, and it resists copying when it is socially complex, built up over a long history, or hard to explain even for the firm that holds it (Barney, 1991). Automation adds a sharper edge. For decades computers took over tasks that could be written down as rules, while the things people know but cannot fully explain resisted, a limit Autor (2015) calls Polanyi's paradox. Generative AI has pushed that frontier back, because it learns from examples rather than rules. It has not abolished it. Data also shows decreasing returns to scale, with skill rather than volume the usual bottleneck (Varian, 2019).

Three questions sort any SME advantage. Is it information, or presence? Information can be copied at scale; a person in the room cannot. Was it built in months, or over decades? A large firm can buy a model this year, but not twenty years of standing in a town. Can someone be held to account for it? An owner can. A chatbot cannot, as Air Canada learned when it argued otherwise (Section 6).

Table 3 applies the three questions to the advantages SMEs usually claim against large firms, item by item.[4]

TABLE 3 — *The SME advantage matrix: how a large firm can use AI against each advantage, and how well it is protected*

SME advantage	How a large firm can use AI to mimic or attack it	Protection	Why
Knowing customers by name and history	CRM memory and generative agents that recall every past order and draft personal messages; remote advisers briefed by AI before each call, as at IKEA.	Low	Mostly information, and information copies at scale.
Fast answers to routine questions	Agents answering in seconds, around the clock, across chat, email and phone.	Low	Large firms will soon answer faster than any small team can.
Specialist knowledge of a niche	Models trained on manuals, standards and case histories; expert copilots for the large firm's front line.	Low to medium	Codified expertise becomes a commodity; judgement built on experience does not.

[4]The protection ratings in Table 3 are the author's judgement, based on the three questions above and on field experience with SMEs. They are meant to start a discussion inside a management team, not to replace one.

SME advantage	How a large firm can use AI to mimic or attack it	Protection	Why
Serving small orders and the long tail	AI cuts the cost to serve, so orders once too small to bother with become worth taking.	Low to medium	The shelter of "too small for them" shrinks.
Authority to bend the rules	Policy engines let agents grant goodwill gestures within set limits; escalation to empowered advisers.	Medium	Routine exceptions can be coded; the unprecedented one still needs an owner's judgement.
Tailoring and customisation	Configurators, generative design and mass-customisation lines that tailor at near-mass cost.	Medium	A strong copy for standard variants; a weak one for one-offs that need a craftsperson.
The same person, year after year	AI memory gives rotating staff the full history, simulating continuity.	Medium	History transfers. Rapport does not.
Short decision chains, no legacy systems	Agentic platforms bought centrally and pushed into front-line processes.	Medium, and time-limited	Legacy systems and governance slow the large firm (Section 11), but the window closes as it pays down that debt.
Physical presence and same-day service	Route optimisation, demand forecasting, local stock points, partner networks of independent fitters.	Medium to high	AI sharpens logistics. It does not put a skilled person on site.
Craft and provenance	Generative design and robotic finishing imitate the look of hand-made goods; AI-written origin stories.	High	The look can be copied. The maker, the material's origin and the history cannot, and pretending otherwise is misrepresentation.
Local trust and community roots	Hyper-local targeting, local content, sponsorship bought at scale.	High	Earned over years of visible conduct; it cannot be bought in a quarter.
Owner accountability	Named "personal" advisers and AI avatars that simulate a single point of contact.	High	Somebody with a name answers for the job. A simulated contact cannot carry that responsibility.
Emotional resonance and shared values	Emotion-aware chat, brand storytelling, personalisation at scale.	High, if made explicit	Customers sense simulated care, but values an SME never states cannot resonate.

3.2 Reading the matrix

Read down the protection column and the pattern of Table 1 returns. Everything an SME does that is really information (knowing, remembering, answering, advising from codified knowledge) sits at the top, exposed. Everything that is presence, judgement, history or accountability sits at the bottom, protected.

The uncomfortable part is the top. Ask SME owners what makes them special and they name the exposed items first: "we know our customers", "we pick up the phone", "we're the experts". True today. Copyable tomorrow. The local plumber is safe from AI; the local accountant, whose value is largely codified rules applied to clients' figures, is not.

Two conclusions follow. The exposed advantages are not a moat but a race: they keep their value only if the SME automates them first and spends the time it gets back deepening the protected

ones. And protection is not automatic. Local trust that is never maintained decays. A craft that is never explained is invisible to the buyer comparing two websites. The matrix shows where the walls are; somebody still has to keep them up.

3.3 The SME's head start

There is also an asymmetry in the SME's favour. IKEA had to *build* proximity, retraining thousands of people and creating a new channel from scratch. The SME already *owns* it. What it lacks is time, because its people are buried in routine work, rework and Failure Demand, and AI can give that time back. So the race is not about who has the better technology. It is about who first turns automation into human closeness to the customer: the large firm learning proximity, or the small firm learning scale.

4. TODAY'S DELIGHTER, TOMORROW'S BASIC

Noriaki Kano and his colleagues gave quality managers a way to see why customer expectations move (Kano *et al.*, 1984). Customers do not judge every attribute of a product or service the same way. *Basic* (must-be) attributes cause anger when missing and earn no thanks when present; nobody praises a hotel for clean sheets. *Performance* attributes satisfy in proportion: the faster, the better. *Delighters* surprise. Their absence costs nothing, because nobody expected them, and their presence wins loyalty.

The part of Kano's theory that matters here is about time. Attributes do not stay in their category. Delighters slide into performance and then into basics, as competitors copy them and customers get used to them (Löfgren, Witell and Gustafsson, 2011). Parcel tracking on a website was a novelty in the 1990s. Today a delivery without a tracking link generates a phone call.

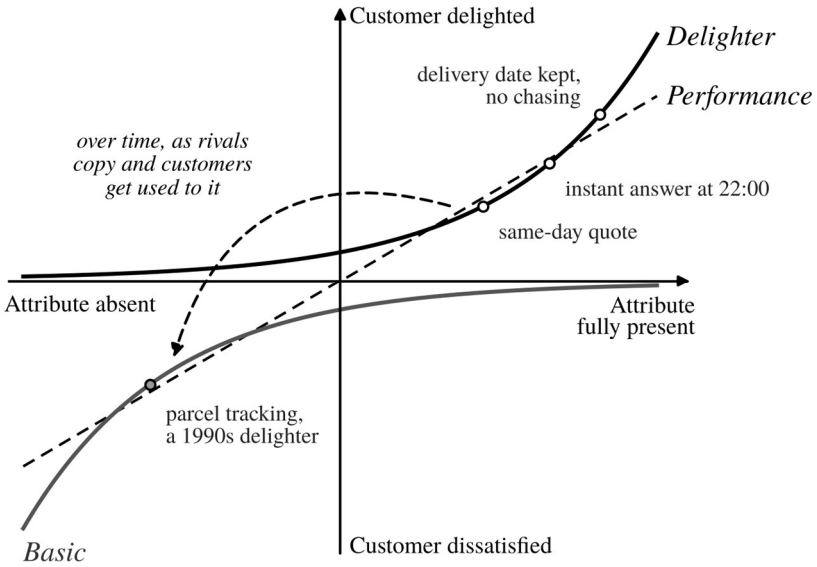


FIGURE 3. *Today's delighter, tomorrow's basic: AI-enabled service attributes on the Kano clock.*

AI-enabled customer value sits on that clock now (Figure 3). An instant answer at ten at night. A quote the same day instead of next week. A delivery date that is confirmed, kept, and updated without anyone chasing. A firm that remembers what you bought three years ago. Today these delight, because most firms, large or small, do not offer them.

They will not delight for long. Every large firm that deploys them teaches the SME's customers what to expect, and expectations take no account of company size. A customer answered in seconds by IKEA at 22:00 does not lower the bar for the kitchen studio the next morning.

This is the mechanism behind Stages 2 and 3 of the spiral. The non-adopting SME does not get worse. It stands still while the definition of basic moves past it, and every gap between what the customer now expects and what the firm still does arrives as a chasing call: Failure Demand created by somebody else's innovation.

Two consequences follow. The first is timing. In the delighter phase the early adopter is paid in loyalty and margin for the same AI that will later be compulsory; in the basic phase it merely stops losing customers. Same investment, smaller return.

The second concerns the protected advantages of Table 3. They exempt nobody from the new basics. In Kano's logic the customer judges the basics first, and no delighter makes up for a basic that is missing. The artisan cabinet-maker's craft remains a delighter no algorithm will commoditise. If her customers need three phone calls to get a delivery date, though, the craft will be judged through the irritation.

Human-only can still be a strategy. Slow cannot.

5. WHAT THE LEAN CLASSICS ALREADY TOLD US

In every technology shift of the past forty years, technology alone never decided who survived. When small firms adopted computers, the owner's knowledge and appetite and the skills of staff counted for more than firm size (Thong, 1999). With the internet, the payoff came from reorganising: decentralised decisions, redesigned processes, new incentives (Brynjolfsson and Hitt, 2000). Firms that skipped the reorganisation got the cost without the benefit.

Lean practitioners know this pattern well. Each of their founding texts turns into a rule for AI.

Deming (1982), watching American industry lose ground to Japan, listed among the obstacles to transformation

> "the supposition that solving problems, automation, gadgets, and new machinery will transform industry."
>
> (Deming, 1982, pp. 127–128)

Automate an unstable process and you get the same defects, faster. Hammer (1990) reached the same verdict from the IT side: computers that speed up obsolete processes rarely pay, so redesign the work first.

Imai (1986) separated innovation, the big step bought with capital, from kaizen, the daily improvement owned by everyone. Innovation without kaizen starts to decay the day it is installed. A firm that buys point-solution AI without reworking its *gemba* ends up relying on the vendor to keep it running.

Goldratt's *The Goal* (Goldratt and Cox, 1984) opens with expensive robots that lift one section's efficiency figures while inventory stays put and throughput does not rise. To put it bluntly, the plant paid more and shipped the same. The parallel today is exact: SMEs using AI to write emails faster while the real constraint (hand-offs, slow decisions, manual re-keying) sits untouched.

Seddon (2003) gives the distinction that matters most in the AI era. *Value Demand* is what the customer actually wants. *Failure Demand* is demand caused by a failure to do something, or to do it right, for the customer: the chasing call, the complaint, the correction. Command-and-control

firms answer backlogs by adding checks, which creates more Failure Demand. AI deployed with the same logic does the same thing at machine speed (Table 4).

TABLE 4 — *Two logics of AI deployment*

	Command and control	**Systems thinking (Seddon)**
Role of AI	A cost-cutting tool to process transaction volume and reduce headcount.	Deployed only after end-to-end flows are redesigned around Value Demand.
Design unit	Functions, activities and targets.	The customer's end-to-end flow.
Released capacity	Removed from the organisation.	Redeployed to the demand the system cannot yet absorb.
Outcome	Amplifies Failure Demand: errors, customer friction and rework.	Removes processing delay; throughput scales with demand.

Devine (2023) shows that change stalls whenever the front line is left out of process design. In a technology shift management cannot foresee every consequence, so the people doing the work have to adapt the tools as they go. Bicheno and Holweg (2016) supply the test for any tool: does it serve takt time and cut lead time? If not, it is a cost. And Liker (2004) reminds us that Toyota's Principle 8 asks for reliable, thoroughly tested technology that serves people and process. That is a rule about where and how technology enters the system. It is not a licence to wait.

Clean flow, a habit of PDCA, money spent at the constraint. A Lean SME with all three absorbs a technology shock; a firm without them spirals.

6. THE THIRD DOOR: AI THAT MANUFACTURES FAILURE DEMAND

Section 1 named two ways into the spiral: waiting, and cutting people. There is a third, and many SMEs are walking through it now. The firm adopts AI quickly, at the customer interface, without redesigning the work behind it. It buys a chatbot.

Seddon's arithmetic shows why that is dangerous. When his team studied 1,200 calls into a council housing call centre, 80% of them turned out to be Failure Demand (The Open University, n.d.). A bot placed in front of a flow like that does not remove the failures. It answers them faster, or fails to answer them, and either way the customer comes back.

The evidence is no longer anecdotal. The US Consumer Financial Protection Bureau has described customers trapped in "doom loops" of repetitive, unhelpful answers with no route to a person, and warned that failed chatbots can break trust and cause widespread harm (CFPB, 2023). In a Gartner survey of 5,728 customers conducted in December 2023, 64% said they would prefer companies not to use AI in customer service; the top concern was that it would become harder to reach a person, and more than half said they would consider switching to a competitor on learning that a company planned to use it (Gartner, 2024).

Then the courts. In 2024 a Canadian tribunal ordered Air Canada to compensate a passenger whose bereavement-fare claim had followed advice from the airline's chatbot, advice that contradicted the airline's own policy. Air Canada argued, in effect, that the chatbot was a separate legal entity responsible for its own actions. The tribunal called this "a remarkable submission" (*Moffatt v. Air Canada*, 2024).

And the cautionary arc. In February 2024 Klarna announced that its AI assistant was handling two-thirds of its customer service chats, the work of 700 full-time agents, with repeat enquiries

down by a quarter (Euronews, 2024). On the numbers chosen, it worked. Fifteen months later the chief executive said the cost-cutting drive had gone too far, and the company began recruiting people so that customers could always reach a human (Daly, 2025).

For a large firm a bad bot is a cost. For an SME it is worse, because it attacks the one thing the firm sells. A customer who has spent twenty minutes in a loop does not remember that the firm is small, local and caring. She remembers the loop. In Kano's terms the bot has become a *reverse* attribute, one whose presence dissatisfies (Kano *et al.*, 1984). In Seddon's terms every turn of the loop is Failure Demand with a timestamp. The SME that installs one has done the large firm's work for it: blade two, self-inflicted.

The rule is not "no bots". IKEA's Billie handled nearly half of all enquiries, and the people it released moved closer to customers (Section 2.4). The rule is that an agent sits behind a stable process and in front of a person, never in place of one. Table 5 lists the usual failures and their countermeasures.

TABLE 5 — *How customer-facing AI manufactures Failure Demand, and the countermeasure*

Failure	What the customer experiences	Countermeasure
The loop without an exit	The same three options again, and no route to a person.	One step to a named human, always visible.
The confident wrong answer	A plausible reply that turns out to be false, as with Air Canada's refund advice.	Agents answer only from the system of record; anything else is handed over.
Deflection as the target	A bot rewarded for keeping customers away from staff does exactly that.	Measure first-contact resolution and repeat contacts, never deflection.
Automating the failure, not the fix	"Where is my order?" answered politely ten thousand times; the late order never fixed.	Remove the cause first (Stage II); automate only what remains.
No memory across channels	The problem explained to the bot, then to a person, then again by email.	One case history, passed on with the hand-over.
The bot as gatekeeper	The best customers, who used to ring the owner, now meet a menu.	Key customers keep a direct line. The bot serves; it does not screen.

If the bot cannot fix it, the bot must not stand between the customer and someone who can.

7. THE COST OF WAITING

Gains from new general-purpose technologies arrive unevenly; they pool where the skills and organisation to use them already exist (Forman, Goldfarb and Greenstein, 2012). Early AI evidence fits the pattern. Adoption is concentrated among large firms (McElheran *et al.*, 2024). Those who do adopt see large task-level gains in controlled studies: around 15% more customer-support issues resolved per hour (Brynjolfsson, Li and Raymond, 2025) and roughly 40% less time on professional writing tasks (Noy and Zhang, 2023). Table 6 sets out how adapting and non-adapting SMEs then diverge.[5]

[5]The scenario shares, ranges and horizons in this section, and the 24-month threshold in Section 7.2, are the author's estimates, drawn from the literature cited and from field experience with SMEs. They are illustrative and are not the output of a fitted statistical model.

TABLE 6 — *Divergent operational trajectories*

Horizon	Non-adapting SME (SAS trajectory)	Lean AI adapted SME
Years 0–2	False security; margins stable.	Customer value and value streams mapped; AI pilots at the constraint.
Years 3–5	25–40% margin compression; talent loss.	Near-zero Failure Demand; 3× output.
Years 5–10	Severe talent drain; market exit.	Category leadership and scale.
Years 10–20	Terminal liquidation or fire-sale.	Ecosystem leadership; acquisitive growth.

7.1 Three ways to fail

Sector exposure, barriers to entry and the depth of customer relationships decide which of three paths a non-adapter takes.

The commonest is the slow leak (we estimate 50–60% of non-adapters). Competitors automate administration, quoting and dispatch; the industry's cost-to-serve falls while the non-adapter's stays put. Investment and benefits get cut, the best people leave for AI-augmented firms where they produce and earn more, and within a decade the firm serves mainly price-sensitive clients, has nothing a buyer wants, and winds down when the founder retires.

The second is the commoditisation shock (15–20%). It hits low-barrier, highly digital sectors such as accounting, legal services, digital marketing, freight brokering and light assembly. AI-native entrants offer comparable work at 10–20% of the old price, almost instantly. Principal accounts go. Cash flow turns negative within 36 to 48 months, and the firm is bought cheaply for its licences, client list or equipment.

The third is the digital artisan (20–25%), the firm that repositions as human-only, high-touch or hyper-local. It survives at premium prices in a smaller market, competing on trust and local relationships when cost advantage is out of reach (Barney, 1991). But the double squeeze makes this niche narrower than in past transitions. Human-only works where closeness is physical or relational. Where it is merely informational — advice, reassurance, tailored recommendations — it is precisely the ground that AI-augmented large firms are now taking (Table 3). And the digital artisan still has to meet the new basics of Section 4: human-only cannot mean slow.

7.2 The 24-month line

Why 24 months? Because the gap does not grow in a straight line. In the ERP era a firm three years behind could catch up by buying the same licences. With AI that learns from process data, the leader's advantage compounds: two years of cleaner processes, faster loops, and accumulated knowledge of each customer. A firm's ability to absorb new knowledge depends on what it already knows (Cohen and Levinthal, 1990), and firm-specific process know-how is exactly what rivals cannot buy (Barney, 1991).

The laggard, meanwhile, falls into what Repenning and Sterman (2001) call the capability trap. Under pressure, managers pull time out of improvement and put it into firefighting. Capability erodes; pressure rises; repeat.

Our field estimate is blunt. After about 24 months of inaction, the cost of re-engineering an SME's accumulated operational debt exceeds the present value of its future cash flows.

Past that point, rescue no longer pays.

8. THE LEAN AI ROADMAP

Survival means rejecting two things at once: doing nothing, and spending recklessly. The roadmap in Figure 4 puts customer value before process, process stability before digital integration, and ends not in automation but in proximity. It runs one value stream at a time, in PDCA cycles short enough to finish inside a quarter.



FIGURE 4. *The Lean AI Roadmap.*

To keep it concrete, picture a 40-person kitchen and bathroom business, the kind of firm most exposed to IKEA's remote design advisers. (It is an illustration, not a case study.)

8.1 Stage 0 — Define what the customer values

Lean starts from value as the customer defines it (Womack and Jones, 1996). Most SMEs skip that clause. They map the flow they have, assume they know what customers pay for, and automate accordingly. That is how firms end up automating the very thing customers came to them for.

So before drawing a single process box, map customer value. Talk to fifteen recent customers and five lost ones. Ask what made them choose the firm, what nearly made them leave, and what they would miss if it closed. Then sort every attribute they mention three ways. Its Kano category: basic, performance or delighter (the paired questions of a Kano survey, "how would you feel if it were there?" and "if it were not?", do this cleanly; Matzler and Hinterhuber, 1998). Its position against the large competitor: better, level or worse, in the manner of Gale's (1994) customer value map. And its exposure, from Table 3.

The result fits on one page, and it changes what gets mapped next. In our illustration, customers rarely mention price first. They mention the designer who came to the house, measured twice and spotted the boiler flue nobody else saw; the fitter who left the site clean; the owner who picked up the phone when the worktop cracked. Those are protected advantages: presence, craft, accountability. They also mention, with some irritation, the week it took to receive a quote and the three calls it took to get a delivery date. Those are basics in the making, and exactly what a large competitor's AI already delivers.

That page tells the firm what never to automate away and where automation must go first. The value stream map comes after it, not before.

8.2 Stage I — Map the value stream

Now map the flow that delivers what Stage 0 found. For two weeks, log every inbound contact (calls, emails, messages, walk-ins) and sort each one: is the customer asking for something they value, or chasing something we failed to do? In practice, Failure Demand in a business like this has a familiar face. "Where is my delivery?" "The quote shows the wrong worktop." "Nobody called me back." That log becomes the first fix list and the first automation backlog, in that order.

8.3 Stage II — Stabilise before you automate

Remove the causes of the largest Failure Demand categories before an agent ever touches them. Write the standard for the flow — survey, design, quote, order, fit, sign-off — and make it the way everyone works. The test is simple. If three designers price the same kitchen three different ways, an agent will give the customer a fourth. This is Deming's warning, and a senior staffing executive put it in eight words this year: "You cannot have AI in a bad process" (Lee, cited in He and Zumbühl, 2026).

8.4 Stage III — Automate at the constraint

Find the step that sets the pace of the whole business. In our illustration it is design and quoting, where two experienced designers' diaries decide how many jobs the firm can sell. Put the agent there: drafting quotes from survey notes, checking them against the price list, answering "where is my order?" straight from the order system, chasing suppliers. Keep a person one step behind every agent the customer can see (Section 6). Let the people who run the process build or configure it, since they know where the exceptions hide. Then measure what Goldratt would measure — lead time from survey to signed order, and jobs completed per month. Hours saved anywhere else count for nothing if the constraint has not moved (Bicheno and Holweg, 2016).

8.5 Stage IV — Redeploy to proximity

The released hours go back to the customer, and specifically to the protected advantages Stage 0 identified: longer design consultations, home visits, a call a month after fitting to check that the drawers still close. Stages 0 to III answer the first blade of the squeeze. Stage IV answers the second.

It is also the stage where most AI programmes quietly fail, because it turns on one question. What happens to the people whose work the agents absorb? Before answering, look at a country now asking that question at national scale.

9. THE VIEW FROM CHINA

China is pushing AI into the real economy through its "AI+" policy, while building employment and skills systems to absorb the disruption. An August 2026 World Economic Forum piece drew on business leaders and researchers there to describe what integration looks like inside firms (He and Zumbühl, 2026). This is practitioner testimony, not a controlled study. It still counts. The people quoted are watching AI pushed through a whole economy by policy, and what they report matches the Lean argument almost point for point: on their early evidence, AI changes how tasks, capabilities and systems are organised more than it replaces individual jobs. Table 7 maps the findings onto the roadmap.

TABLE 7 — *Lessons from Chinese firms, mapped to the Lean AI Roadmap*

Finding (He and Zumbühl, 2026)	Who reports it	What it means for the SME
"You cannot have AI in a bad process."	Ian Lee, The Adecco Group	Stabilise the flow before automating it (Stage II).
Firms of 20 or 30 people still have to decide which processes are ready for AI.	The authors	A value-stream map is enough to decide; no transformation team needed (Stage I).

Finding (He and Zumbühl, 2026)	Who reports it	What it means for the SME
Experienced workers are being brought back to guide AI-enabled work.	Ian Lee	Keep senior people close to the work; their judgement covers what agents cannot (Stage IV).
In a finance case, business professionals built the AI workflow themselves.	Former Huawei recruitment specialist (unnamed)	Let process owners configure the agents (Stage III).
An AI recruitment system or an AI course is not workforce transformation.	The authors	Access is not absorption (Section 11).
Automating junior tasks can strip out the work that builds judgement.	AIGW Policy Observatory	Name what a junior task teaches before automating it (Section 10).
A youth skills initiative for 1 million young people and around 1 million internships.	Employment Priority Strategy, 15th Five-Year Plan	Build a small in-house version: retraining tied to each automated task.

Several findings confirm Sections 5 and 8. Ian Lee, President of Geographic Regions at The Adecco Group, puts process first: understand the workflow, who is accountable, and where human judgement cannot be dropped, before any AI goes in. A former Huawei executive recruitment specialist adds that AI lowers the technical threshold, so the scarce asset becomes domain expertise combined with AI fluency and the agility to keep learning. In one finance case she cited, the business professionals built the AI-enabled workflow themselves, because they understood the process best.

Two others extend the argument. Both are about people.

Lee has seen experienced workers brought back into organisations to guide AI-enabled work, because they carry context and business judgement that task-level automation cannot supply. Some firms, in other words, are already buying experience back in.

The sharper warning concerns juniors. Routine and junior tasks are where people have always learnt how clients, decisions and professional judgement work. Exploratory research from the AIGW Policy Observatory warns that firms may automate exactly the tasks that build junior judgement, leave juniors with work that teaches little, and accept polished AI output that hides whether the employee could explain, verify or challenge it unaided. The article's advice is to identify the capability a junior task builds before automating it. Wenxia Zhou of Renmin University calls this a "mapless career era", in which the ladder gives way to something closer to a forest: people keep learning and keep moving across roles.

For a 40-person firm the risk is sharper than for IKEA. No graduate scheme. No training academy. The junior who fields the chasing calls this year is the estimator of 2029, and those calls are how she learns which jobs go wrong.

At national level, the Employment Priority Strategy under China's 15th Five-Year Plan (2026–2030) calls for new forms of human–AI collaborative work, a youth employment skills initiative aimed at 1 million young people, and around 1 million internship places. The authors conclude that workforce infrastructure is becoming part of AI infrastructure. No SME has a five-year plan or a ministry behind it; it needs a small version of the same thing.

10. WHAT TO DO WITH THE PEOPLE

The command-and-control instinct treats released capacity as cost to be cut. Brynjolfsson (2022) calls the habit of building AI to replace people rather than augment them the "Turing Trap": it concentrates gains and gives up the larger value augmentation creates. Labour economics agrees. Automation displaces workers from existing tasks, but technology also *reinstates* labour through new tasks where people hold the advantage, and the long-run outcome depends on the balance (Acemoglu and Restrepo, 2019).

For an SME, substitution is doubly dangerous. It throws away the tacit process knowledge future improvement depends on (Cohen and Levinthal, 1990). And it pushes customers into automated channels for exactly the complex problems those channels cannot resolve, creating new Failure Demand while cutting cost. The firm that automates to eliminate people enters the spiral at Stage 5 and gives away its proximity of its own accord — finishing the second blade's work for it.

Read through Seddon's lens, the IKEA case is the template. The agent absorbed the predictable part of the flow (order status, product information, routine queries), much of which is likely to be Failure Demand or low-value demand. What the agent could not resolve showed where the customer's real Value Demand lay: help designing a home. The released people were moved to precisely that work: reinstatement, carried out on purpose at firm level. In customer support, AI assistance gave its largest gains to novices and minimal gains to the most experienced agents, whose value lies in judgement the tool does not supply (Brynjolfsson, Li and Raymond, 2025). Read alongside the Chinese evidence, that finding carries a warning: if the tool does the hard part for novices, the firm has to make sure they still learn it.

SMEs cannot copy IKEA's scale. They can copy its sequence, and they start ahead, because they already own the proximity IKEA had to build — with far thinner buffers, admittedly. Five operating rules follow.

1. Let agents absorb routine work and Failure Demand.
2. Study, systematically, what the agents cannot resolve. That residue is where Value Demand lives.
3. Retrain people onto that residual demand.
4. Before automating a junior task, write down what it teaches, and decide how that will now be taught.
5. Turn the redeployed hours into something customers will pay for: advice, design, aftercare.

A 40-person firm that follows this sequence turns AI from a threat to its workforce into a multiplier of it. A firm that cuts the same people instead loses the tacit knowledge it needs for Stage II, damages its culture, and hands its best customers to competitors who still answer the phone with an expert.

None of this is comfortable. Ingka could absorb a multi-year transition from a position of scale; most SMEs cannot, and the choice described here has to be made inside the same 24-month window.

11. THE OBJECTION: CHEAP AI FAVOURS THE SMALL

11.1 The case

The first claim faces a serious objection. The enterprise systems of the last wave were expensive, slow to implement and imposed their own logic on the firms that bought them

(Davenport, 1998): technology priced for large balance sheets. Agentic AI runs on models rented by the call, and low-code platforms such as n8n let people with little programming background assemble serious workflows from configurable parts (Bock and Frank, 2021). At least on paper, the fixed cost of operational sophistication has collapsed.

The evidence behind the objection is real. IT investment has been followed, two or three years later, by a fall in average firm size (Brynjolfsson *et al.*, 1994). IT bought as a service was associated with lower failure rates for young US plants, owned IT capital with higher ones (Jin and McElheran, 2017). Cloud adoption lifted revenue, employment and productivity in young firms, with weaker gains for incumbents (DeStefano, Kneller and Timmis, 2025). Large firms carry technical debt (Cunningham, 1992; Sculley *et al.*, 2015) and the inertia that comes with being built for reliability (Hannan and Freeman, 1984). And if the gains from AI depend on redesigning the organisation as a system rather than task by task (Agrawal, Gans and Goldfarb, 2024), forty people can do that far more easily than a matrixed enterprise. On this view AI lowers the minimum efficient scale (Bain, 1956) of operational sophistication, and our first claim is overstated.

11.2 What it gets right

Much of it. The capital barrier has largely fallen: for a monthly fee an SME can rent capability no firm could have bought at any price a decade ago, and the Chinese finance case shows business staff, not technicians, building with it. The technical-debt point also adds a second asymmetry to the one in Section 3.3. The large firm has to learn proximity while dragging its legacy systems behind it; the SME can learn scale without them.

11.3 Where it falls short

Access is not absorption. Three facts keep them apart.

First, cheap tools have not closed the adoption gap. In 2025, 55.0% of large EU enterprises used AI technologies, against 30.4% of medium-sized and 17.0% of small ones. Among enterprises that had considered AI but did not use it, the reason most often given was lack of relevant expertise (70.9%), not cost (Eurostat, 2025). SMEs using generative AI report that it raises their need for highly skilled workers (OECD, 2025). If money were the binding constraint, falling prices would have closed the gap.

They have not.

Second, the returns depend on things that cannot be rented. Across more than 30,000 US manufacturing establishments, predictive analytics improved performance only alongside substantial IT capital, an educated workforce, or workplaces designed for high-efficiency production (Brynjolfsson, Jin and McElheran, 2021) — the last being close to what Lean practitioners call flow. In a 2025 benchmark simulating a small software company, the best agent tested completed only 30% of its tasks on its own (Xu *et al.*, 2025). That figure will rise. The exceptions, hand-offs and judgement calls will still land on the process and the people around the agent, and on an unstable process an agent just executes the variation faster.

Third, cheap access at the bottom of the stack sits alongside concentration at the top (Korinek and Vipra, 2025), much as large firms' proprietary IT drove most of the rise in US industry concentration in the last wave (Bessen, 2020). Whatever an SME can rent, its competitors can rent too. That buys parity. The advantage lies in the proprietary layer of data, integration and process knowledge around the tools, which large firms can still build and most SMEs cannot. Nor does low-code escape debt: dozens of ungoverned workflows soon carry the glue code and hidden

dependencies Sculley *et al.* (2015) found in machine-learning systems. Legacy debt, acquired at start-up speed.

11.4 Process maturity is the dividing line

AI lowers the minimum efficient scale of *access* to operational sophistication. It does not lower the organisational capability needed to *absorb* it (Cohen and Levinthal, 1990). The dividing line therefore runs less between large and small firms than between process-mature and process-immature ones, as Table 8 shows.

TABLE 8 — *Barriers to operational sophistication in two technology eras*

Barrier	Legacy enterprise IT (ERP era)	Agentic AI and low-code era	Change
Capital outlay	Large licence and integration budgets; favoured large balance sheets.	Usage-based pricing; minimal upfront investment.	Largely removed
Time to deploy	Multi-year implementation programmes.	Weeks for a single workflow.	Largely removed
Dedicated IT department	Required to implement and maintain.	Optional for building and piloting workflows.	Reduced
Stable, standardised processes	Required; the system imposed its own process logic.	Still required; agents execute variation faster.	Unchanged
Complementary skills, data and management practices	Required.	Still required; the most frequently cited barrier to adoption.	Unchanged
Technical debt	Accumulated in central legacy systems.	Accumulates in ungoverned workflows unless managed.	Displaced, not removed

The first claim therefore needs restating, in conditional form:

1′. AI widens the cost gap between large firms and process-immature SMEs, but can narrow it for process-mature SMEs, because it lowers the capital barrier to operational sophistication while leaving the organisational barrier intact.

So the objection does not refute the argument; it names the condition under which the third claim succeeds. The SME advantage is real but not automatic. It goes to firms that turn cheap access into capability through the roadmap in Section 8. For the rest, cheap agents automate disorder faster. The large firm's technical debt opens the SME's window, and process discipline decides whether the SME gets through it before that debt is paid down. Section 12 looks at what lies on the other side.

12. ROOM TO GROW

Everything so far has been about defence. The same tools also attack.

The tools that let IKEA learn proximity also let a 40-person firm learn scale, and scale used to be the one thing a small firm could not learn quickly. Growth meant hiring ahead of demand, and every new market came with three walls: language, regulation and logistics. Agentic AI lowers all three.

Language first. When eBay introduced machine translation on its platform, exports through it rose by 17.5%, and the pattern of gains pointed to lower translation-related search costs as the cause (Brynjolfsson, Hui and Liu, 2019). That was an earlier generation of the technology, on a single marketplace. An agent can now draft the quote, answer the follow-up and handle the after-sales query in the customer's language, overnight.

Regulation next. Customs codes, labelling rules, a neighbouring country's installation standards: work that once needed a specialist the SME could not afford is now work an agent can draft and a competent person can check. Logistics third. Forecasting, routing and the paperwork of cross-border shipping are the high-volume, well-documented tasks AI does best (Section 2.3). WTO economists estimate that AI could raise global trade by 34–37% by 2040, partly through lower trade costs (Bacchetta *et al.*, 2025).

Three routes open up.

Scale without hiring ahead. Agents absorb the routine growth in enquiries, quotes and order administration, so headcount follows value-adding work rather than paperwork. The firm can say yes to a larger contract before it has the back office for one.

Markets beyond the language border. For a firm in Lorraine the natural market already crosses three frontiers. A kitchen studio there can design for customers in Luxembourg and the Saarland, quote in German or French, and answer in either, without hiring a translator. The fitting still needs the fitter, and that sets a radius. The design consultation does not, and that widens it.

Niches too thin for one country. A cutler in Thiers who makes knives for left-handed professional cooks has a few hundred potential customers in France and enough across Europe, North America and Japan to build a business on. Language, export paperwork and the cost of serving distant buyers used to put that niche out of reach. They no longer do. And the protected advantages of Table 3 (craft, provenance, the maker's name) travel well, because they were never about distance.

New markets also reset the Kano clock. An AI-enabled delighter that has become basic at home may still delight abroad, for a while.

The door swings both ways. The same agents let a firm from Porto or Poznań sell into the Lorraine studio's home market. Growth is what the roadmap pays for, and the firms that skip it will meet competitors who followed it.

13. WHAT FOLLOWS

Once capital stops being the scarce input, the competitive gap opens along whatever is. In the AI era that is process discipline.

The framework makes five testable predictions. Returns to AI among SMEs should depend less on firm size than on measured process maturity and management practices (Bloom, Sadun and Van Reenen, 2012). The gap between process-mature SMEs and large firms should narrow, while the gap between process-immature SMEs and both should widen. Firms that redeploy released staff to customer-facing work should keep more customers than firms that cut them. SMEs whose customer-facing AI is judged by deflection should see more repeat contacts and lower retention than those that judge it by first-contact resolution. And firms that automate junior tasks without replacing what those tasks taught should, within a few years, find their supply of experienced people drying up. Firm-level data can test all five.

For owners, the implication fits on a whiteboard. This quarter, pick one value stream. Talk to fifteen customers and write down what they value, and which of those things a machine could copy. Log two weeks of inbound demand and separate value from failure. Fix the main causes of Failure Demand and write the standard. Find the constraint and put an agent there, built with the people who run it, with a person one step behind it. Then decide, before the first hour is freed, where that hour will go — and what each automated junior task used to teach.

REFERENCES

Acemoglu, D. and Restrepo, P. (2019) 'Automation and new tasks: how technology displaces and reinstates labor', *Journal of Economic Perspectives*, 33(2), pp. 3–30. Available at: https://doi.org/10.1257/jep.33.2.3.

Agrawal, A., Gans, J.S. and Goldfarb, A. (2024) 'Artificial intelligence adoption and system-wide change', *Journal of Economics & Management Strategy*, 33(2), pp. 327–337. Available at: https://doi.org/10.1111/jems.12521.

Autor, D.H. (2015) 'Why are there still so many jobs? The history and future of workplace automation', *Journal of Economic Perspectives*, 29(3), pp. 3–30. Available at: https://doi.org/10.1257/jep.29.3.3.

Bacchetta, M., Bekkers, E., Ganne, E. and Xu, A. (2025) 'Harnessing AI for inclusive growth', *eTrade for all*, 29 September (summary of the WTO *World Trade Report 2025*). Available at: https://etradeforall.org/node/14380 (Accessed: 14 October 2026).

Bain, J.S. (1956) *Barriers to New Competition: Their Character and Consequences in Manufacturing Industries*. Cambridge, MA: Harvard University Press.

Barney, J. (1991) 'Firm resources and sustained competitive advantage', *Journal of Management*, 17(1), pp. 99–120.

Bessen, J. (2020) 'Industry concentration and information technology', *Journal of Law and Economics*, 63(3), pp. 531–555. Available at: https://doi.org/10.1086/708936.

Bicheno, J. and Holweg, M. (2016) *The Lean Toolbox*. 5th edn. Buckingham: PICSIE Books.

Bloom, N., Sadun, R. and Van Reenen, J. (2012) 'Americans do IT better: US multinationals and the productivity miracle', *American Economic Review*, 102(1), pp. 167–201.

Bock, A.C. and Frank, U. (2021) 'Low-code platform', *Business & Information Systems Engineering*, 63(6), pp. 733–740. Available at: https://doi.org/10.1007/s12599-021-00726-8.

Bresnahan, T.F., Brynjolfsson, E. and Hitt, L.M. (2002) 'Information technology, workplace organization, and the demand for skilled labor: firm-level evidence', *Quarterly Journal of Economics*, 117(1), pp. 339–376.

Brynjolfsson, E. (2022) 'The Turing Trap: the promise and peril of human-like artificial intelligence', *Daedalus*, 151(2), pp. 272–287.

Brynjolfsson, E. and Hitt, L.M. (2000) 'Beyond computation: information technology, organizational transformation and business performance', *Journal of Economic Perspectives*, 14(4), pp. 23–48. Available at: https://doi.org/10.1257/jep.14.4.23.

Brynjolfsson, E., Hui, X. and Liu, M. (2019) 'Does machine translation affect international trade? Evidence from a large digital platform', *Management Science*, 65(12), pp. 5449–5460. Available at: https://doi.org/10.1287/mnsc.2019.3388.

Brynjolfsson, E., Jin, W. and McElheran, K. (2021) 'The power of prediction: predictive analytics, workplace complements, and business performance', *Business Economics*, 56(4), pp. 217–239. Available at: https://doi.org/10.1057/s11369-021-00224-5.

Brynjolfsson, E., Li, D. and Raymond, L. (2025) 'Generative AI at work', *Quarterly Journal of Economics*, 140(2), pp. 889–942. Available at: https://doi.org/10.1093/qje/qjae044.

Brynjolfsson, E., Malone, T.W., Gurbaxani, V. and Kambil, A. (1994) 'Does information technology lead to smaller firms?', *Management Science*, 40(12), pp. 1628–1644. Available at: https://doi.org/10.1287/mnsc.40.12.1628.

Brynjolfsson, E., Rock, D. and Syverson, C. (2021) 'The productivity J-curve: how intangibles complement general purpose technologies', *American Economic Journal: Macroeconomics*, 13(1), pp. 333–372. Available at: https://doi.org/10.1257/mac.20180386.

CFPB (2023) *Chatbots in Consumer Finance*. Issue Spotlight. Washington, DC: Consumer Financial Protection Bureau, June. Available at: https://files.consumerfinance.gov/f/documents/cfpb_chatbot-issue-spotlight_2023-06.pdf (Accessed: 14 October 2026).

Cohen, W.M. and Levinthal, D.A. (1990) ‘Absorptive capacity: a new perspective on learning and innovation’, *Administrative Science Quarterly*, 35(1), pp. 128–152. Available at: https://doi.org/10.2307/2393553.

Cunningham, W. (1992) ‘The WyCash portfolio management system’, in *Addendum to the Proceedings on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA ’92)*. New York: ACM, pp. 29–30. Available at: https://doi.org/10.1145/157710.157715.

Daly, C. (2025) ‘Klarna slows AI-driven job cuts with call for real people’, *Bloomberg*, 8 May. Available at: https://www.bloomberg.com/news/articles/2025-05-08/klarna-turns-from-ai-to-real-person-customer-service (Accessed: 14 October 2026).

Davenport, T.H. (1998) ‘Putting the enterprise into the enterprise system’, *Harvard Business Review*, 76(4), pp. 121–131.

Deming, W.E. (1982) *Out of the Crisis*. Cambridge, MA: Massachusetts Institute of Technology, Center for Advanced Engineering Study. Page references are to the MIT Press reissue (2000).

DeStefano, T., Kneller, R. and Timmis, J. (2025) ‘Cloud computing and firm growth’, *Review of Economics and Statistics*, 107(6), pp. 1638–1651. Available at: https://doi.org/10.1162/rest_a_01393.

Devine, F. (2023) *Rapid Mass Engagement: Driving Continuous Improvement through Employee Culture Creation*. New York: McGraw Hill.

Dunbar, R.I.M. (1993) ‘Coevolution of neocortical size, group size and language in humans’, *Behavioral and Brain Sciences*, 16(4), pp. 681–694. Available at: https://doi.org/10.1017/S0140525X00032325.

Euronews (2024) ‘Klarna’s AI bot is doing the work of 700 employees. What will happen to their jobs?’, *Euronews Next*, 28 February. Available at: https://www.euronews.com/next/2024/02/28/klarnas-ai-bot-is-doing-the-work-of-700-employees-what-will-happen-to-their-jobs (Accessed: 14 October 2026).

European Commission (2003) ‘Commission Recommendation of 6 May 2003 concerning the definition of micro, small and medium-sized enterprises (2003/361/EC)’, *Official Journal of the European Union*, L 124, 20 May. See also: https://single-market-economy.ec.europa.eu/smes/sme-fundamentals/sme-definition_en (Accessed: 14 October 2026).

Eurostat (2025) ‘Use of artificial intelligence in enterprises’, *Statistics Explained*. Luxembourg: Eurostat (data extracted December 2025). Available at: https://ec.europa.eu/eurostat/statistics-explained/index.php/Use_of_artificial_intelligence_in_enterprises (Accessed: 5 October 2026).

Evans, P.B. and Wurster, T.S. (1997) ‘Strategy and the new economics of information’, *Harvard Business Review*, 75(5), pp. 70–82.

Forman, C., Goldfarb, A. and Greenstein, S. (2012) ‘The Internet and local wages: a puzzle’, *American Economic Review*, 102(1), pp. 556–575. Available at: https://doi.org/10.1257/aer.102.1.556.

Gale, B.T. (1994) *Managing Customer Value: Creating Quality and Service That Customers Can See*. New York: Free Press.

Gartner (2024) ‘Gartner survey finds 64% of customers would prefer that companies didn’t use AI for customer service’, press release, 9 July. Available at: https://www.gartner.com/en/newsroom/press-releases/2024-07-09-gartner-survey-finds-64-percent-of-customers-would-prefer-that-companies-didnt-use-ai-for-customer-service (Accessed: 14 October 2026).

Goldratt, E.M. and Cox, J. (1984) *The Goal: A Process of Ongoing Improvement*. Croton-on-Hudson, NY: North River Press.

Hammer, M. (1990) ‘Reengineering work: don’t automate, obliterate’, *Harvard Business Review*, 68(4), pp. 104–112.

Hannan, M.T. and Freeman, J. (1984) ‘Structural inertia and organizational change’, *American Sociological Review*, 49(2), pp. 149–164.

He, Q. and Zumbühl, Y. (2026) ‘Lessons from Chinese business leaders about integrating AI in the workplace’, *World Economic Forum Stories*, 27 August. Available at: https://www.weforum.org/stories/artificial-intelligence/lessons-china-ai-workplace/ (Accessed: 8 October 2026).

Huang, M.-H. and Rust, R.T. (2021) ‘A strategic framework for artificial intelligence in marketing’, *Journal of the Academy of Marketing Science*, 49(1), pp. 30–50. Available at: https://doi.org/10.1007/s11747-020-00749-9.

Imai, M. (1986) *Kaizen: The Key to Japan’s Competitive Success*. New York: McGraw-Hill.

Ingka Group (2023) *AI and remote selling bring IKEA design expertise to the many*. Ingka Group Newsroom, 29 June. Available at: https://www.ingka.com/newsroom/ai-and-remote-selling-bring-ikea-design-expertise-to-the-many/ (Accessed: 29 September 2026).

Jin, W. and McElheran, K. (2017) ‘Economies before scale: survival and performance of young plants in the age of cloud computing’, *Rotman School of Management Working Paper*, No. 3112901. Available at: https://doi.org/10.2139/ssrn.3112901.

Kano, N., Seraku, N., Takahashi, F. and Tsuji, S. (1984) ‘Attractive quality and must-be quality’, *Journal of the Japanese Society for Quality Control*, 14(2), pp. 39–48 (in Japanese).

Korinek, A. and Vipra, J. (2025) ‘Concentrating intelligence: scaling and market structure in artificial intelligence’, *Economic Policy*, 40(121), pp. 225–256. Available at: https://doi.org/10.1093/epolic/eiae057.

Latzko, W. J., & Saunders, D. M. (1995). *Four Days with Dr. Deming: A Strategy for Modern Methods of Management*. Addison-Wesley.

Liker, J.K. (2004) *The Toyota Way: 14 Management Principles from the World’s Greatest Manufacturer*. New York: McGraw-Hill.

Löfgren, M., Witell, L. and Gustafsson, A. (2011) ‘Theory of attractive quality and life cycles of quality attributes’, *The TQM Journal*, 23(2), pp. 235–246. Available at: https://doi.org/10.1108/17542731111110267.

Matzler, K. and Hinterhuber, H.H. (1998) ‘How to make product development projects more successful by integrating Kano’s model of customer satisfaction into quality function deployment’, *Technovation*, 18(1), pp. 25–38. Available at: https://doi.org/10.1016/S0166-4972(97)00072-2.

McElheran, K. *et al.* (2024) ‘AI adoption in America: who, what, and where’, *Journal of Economics & Management Strategy*, 33(2), pp. 375–415. Available at: https://doi.org/10.1111/jems.12576.

Moffatt v. Air Canada (2024) BCCRT 149. Decision of the British Columbia Civil Resolution Tribunal, 14 February. Available at: https://decisions.civilresolutionbc.ca/crt/crtd/en/525448/1/document.do (Accessed: 14 October 2026).

Nooteboom, B. (1994) ‘Innovation and diffusion in small firms: theory and evidence’, *Small Business Economics*, 6(5), pp. 327–347. Available at: https://doi.org/10.1007/BF01065137.

Noy, S. and Zhang, W. (2023) ‘Experimental evidence on the productivity effects of generative artificial intelligence’, *Science*, 381(6654), pp. 187–192. Available at: https://doi.org/10.1126/science.adh2586.

OECD (2021) *The Digital Transformation of SMEs*. OECD Studies on SMEs and Entrepreneurship. Paris: OECD Publishing. Available at: https://doi.org/10.1787/bdb9256a-en.

OECD (2025) *Generative AI and the SME Workforce: New Survey Evidence*. Paris: OECD Publishing. Available at: https://doi.org/10.1787/2d08b99d-en.

The Open University (n.d.) ‘Is some demand unnecessary?’, in *Capacity and Demand Management*. OpenLearn, section 12. Available at: https://www.open.edu/openlearn/society-politics-law/criminology/capacity-and-demand-management/content-section-12 (Accessed: 14 October 2026).

Porter, M.E. (1980) *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. New York: Free Press.

Repenning, N.P. and Sterman, J.D. (2001) ‘Nobody ever gets credit for fixing problems that never happened: creating and sustaining process improvement’, *California Management Review*, 43(4), pp. 64–88. Available at: https://doi.org/10.2307/41166101.

Retail Gazette (2023) ‘Ikea embraces remote interior design as AI transforms sales tactics’, *Retail Gazette*, 13 June. Available at: https://www.retailgazette.co.uk/blog/2023/06/ikea-remote-interior-design/ (Accessed: 29 September 2026).

Retail Gazette (2026) ‘Ikea owner Ingka to cut 800 roles as CEO says it has “grown too complex”’, *Retail Gazette*, 20 March. Available at: https://www.retailgazette.co.uk/blog/2026/03/ikea-ingka-cut-roles-too-complex/ (Accessed: 29 September 2026).

Sculley, D. *et al.* (2015) ‘Hidden technical debt in machine learning systems’, *Advances in Neural Information Processing Systems*, 28, pp. 2503–2511.

Seddon, J. (2003) *Freedom from Command and Control: A Better Way to Make the Work Work*. Buckingham: Vanguard Press.

Thong, J.Y.L. (1999) ‘An integrated model of information systems adoption in small businesses’, *Journal of Management Information Systems*, 15(4), pp. 187–214. Available at: https://doi.org/10.1080/07421222.1999.11518227.

Treacy, M. and Wiersema, F. (1993) ‘Customer intimacy and other value disciplines’, *Harvard Business Review*, 71(1), pp. 84–93.

Varian, H. (2019) ‘Artificial intelligence, economics, and industrial organization’, in Agrawal, A., Gans, J. and Goldfarb, A. (eds) *The Economics of Artificial Intelligence: An Agenda*. Chicago: University of Chicago Press, pp. 399–419.

Womack, J.P. and Jones, D.T. (1996) *Lean Thinking: Banish Waste and Create Wealth in Your Corporation*. New York: Simon & Schuster.

Xu, F.F. *et al.* (2025) ‘TheAgentCompany: benchmarking LLM agents on consequential real world tasks’, *Advances in Neural Information Processing Systems*, 38 (Datasets and Benchmarks Track). Available at: https://arxiv.org/abs/2412.14161.