APPLIED FUTURES RESEARCH
REALITY REGISTER · REA · 007

The Crew and the Machine

Workforce Response as the Hidden Variable in Automation ROI

Prepared by Kevin Nussbaum / Applied Futures Research
August 2026

FOURTH PAPER IN THE ROBOTICS DEPLOYMENT SERIES

REA·004 *The Humanoid as Capital Asset — the asset thesis.* **REA·005** *The Pilot Trap — deployment reality.* **REA·006** *Own, Lease, or Rent the Robot — capital structure.* **REA·007** *this paper — the human factor.*

NOTE ON METHOD
Epistemic Tagging and Source Discipline

Claims in this paper carry one of four tags. The tag is not decoration; it is a statement about what would have to be true for the claim to be wrong, and it tells the reader how much weight the claim can bear.

Tag	What it means
GROUNDED	Reported in a primary source, peer-reviewed study, or published survey with a stated method and sample. Wrong only if the source is wrong.
ANALYTIC	Derived by the author from stated assumptions and arithmetic. The logic is auditable; the inputs are estimates. Change the inputs and the output changes.
PROJECTION	A forward statement about conditions that have not yet obtained. Directionally defensible, numerically soft.
SPECULATIVE	Reasoned conjecture offered because the question matters and the evidence is thin. Held loosely and labelled as such.

Two additional disciplines apply throughout. First, where a widely circulated statistic does not survive scrutiny, this paper says so and declines to use it, even where the number would have been rhetorically convenient. Two such numbers appear in §2.3 and §4.1. Second, where the paper reaches across domains — for example, applying findings from relationship research to a shop floor — the inference is tagged ANALYTIC rather than GROUNDED, because the study populations are not shop-floor crews and the extrapolation is the author's, not the researchers'.

ABSTRACT
The Denominator Nobody Underwrites

Every capital model for an automation deployment rests on an assumption about the duty cycle: the machine will run the hours it was designed to run. That assumption is not a property of the machine. It is a property of the crew around it — the operators who stage its work or let it starve, the technicians who move it up or down the maintenance queue, the supervisors who route exceptions through it or around it, and the veterans who either hand over their recovery knowledge or keep it. Workforce response is therefore not a soft factor sitting outside the spreadsheet. It is an unpriced liability inside the denominator, and in the model developed here it converts a twenty-one-month payback into a fifty-six-month payback without a single act of overt resistance.

This paper argues that the standard remedy fails for a diagnosable reason. Deployment programs treat resistance as a psychology problem and answer it with communication. But two distinct phenomena wear the same uniform on a shop floor. One is an uncalibrated threat response, and communication is the right remedy for it. The other is an accurate reading of a real personal distribution — the firm holds a portfolio and can absorb variance, while the worker holds a single position and cannot — and communication applied to that fear does not reduce resistance. It converts resistance from open to covert, which is strictly worse for the asset, because covert resistance cannot be managed by a supervisor who cannot see it.

The argument proceeds in eight steps. §2 establishes what the loss-aversion literature actually supports and where the popular version outruns the evidence. §3 develops the paper's central distinction between calibrated and uncalibrated fear. §4 explains why competently executed change management fails when it lands on the wrong one. §5 sets out a six-mechanism taxonomy of utilization leakage with observable signatures. §6 examines the German robot-adoption record as a natural experiment in incumbent protection. §7 prices the liability in a sixty-month model. §8 specifies the commitment architecture that changes a worker's actual distribution rather than his perception of it. §9 hands the buyer a checklist of what the vendor's model omits.

The conclusion is not that firms should be kind to their crews, though there are reasons for that which have nothing to do with payback periods. The conclusion is narrower and harder: the commitment that buys adoption is expensive, it belongs in the capital model as a line item rather than in the HR budget as an afterthought, and once it is in the model the payback math changes. Deployments that fail are, with some regularity, deployments that were approved on a number that had quietly assumed the commitment away.

§ 1
The Variable That Is Not in the Model

A humanoid or mobile-manipulation cell is specified, quoted, and justified on rated productive hours. The vendor's model takes a duty cycle — two shifts, two hundred and fifty days, some allowance for changeover and planned maintenance — multiplies it by a displaced labor rate, subtracts an operating cost, and returns a payback period. Every number in that chain is defensible. The chain is nonetheless incomplete, because rated hours are a design property and realized hours are a social outcome.

The gap between the two has a name in the reliability literature. Overall equipment effectiveness decomposes realized output into availability, performance, and quality losses, and every serious plant tracks it. What OEE does not do — cannot do, by construction — is attribute those losses to their upstream causes. A cell that is starved because the forklift operator prioritized a legacy line records an availability loss. A cell that runs at reduced rate because nobody transferred the parameter knowledge that prevents a recurring jam records a performance loss. In both cases the metric is accurate and the diagnosis is absent. The number lands in the denominator and the reason stays on the floor.

1.1 · Who is actually holding the duty cycle

Consider what has to happen for a cell to hit its rated hours on any given Tuesday. Work has to be staged in front of it before it runs dry, which is somebody's discretionary sequencing decision. Exceptions — the part presented at the wrong angle, the tote that arrived short, the tolerance that drifted — have to be resolved at the cell rather than routed around it, which is somebody's judgement about whether it is faster to fix the robot's problem or to do the job by hand. When a robot faults, it has to reach the top of the maintenance queue, which is somebody's ranking against six other work orders. And the knowledge of how to recover it in ninety seconds rather than forty minutes has to have moved out of one veteran's head into three or four others, which is a transfer nobody can compel.

Every one of those is a human decision, and every one of them is invisible in the capital model. None of them requires bad faith to go the wrong way. A supervisor who is neutral about the robot but under pressure on a legacy line will starve the robot, and will be right to, given his incentives. This is the first thing worth being precise about: workforce-mediated utilization loss does not require sabotage, hostility, or even disapproval. It requires only that nobody's interests are served by the machine running well.

1.2 · What the workforce is currently telling us

The 2026 survey record on this point is unusually consistent, and it points in a direction that should concern anyone underwriting a deployment. ADP Research's Today at Work program, drawing on a global workforce survey of more than thirty-nine thousand workers across thirty-six markets, found that only twenty-two per cent of workers strongly agreed their job was safe from elimination. The distribution by rank is the part that matters here: eighteen per cent among individual contributors and twenty-one per cent among frontline managers, rising to thirty-one per cent among upper managers and thirty-five per cent in the C-suite. **[GROUNDED]**

22%

Feel secure

of workers globally strongly agree their job is safe from elimination — ADP Research, 39,000+ workers, 36 markets

18 → 35%

Confidence by rank

feeling secure, individual contributors to C-suite. The people nearest the machine are the least confident

43%

Two-year horizon

expect automation to replace their job within two years — up five points year on year, ManpowerGroup 2026

Three other 2026 datasets triangulate it. Mercer's Global Talent Trends survey of twelve thousand workers and leaders put the share fearing their job would be made obsolete by AI at forty percent, against twenty-eight percent two years earlier. ManpowerGroup's Global Talent Barometer recorded forty-three per cent expecting automation to replace their job within two years, a five-point rise in a single year, alongside a striking pairing: eighty-nine percent were confident they had the skills for the job they hold today, and confidence in using new technology fell eighteen percent. Qualtrics found frontline workers materially less optimistic about AI than knowledge workers — thirty-nine percent hopeful or excited against fifty-three — and attributed the gap not to awareness but to experience, on the grounds that frontline workers have watched automation reshape warehouses and production floors before. **[GROUNDED]**

Read together, these describe a workforce that is confident in its present competence and unconfident in its future position. That is precisely the profile that produces high adoption friction, because the thing at risk is not the worker's ability to do the new job. It is his standing in a place he has already invested in.

> *A cell can be starved to fifty-five percent of its rated duty cycle by people who are individually reasonable, locally correct, and entirely within their job descriptions. Nothing in a purchase order prevents this. Nothing in an OEE dashboard explains it.*
>
> § 1.2

§ 2
The Asymmetry Engine

The intuition behind workforce resistance is old and roughly correct: people weigh what they might lose more heavily than what they might gain. It is worth getting the actual state of that evidence right, partly because it is more interesting than the popular version and partly because the popular version contains two numbers that should not be used.

2.1 · What the loss-aversion literature supports

Tversky and Kahneman's 1992 formulation of cumulative prospect theory reported a loss-aversion coefficient of $\lambda = 2.25$ — losses weighted roughly two and a quarter times an equivalent gain.

That figure entered management writing as a constant, which it is not. In 2024 two meta-analyses of the same construct arrived at different answers. Brown, Imai, Vieider and Camerer assembled six hundred and seven estimates from a hundred and fifty articles across economics, psychology and neuroscience and reported a mean λ of 1.955, with a ninety-five per cent interval of 1.820 to 2.102. Walasek, Mullett and Stewart took a stricter route, re-estimating parameters from the raw data of seventeen studies rather than pooling published coefficients, and reported a median λ of 1.31, interval 1.10 to 1.53. [GROUNDED]

A 2025 re-analysis by Yechiam and Zeif is the most useful of the three for present purposes. Working through the Brown corpus with moderators drawn from the methodological critiques, they found that strong loss aversion replicates in some conditions and attenuates in others — notably, it holds when losses are smaller than gains and when gains and losses are presented in an ordered fashion and weakens under symmetric conditions. Their earlier work made the same point more directly: pronounced loss aversion emerges for large losses, and not for all losses. [GROUNDED]

2.2 · The useful reframe: λ is a regime, not a constant

The apparent disagreement in that literature dissolves once λ is treated as a function rather than a number. The laboratory paradigm that produces the contested estimates is a small monetary gamble: bounded, quantified, impersonal, and repeatable. Under those conditions the asymmetry is modest and hard to pin down — which is exactly what one would expect if the mechanism were calibrated to consequence rather than to arithmetic. Losing eleven dollars in a mixed lottery is not a threat. It is an accounting entry.

Push the stake up in magnitude and in personal relevance and the asymmetry steepens. That is the finding the large-loss analyses converge on, and it has a direct implication for deployment work: the coefficient that governs how a crew weighs an automation announcement is not the laboratory coefficient. It is whatever λ obtains when the stake is a job, a skill built over fifteen years, a shift schedule that fits a childcare arrangement, and a position in a status hierarchy the worker has been climbing since he was hired. Any change-management effort budgeted against a two-to-one intuition is under-resourced by an unknown but material factor. [ANALYTIC]

Figure 1 — λ is a regime, not a constant

Laboratory gambles cluster near 2:1 and even that is contested. The asymmetry steepens as the stake grows large, personal and relational.

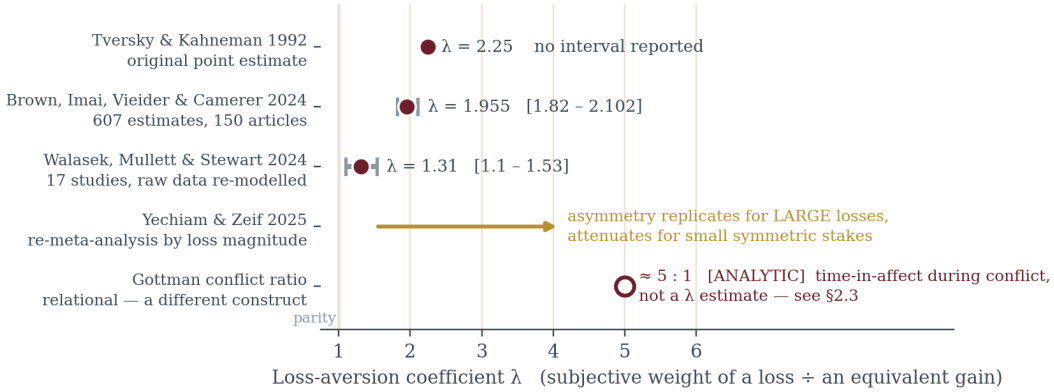

Figure 1 — Published λ estimates with intervals, plus the two constructs commonly conflated with them. The Gottman marker is included to show where it sits relative to the λ literature, not to imply it belongs to it.

2.3 · The relational coefficient, and one number to retire

A different literature handles the case where the interaction is relational rather than transactional, and it reports a steeper asymmetry. Baumeister, Bratslavsky, Finkenauer and Vohs surveyed the evidence in 2001 under a title that doubled as the finding: bad is stronger than good. Their review found the greater power of bad events over good ones in everyday events, major life events, close-relationship outcomes, social networks, interpersonal interactions, and learning — and noted that bad impressions and bad stereotypes form faster and resist disconfirmation more stubbornly than good ones. That last clause is the operationally important one for a plant manager. **[GROUNDED]**

Gottman's much-cited five-to-one ratio belongs to this literature, and it is worth stating what it actually measured, because the popular version has drifted. The original finding concerned time spent in emotional states during conflict conversations: stable couples maintained roughly five seconds in a positive or neutral state for every second in a negative one. The widely repeated translation — that five positive interactions are needed to offset one negative — is a reading of that result, not the result itself. The direction is well supported; the tidy arithmetic is a paraphrase. **[GROUNDED]**

And one number should be retired outright. Fredrickson and Losada's 2005 claim of a critical positivity ratio of 2.9013 was shown by Brown, Sokal and Friedman in 2013 to rest on a misapplication of differential equations borrowed from fluid dynamics. American Psychologist issued a formal withdrawal of the modelling section and of the specific ratios in September of that year; Losada did not defend the model. The number nevertheless continues to circulate in management and coaching literature. This paper does not use it, and a reader who encounters it in a vendor's change-management deck has learned something useful about the deck. **[GROUNDED]**

Regime	Character of the stake	Asymmetry
Transactional	Bounded monetary gamble; impersonal; repeatable; no implication for standing.	λ contested, roughly 1.3–2.1 depending on method. [GROUNDED]

Regime	Character of the stake	Asymmetry
Large-loss transactional	Monetary but consequential; magnitude sufficient to register as a threat rather than an entry.	Asymmetry replicates and steepens. [GROUNDED]
Relational	Trust, standing, and the reliability of a counterparty inferred from accumulated evidence.	Materially steeper; bad forms faster and resists disconfirmation. [GROUNDED]
Automation announcement	All three at once: income, skill, schedule, standing — delivered as a signal about the employer relationship.	Governed by the relational regime, not the laboratory one. [ANALYTIC]

There is a rational core to the relational steepness that is easy to miss when it is framed as bias. In a repeated game, a defection is more diagnostic than a cooperation. Cooperations are cheap and consistent with several different underlying dispositions; a defection narrows the field. A worker who has watched an employer behave well for eleven years and then sees one announcement that contradicts the pattern is not being irrational when he reweights. He is updating on the informative observation. This matters for the argument that follows, because it means the steep coefficient is not simply an obstacle to be managed around. Part of it is signal processing working as designed.

§ 3

Two Fears in One Uniform

This is the section the rest of the paper depends on. The claim is that shop-floor resistance to automation comprises two phenomena that are behaviorally similar, causally different, and responsive to different remedies — and that the failure to distinguish them accounts for a large share of deployment disappointment.

3.1 · The portfolio and the position

Start with the structural asymmetry, because it is not psychological at all. A firm evaluating an automation program is running an expected-value calculation across a portfolio. It may deploy forty cells across six sites over four years. Some will overperform, some will disappoint, one will be written off. The average carries the program, and variance at the individual-cell level is absorbable by construction. That is what a portfolio is for.

The worker has no portfolio. He has one job, one skill stack, one commute, one mortgage, and one shift pattern that the household has been organized around. If the deployment produces a ninety-five per cent probability of a better outcome for the plant and a thirty per cent probability that his particular role is consolidated, those two numbers are not in tension and neither of them is wrong. They are statements about different distributions. The firm is drawing many times from a favorable distribution. He is drawing once from a different one.

Figure 2 — The same deployment, two different distributions

A positive expected value at the plant level is fully compatible with a concentrated, correctly anticipated loss on the floor.

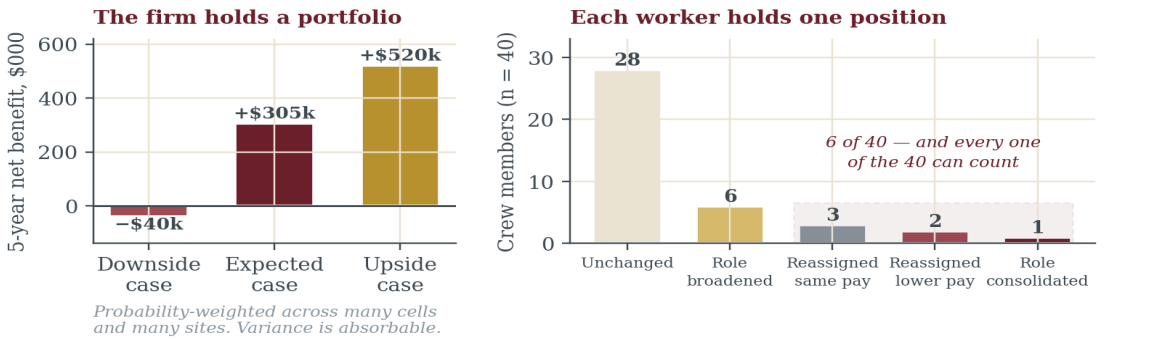

Figure 2 — Illustrative distributions for a single deployment. [ANALYTIC] The left panel is the firm's decision object; the right panel is the crew's. Note that the concentrated outcomes on the right are not a tail risk to the individuals in them — for those six people they are the whole distribution.

Two consequences follow. The first is arithmetic: gains from automation accrue to a legal entity and losses concentrate on a small number of identifiable people. The second is informational, and it is the one managers consistently underestimate. On a floor of forty, everyone can count. They know which two roles are thin, they know who is closest to the cell, and they will have worked out the likely consolidation candidates before the announcement meeting ends. There is no version of this conversation in which the crew is less informed about its own headcount than the person addressing it.

3.2 · Type I and Type II

With that structure in place the two fears can be separated cleanly.

Type I — uncalibrated threat response. The worker's actual distribution is unaffected or improved by the deployment, but he registers the change as a threat anyway. The veteran welder whose work the cell cannot touch resists it because change itself reads as loss; the negativity asymmetry of §2.3 is doing the work, and it is doing it in excess of the facts. This is a genuine psychology problem, and psychology remedies address it: exposure to the machine, early involvement in cell design, transparency about scope, visible competence on the part of the integrator. Type I fear responds to information because information is what it lacks.

Type II — calibrated threat assessment. The worker's actual distribution is worse after the deployment, and he has read it correctly. His role is one of the ones that consolidates, or his overtime disappears, or his skill premium is retired, or his shift moves. Type II fear is not a cognitive error and does not lack information. It is a correct inference from available evidence, and no amount of communication will alter the inference, because the inference is true.

The operational problem is that from the supervisor's chair, Type I and Type II look identical. Both present as questions in meetings, slow uptake, a preference for the manual method, and a general reluctance that is difficult to point at. Both are attributed to the same cause — 'the crew is resistant' — and both receive the same treatment. The following table offers a working diagnostic; it is offered as a heuristic, not an instrument.

Signal	Suggests Type I	Suggests Type II	How to test it
Who is resisting	Distributed across the crew, including people the cell cannot affect.	Concentrated in roles adjacent to the automated task.	Map resistance onto the task inventory. If the map has a shape, it is Type II.
What is asked in meetings	Questions about the technology — reliability, safety, whether it will work.	Questions about consequence — headcount, bidding rights, overtime, seniority.	Log the questions verbatim for two weeks and sort them.
Response to detail	Eases as detail improves. More specificity reduces the objection.	Hardens as detail improves, because detail confirms the inference.	Publish the task-level scope and watch which direction the temperature moves.
Response to reassurance	Accepted; the worker was looking for a reason to relax.	Read as evasion, and costs the speaker credibility for the next announcement.	If reassurance lowers trust, the fear was calibrated.
Behavioral expression	Open — voiced, argued, sometimes loudly. Visible and therefore manageable.	Increasingly covert once open objection has been answered with reassurance.	Watch for the objections stopping while utilization stays flat.

> *The last row is the whole problem in one line. Type I resistance that has been handled goes quiet because it has resolved. Type II resistance that has been handled goes quiet because it has learned that speaking is unproductive. Both look like progress on a dashboard.*
>
> § 3.2

§ 4
Why Competent Change Management Still Fails

4.1 · A statistic this paper will not use

Almost every discussion of this topic opens with the claim that seventy per cent of change initiatives fail. It is worth dealing with that number before going further, because it will not appear again. Hughes reviewed five published instances of the claim in the Journal of Change Management in 2011 and traced them to their origins — Hammer and Champy's 1993 reengineering work, Beer and Nohria's 2000 assertion in Harvard Business Review, Kotter's later estimate, and two others — finding in each case an absence of valid and reliable empirical evidence for the specific figure. Subsequent syntheses have made the further point that reported failure rates are not comparable across studies unless definitions, denominators, and time horizons are specified, and that much academic repetition of the number rests on secondary citation. **[GROUNDED]**

The reason to raise this is not pedantry. It is that the seventy per cent figure has a specific rhetorical function: it establishes change as inherently, mysteriously difficult, which is precisely the framing that lets a failed deployment be filed under a general truth about human nature instead of investigated. If failure is a background rate, nobody has to identify a mechanism. This paper's position is the opposite. Deployment disappointment has mechanisms, they are enumerable, and §5 enumerates six of them.

4.2 · The misapplied remedy

The change-management toolkit is not defective. Early involvement, transparent scoping, training that builds competence rather than compliance, recognition for operators who master the new system, embedded technical support through the learning curve rather than only at go-live — these are sound practices and the industrial literature endorses them for good reason. They are also, almost without exception, instruments for reducing Type I fear. They supply information, exposure, and agency to people whose objection is grounded in uncertainty.

Applied to Type II fear, the same instruments do not merely fail to help. They actively damage the thing they were deployed to protect. Consider the standard messaging: the robot is here to take the dull and dangerous work, not your job; automation changes roles rather than eliminating them. Delivered to a crew in which three roles will in fact be consolidated, that message is heard by three people who know it is false about them and by thirty-seven who know those three people. The message does not reduce their fear. It relocates the object of their fear from the machine to the management, and it converts a specific worry into a general one about whether anything they are told is reliable.

Figure 3 — Why competently executed change management still fails

The remedy is not wrong; it is misapplied. Only the lower-left cell destroys value — and it is the cell most programmes land in.

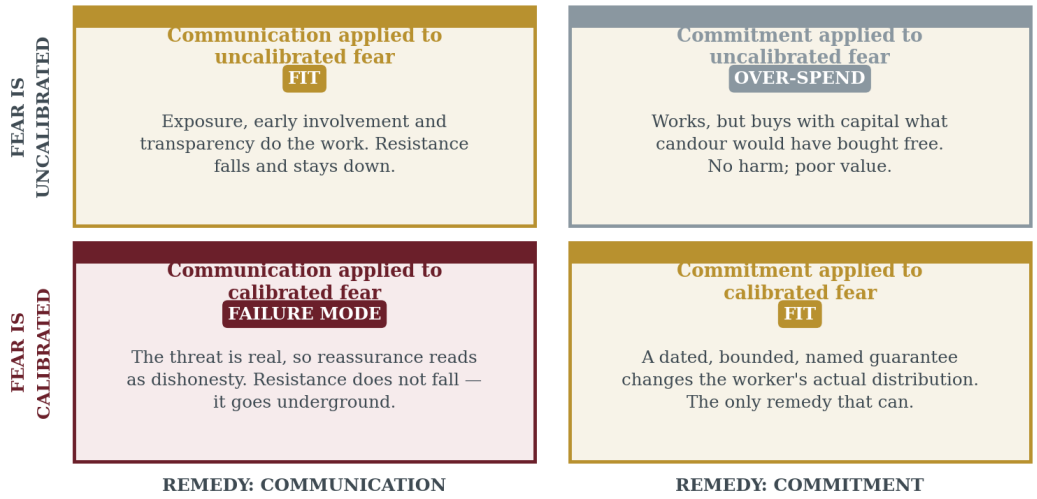

Figure 3 — Remedy fit against fear type. [ANALYTIC] Three of four cells are benign; the lower-left cell is where value is destroyed, and it is the cell a well-executed communication program lands in whenever the underlying fear was calibrated.

4.3 · The conversion problem

The specific harm in the lower-left cell deserves naming, because it is not simply 'the program didn't work.' It is a conversion of resistance from a visible to an invisible form, and the invisible form is more expensive.

Open resistance has a peculiar operational virtue: it is legible. A supervisor who is told to his face that the crew thinks the cell is a mistake knows what he is dealing with, knows who to talk to, and can escalate an accurate description of the problem. Open resistance is also self-limiting, because it is socially costly to sustain and because it invites a response. What it is not, is quiet.

Covert resistance produces no such signal. It looks like compliance. The objections stop, the training attendance is complete, the survey scores improve, and the cell runs at sixty-one per cent of rated hours for reasons that each have an innocent local explanation. The forklift was needed elsewhere. The part presented oddly so it was done by hand, as anyone would. The work order was queued behind two others that were more urgent. Nobody lied, nobody refused, and the asset is underwater. This is the failure mode that generates post-mortems attributing the shortfall to integration difficulty, because integration difficulty is the only category available that fits the observable facts.

4.4 · Job hugging, and why exit is not the release valve

There is a further 2026 finding that sharpens this. ManpowerGroup's Barometer recorded sixty-four per cent of workers planning to stay with their current employer — a pattern the report labels job hugging — alongside the rising automation fear already cited, and alongside fifty-six per cent reporting no recent training and fifty-seven per cent reporting no mentorship. [GROUNDED]

The interaction is unfavorable. In a loose labor market, a worker whose position is threatened has an exit option, and a meaningful share of Type II resistance leaves the building rather than staying to

express itself. In a market where workers judge exit to be risky, the same fear stays inside the organization and has to be expressed through the only channels remaining, which are discretionary effort and information sharing. The tighter the worker's perceived exit, the more of his threat response is routed through the duty cycle of the asset standing next to him. A deployment planned during a period of high job-hugging should expect a higher covert-resistance load than the same deployment planned five years earlier, holding the crew constant. **[ANALYTIC]**

§ 5
The Failure Taxonomy: Six Signatures of Utilization Leakage

If workforce response is to be managed as a capital risk rather than a mood, it has to be decomposed into mechanisms with observable signatures. Six recur. They are presented with their characteristic misdiagnosis, because the misdiagnosis is what determines whether the money spent in response does any good.

Figure 4 — Where the duty cycle actually goes [ANALYTIC]

Six workforce-mediated leakage paths in an unmanaged deployment. Not one of them appears as a line item; all of them appear in the denominator.

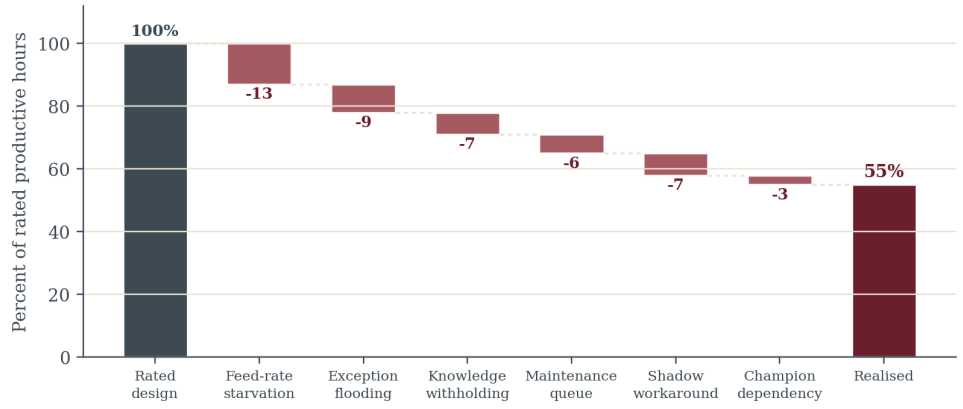

Figure 4 — Illustrative decomposition of the gap between rated and realized utilization in an unmanaged deployment. [ANALYTIC] The individual magnitudes are the author's estimates; the mechanisms and their signatures are drawn from the deployment and tacit-knowledge literature.

Mechanism	Observable signature	Usually misdiagnosed as	Drag	Fear type
Feed-rate starvation	Cell idle waiting on staged work; upstream resources consistently allocated elsewhere at the margin.	Layout problem; buffer sizing; 'islands of automation'.	−13 pts	Type II, and rational given local incentives
Exception flooding	High share of edge cases resolved manually and never fed back; exception log grows but never shortens.	Immature perception stack; vendor capability shortfall.	−9 pts	Either; distinguishable by whether fixes are proposed
Knowledge withholding	Recovery takes forty minutes with one operator and four with another; nothing written down.	Training gap; documentation debt.	−7 pts	Type II — the knowledge is the job security
Maintenance deprioritisation	Cell work orders consistently ranked behind legacy assets of lower economic value.	Maintenance capacity constraint.	−6 pts	Type I or institutional inertia
Shadow workaround	A parallel manual path is retained past the point of need and quietly preferred under time pressure.	Change-management shortfall; insufficient training.	−7 pts	Either; the path itself is the tell
Champion dependency	Utilization tracks one operator's shift pattern with a visible weekly signature.	Deployment success — until the champion transfers.	−3 pts	Neither; a structural single point of failure

5.1 · Knowledge withholding deserves separate treatment

Five of the six are straightforward. The third is not, because it sits on top of a problem that predates automation entirely and that automation makes acute.

Polanyi's formulation from 1958 remains the cleanest statement of it: we know more than we can tell. The knowledge that keeps a line running is substantially tacit — judgement calls, recovery sequences, the feel for a drifting parameter, the workaround for a fixture that has been slightly wrong for six years. It is not in the standard operating procedure because it was never articulable enough to put there, and much of it is held by operators who could not fully explain it if they wanted to. The automation literature has converged on this as a primary cause of pilots that work in a fenced corner and fail on the floor: the deployment automates the official process rather than the real one and breaks on the first exception the real process had a name for.

Two things then compound. First, the tacit knowledge that would rescue the deployment is held precisely by the people whose position the deployment threatens, and being the only person who can recover a critical asset is a rational job-security strategy. Withholding here does not require malice; it requires only the absence of a reason to volunteer. Second, a knowledge-capture program launched immediately after an automation announcement is transparently an extraction exercise, and will be read as one. The window for building the documentation that makes a deployment survivable is before the announcement, not after — which is an inconvenient sequencing requirement, because before the announcement there is no budget line for it. [ANALYTIC]

> *Ask a veteran to write down what he knows in the same month you tell him his role is under review, and you have asked him to hand over the only asset he holds that you cannot replace. He will comply, partially, and the part he keeps will be the part that matters.*
>
> § 5.1

§ 6
The German Natural Experiment

The argument to this point has been mechanical. It is worth testing against the largest available body of evidence on what actually happens when robots meet a workforce, and there is a useful comparison available, because two advanced economies automated heavily over the same period under different rules about incumbent protection.

The American record is well established. Acemoglu and Restrepo, using variation in industry-level robot penetration across commuting zones, found robust negative effects on both employment and wages, with one additional robot per thousand workers reducing the employment-to-population ratio by roughly 0.2 percentage points and wages by about 0.42 per cent. Their working-paper estimates spanned 0.18 to 0.34 points and 0.25 to 0.5 per cent respectively. They also established that the effect was distinct from Chinese import competition, offshoring, and other IT capital. [GROUNDED]

The German record over a comparable period looks different in a specific and instructive way. Dauth, Findeisen, Suedekum and Woessner examined industrial robot adoption in German labor markets from 1994 to 2014 and found no effect on total employment in locales specializing in robot-intensive industries: manufacturing job losses were offset by gains in business services. At the individual level, robot adoption did not increase displacement risk for incumbent manufacturing workers. Those workers stayed with their original employers, and many adjusted by switching occupations at the same workplace. Job stability actually rose. The authors attributed the smaller displacement effect — on the order of half the American magnitude — to legislative firing costs, strong unions and works councils, and a skilled workforce that could be retrained. [GROUNDED]

Figure 7 — The trade-off, made visible [GROUNDED]

Protecting incumbents halved displacement and eliminated the net employment effect. The cost moved to wage growth — and to workers not yet hired.

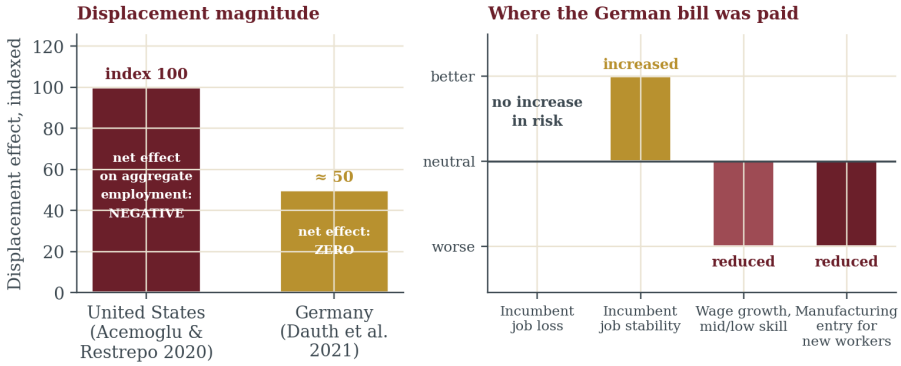

Figure 7 — Displacement magnitude and incidence, United States against Germany. [GROUNDED] Sources as cited in §6; the index on the left panel expresses the German displacement effect relative to the American, per the authors' own comparison.

6.1 · Reading the result honestly

It would be easy to stop there and declare that worker protection is free — that Germany got the productivity and kept the jobs. The authors did not stop there, and neither should a capital allocator. The same study found that the increased job stability came at the cost of lower wage growth, with negative earnings effects for the bulk of medium-skilled and low-skilled workers whose task profiles were most substitutable, while managers and high-skilled workers engaged in abstract tasks benefited. And the equilibrium loss of manufacturing jobs was driven entirely by fewer new jobs for young labor-market entrants.

That is the trade-off in its most legible form, and it is worth stating plainly. Protecting incumbents halved the displacement effect and eliminated the net employment loss. The bill did not disappear; it moved. Part of it was paid by the protected workers themselves, in wage growth foregone. The rest was paid by people who were never in the room — the twenty-two-year-olds who would have entered manufacturing and instead did not, and who had no representation on any works council because they had not been hired. **[GROUNDED]**

> *Every decision is a trade-off. The German record does not show that worker protection is free. It shows that the cost of protection is smaller than the cost of unmanaged displacement, and that it lands on different people.*
>
> § 6.1

For the purposes of this paper the German case establishes three things. First, workforce response is not a fixed constant of human nature; it is a function of the institutional commitment surrounding the deployment, and it can be moved. Second, moving it works — the productivity gains were realized, and the adoption happened. Third, the commitment has a price, it is paid by identifiable parties, and any presentation that describes it as costless is not describing the German case.

§ 7
Pricing the Unpriced Liability

The claim that workforce response belongs in the capital model rather than the HR budget is only useful if it can be priced. What follows is a sixty-month model of a single deployment under four workforce scenarios. It is entirely illustrative: the arithmetic is auditable, the inputs are estimates, and a reader who disagrees with an input should substitute his own and re-run it. That is the point of publishing the assumptions.

7.1 · The deployment

Parameter	Value	Basis
Units deployed	2	Single material-handling and machine-tending cell.
Hardware cost	$190,000	Two platforms at $95,000, consistent with the mid-market band established in REA·004 and REA·006.
Integration, safety survey, fixturing, MES tie-in	$110,000	Author's estimate; integration commonly runs 40–70 per cent of hardware in first deployments.
Total deployed capital	**$300,000**	**Sum of the above. Financed on balance sheet; the lease and RaaS alternatives are treated in REA·006.**
Rated annual gross benefit	$216,000	Approximately 2.9 FTE-equivalents of task time at a fully burdened $36.00/hour.
Baseline annual operating cost	$46,000	Service contract, spares, power, orchestration software license.
Rated payback	**21.2 months**	**$300,000 ÷ ($216,000 − $46,000), annualized. This is the number in the vendor's pro forma.**

7.2 · Four workforce scenarios

The scenarios differ in exactly one respect: what was promised to the crew and whether it was kept. Hardware, integration, duty-cycle design, and labor rates are identical throughout.

Pro forma. Rated utilization is realized in full. No workforce cost is carried because none was contemplated. This scenario is included because it is what was approved, not because it is attainable.

Scenario A — Unmanaged. The deployment is announced, training is delivered, and no commitment of any kind is made regarding headcount, redeployment, or role. The six leakage mechanisms of §5 operate at the magnitudes shown in Figure 4, taking realized utilization to fifty-five

per cent. Operating cost rises to $54,000 to reflect exception handling, elevated supervisor time, and unplanned maintenance. No workforce spend appears anywhere in the capital file.

Scenario B — Communicated. A competent, well-resourced change-management program is executed: early operator involvement in cell design, transparent scoping, structured training, recognition for operators who master the system, embedded support through the learning curve. One-time cost $35,000. Type I fear is substantially resolved and realized utilization reaches seventy-two per cent. Type II fear is untouched, because nothing in the program altered anyone's actual distribution.

Scenario C — Committed. Everything in Scenario B, plus a written commitment architecture per §8: attrition-only headcount reduction over a dated horizon, named landing roles identified before go-live for every affected worker, and a participatively designed gain-share on realized cell productivity. One-time cost $55,000; recurring gain-share pool $19,000 per year. Realized utilization reaches ninety-one per cent. It does not reach one hundred, because some leakage is structural rather than motivational.

Figure 5 — Payback is a workforce variable [ANALYTIC]

Same hardware, same integration, same duty-cycle design. The only variable that moves is what was promised to the crew — and kept.

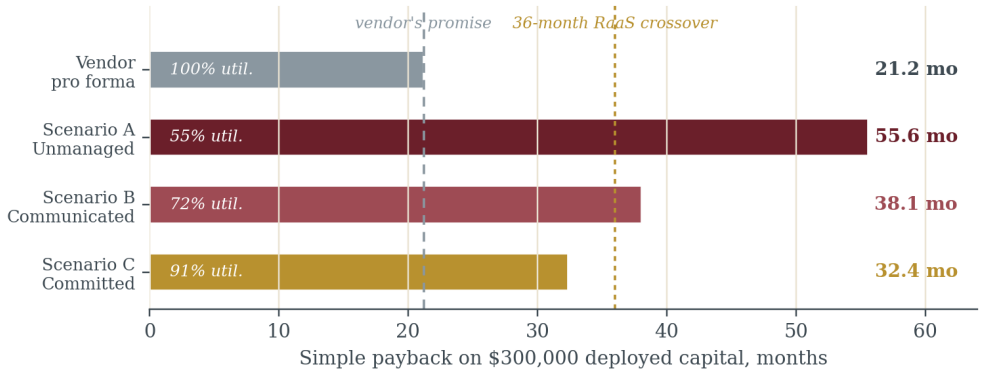

Figure 5 — Simple payback under the four scenarios. [ANALYTIC] The 36-month reference line is the buy-versus-subscribe crossover established in REA·006 and is shown because a payback beyond it changes the financing conclusion, not merely the return.

Line	Pro forma	A · Unmanaged	B · Communicated	C · Committed
Realized utilization	100%	55%	72%	91%
Gross annual benefit	$216,000	$118,800	$155,520	$196,560
Operating cost	($46,000)	($54,000)	($50,000)	($46,000)
Recurring gain-share	—	—	—	($19,000)
Net annual benefit	**$170,000**	**$64,800**	**$105,520**	**$131,560**
Deployed capital	$300,000	$300,000	$300,000	$300,000

Line	Pro forma	A · Unmanaged	B · Communicated	C · Committed
One-time workforce spend	—	—	$35,000	$55,000
Total capital at risk	**$300,000**	**$300,000**	**$335,000**	**$355,000**
Simple payback	**21.2 mo**	**55.6 mo**	**38.1 mo**	**32.4 mo**
Cumulative net at 60 months	**$550,000**	**$24,000**	**$192,600**	**$302,800**

Figure 6 — Sixty months of the same asset [ANALYTIC]

The workforce architecture, not the machine, sets the slope — and no scenario recovers the vendor's pro forma.

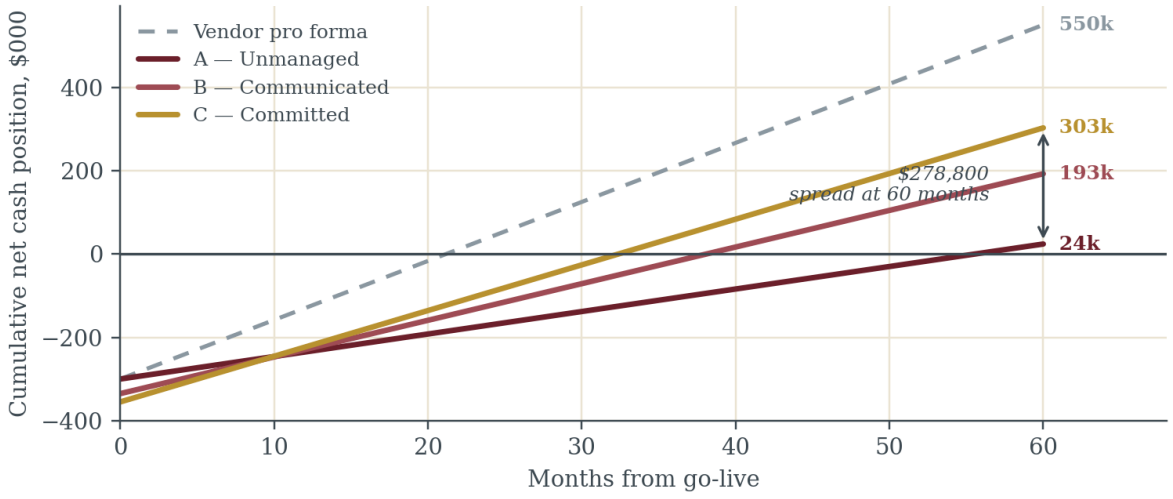

Figure 6 — Cumulative net cash position over sixty months. [ANALYTIC] The four lines describe the same machine. What separates them is a decision made before go-live about what the crew was told and whether it was binding.

7.3 · What the model says

23.2
Months of penalty
additional payback carried by an unmanaged crew versus a committed one, on identical hardware

$278,800
Sixty-month spread
difference in cumulative net position between Scenario A and Scenario C, on $300,000 of capital

1.86×
Return on commitment
additional net benefit per dollar of commitment architecture spend over sixty months

Four readings are worth drawing out. First, the workforce variable dominates. A twenty-three-month swing in payback on identical hardware is larger than the effect of most of the decisions that actually get debated in the approval meeting — vendor selection, financing structure, deployment sequencing.

Second, the commitment architecture pays for itself substantially: a hundred and fifty thousand dollars of workforce spend across sixty months returns roughly two hundred and seventy-nine thousand dollars of additional net benefit in this model, a ratio of about 1.86 to one. Third, and least comfortable, no scenario recovers the pro forma. The twenty-one-month payback was never available. It was a number produced by assuming a duty cycle that no crew delivers. [ANALYTIC]

The fourth reading is the one that should change behavior. In Scenario A the deployment is not a disaster on any single day. It clears break-even at fifty-six months and ends the five-year period marginally positive. It will be reported as a modest success by anyone measuring against zero rather than against the pro forma, and the twenty-three-month gap will be absorbed into a general observation that the technology was less mature than hoped. The mechanism will not be identified, because nobody instrumented it, and the next cell will be approved on the same pro forma.

§ 8
The Commitment Architecture

If Type II fear cannot be talked down, the only remaining lever is to change the distribution the worker is reading. That is a different class of intervention from communication: it costs money, it constrains management discretion, and it must be specific enough to be verifiable. The following five instruments have that property. None of them is novel; what is proposed here is that they be treated as capital-model line items rather than as HR programs.

Instrument	Mechanism	Addresses	Cost character
Attrition-only headcount commitment	A dated, bounded written undertaking that headcount reduction in the affected area occurs only through attrition and voluntary separation for a specified period. Bounded so it is credible; dated so it is falsifiable.	Type II	Carries redundant labor cost during the overlap window. The largest and most real cost in the package.
Named landing roles	Every worker whose task is automated has a specific identified role, with pay and shift stated, before go-live. Not a retraining budget — a named destination.	Type II	Training and temporary double-staffing; modest if planned, expensive if retrofitted.
Participatively designed gain-share	A share of realized cell productivity returns to the crew, with the crew involved in designing the distribution rule. Aligns the operator's interest with the duty cycle directly.	Both	Recurring and self-limiting: it only pays out when the asset performs.
Operator ownership of the cell	The cell is assigned to named operators who hold recovery authority, tuning authority, and public credit for its performance. Converts the machine from a threat to a platform for standing.	Type I	Near zero in cash; real in management attention and in ceding control.

Instrument	Mechanism	Addresses	Cost character
Pre-announcement knowledge capture	Documentation of tacit recovery and parameter knowledge undertaken before any automation announcement, funded as an operating improvement in its own right.	Type II	Must be spent before there is a project to charge it to — the hardest item to get approved.

8.1 · Why the process matters independently of the outcome

The organizational-justice literature supplies a finding that is easy to dismiss as soft and is in fact load-bearing here. Employee reactions to gain-sharing are mediated by perceptions of distributive, procedural, and interactional justice — that is, by judgements about the fairness of the outcome, of the process that produced it, and of the manner of its communication. Work on gain-sharing effectiveness has found fairness and participation to be central to whether the scheme produces its intended behavioral effects, and a case meta-analysis of thirty-three gain-sharing implementations spanning fifty years identified employee involvement in plan design, formal involvement structures, and employee favorability toward the plan as predictors of success — robust across organizational size, union status, and technology. [GROUNDED]

The implication for deployment design is sharper than it first appears. A gain-share designed by finance and announced to the crew is a different instrument from an arithmetically identical gain-share designed with the crew, and the difference is not sentiment. Participation in the design is what makes the commitment credible, because a rule the crew helped write is a rule the crew can predict the application of. Credibility is the entire mechanism: an incredible commitment does not change the worker's read of his distribution, and therefore does not buy any utilisation, while costing exactly as much as a credible one. [ANALYTIC]

8.2 · What does not work

For completeness, and because these are the items most commonly funded:

Reassurance without commitment. 'Automation changes roles, it doesn't eliminate them' spends credibility and buys nothing whenever it is not literally true for everyone hearing it. See §4.2.

Retraining without a destination. A training budget is not a landing role. A worker offered upskilling with no named position at the end of it correctly reads it as being handed the cost of his own transition.

The vendor's change-management module. Vendors are competent at Type I work and structurally incapable of Type II work, because the only instruments that address Type II are commitments the buyer must make about the buyer's own headcount. No vendor can sell that, and a vendor who implies otherwise is selling the Type I module twice.

Undated commitments. 'No one will lose their job because of this' has no expiry, no boundary, and no test, and is therefore heard as rhetoric. 'No involuntary separations in this department through 31 December 2028' is a smaller promise and a vastly more valuable one, because it can be checked.

§ 9
What the Vendor's Model Omits

The register's next paper is a buyer's field guide for evaluating robotics vendors. The workforce items belong in it and are set out here as a checklist. Each is a question that a mid-market buyer can ask in a first or second meeting, and each has a diagnostic answer.

01 **What duty cycle does this pro forma assume, and what did comparable customers realize in months one through twelve?** A vendor who can answer the second half has real deployment history. A vendor who repeats the rated figure is quoting a specification, not an outcome.

02 **Which specific tasks come out of which specific job codes?** If the answer is at the level of 'material handling', the scope has not been worked out and the crew will work it out first.

03 **How many people are affected, and are any of their roles consolidated?** Answer this internally before the first crew meeting. The crew will have an answer within a day of the announcement whether management supplies one or not.

04 **Who owns the exception backlog, and what is the target for shrinking it?** An exception log with no owner and no burn-down target is the signature of Scenario A.

05 **What is the recovery time for the ten most likely faults, and who besides the champion can execute it?** One name is a single point of failure. Three names is a deployment.

06 **Where does the cell sit in the maintenance priority ranking, and who set that ranking?** A high-value asset queued behind low-value legacy equipment is a governance defect, not a capacity problem.

07 **What tacit knowledge does this cell depend on, and is it documented?** If the answer is no, the documentation work belongs before the announcement, funded separately. See §5.1.

08 **What is the written, dated headcount commitment, and who signed it?** Undated is rhetorical. Unsigned is aspirational. Both are read correctly by the crew.

09 **For each affected worker, what is the named landing role, its pay, and its shift?** Retraining budgets are not answers to this question.

10 **Does anyone on the floor gain financially if the cell runs well?** If the honest answer is nobody, the duty-cycle assumption in the pro forma has no support.

11 **What is the workforce line item in this capital request?** A capital request with a zero in that field has not priced the liability; it has assumed it away.

12 **How will realized utilization be reported, to whom, and against which baseline?** Reported against zero, Scenario A is a success. Reported against the pro forma, it is a twenty-three-month miss with an identifiable cause.

§ 10
The Capital Allocator's Verdict

Dimension	Unmanaged	Communicated	Committed
What it costs before go-live	Nothing.	$35,000, one-time.	$55,000 one-time plus $19,000 per year.
What it actually buys	Nothing. The Type I fear it never addressed and the Type II fear it never could.	Resolution of Type I fear. Real, worth paying for, and insufficient on its own.	A change in the crew's actual distribution — the only lever that reaches Type II.
Realized utilization	55%	72%	91%
Payback against a 21.2-month pro forma	55.6 months — a 34-month miss, past the RaaS crossover.	38.1 months — a 17-month miss, still past the crossover.	32.4 months — an 11-month miss, and the only case that competes with subscribing.
How it will be reported	'The technology was less mature than we hoped.'	'Adoption took longer than planned.'	'The cell performed close to plan.'
What it does to the next deployment	Nothing learned; the same pro forma is approved again.	Partial learning; the residual gap stays unexplained.	A defensible baseline for the second cell and a crew that has seen a commitment kept.
Verdict	**Value-destroying relative to available alternatives. Do not fund a deployment in this configuration.**	**Necessary, insufficient, and the most common configuration in the field.**	**The only configuration in which the asset thesis of REA·004 survives contact with a shop floor.**

Closing · Every decision is a trade-off

The temptation in a paper like this is to end on the reconciliation: manage the people well and everybody wins. The evidence does not support that ending, and the German record in §6 is the reason. Protecting incumbents there halved the displacement effect, eliminated the net employment loss, and delivered the productivity gains. It also cost the protected workers wage growth, and it cost a cohort of young entrants the manufacturing jobs they would otherwise have had. The trade-off was not eliminated. It was moved, and it was moved onto people with less standing to object.

The honest formulation is narrower than a reconciliation and more useful. A trade-off is rational when the party bearing the cost is the party collecting the benefit. When those separate — when the gains accrue to a legal entity and the losses concentrate on six identifiable people who can all count — what presents as a psychology problem is an incidence problem, and no quantity of understanding the human threat response will resolve it. The crew's read is not a bias to be corrected. It is a valuation, and it is frequently more accurate than the pro forma it is resisting.

Which leaves the capital allocator with a decision that is straightforward to state and expensive to act on. The commitment that buys adoption is real money, it must be spent before there is any evidence it was necessary, and it belongs in the capital request rather than in a change-management line nobody reads. The alternative is not avoiding the cost. It is paying it in the denominator, over fifty-six months, without ever knowing what it was for.

> *The next paper in this series, REA·008 — 'The Skeptic's Field Guide' — turns from the deployment to the purchase: how to evaluate robotics vendors when you have never bought a robot.*

REFERENCES
Sources, Grouped by Epistemic Weight

Peer-reviewed studies and published surveys with stated methods anchor §2, §6 and §8. Trade and practitioner literature is used only for mechanism description in §1 and §5, never for magnitudes. All figures in §7 are the author's and are tagged ANALYTIC in the text.

Loss aversion and the asymmetry literature

1. Tversky, A. & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 5, 297–323. Source of the λ = 2.25 point estimate.
2. Brown, A. L., Imai, T., Vieider, F. M. & Camerer, C. F. (2024). Meta-analysis of empirical estimates of loss aversion. Journal of Economic Literature. 607 estimates from 150 articles; mean λ = 1.955, 95% interval [1.820, 2.102].
3. Walasek, L., Mullett, T. L. & Stewart, N. (2024). A meta-analysis of loss aversion in risky contexts. Journal of Economic Psychology. Raw-data re-estimation across 17 studies; median λ = 1.31, 95% interval [1.10, 1.53].
4. Yechiam, E. & Zeif, D. (2025). Loss aversion is not robust: A re-meta-analysis. Journal of Economic Psychology. Re-analysis of 84 papers, 163 estimates, n = 149,218, with moderators by loss magnitude and presentation order.
5. Baumeister, R. F., Bratslavsky, E., Finkenauer, C. & Vohs, K. D. (2001). Bad is stronger than good. Review of General Psychology, 5, 323–370.
6. Gottman, J. M. (1994). Why Marriages Succeed or Fail. Simon & Schuster. Source of the conflict-interaction ratio; see §2.3 on what the ratio measures and what it does not.
7. Brown, N. J. L., Sokal, A. D. & Friedman, H. L. (2013). The complex dynamics of wishful thinking: The critical positivity ratio. American Psychologist, 68, 801–813. With the American Psychologist correction of September 2013 formally withdrawing the modelling section of Fredrickson & Losada (2005).

Automation, employment and incidence

8. Acemoglu, D. & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6). One additional robot per thousand workers: employment-to-population −0.2 pp, wages −0.42 per cent.
9. Acemoglu, D. & Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper 23285. Estimate ranges of −0.18 to −0.34 pp and −0.25 to −0.5 per cent, and the discrimination of robot effects from trade and other IT capital.
10. Dauth, W., Findeisen, S., Suedekum, J. & Woessner, N. (2021). The adjustment of labor markets to robots. Journal of the European Economic Association, 19(6), 3104–3153. German evidence 1994–2014: no net local employment effect, no increased displacement risk for incumbents, increased job stability, lower wage growth for medium- and low-skilled workers, and manufacturing job loss concentrated in reduced entry for young workers.
11. Dauth, W., Findeisen, S., Suedekum, J. & Woessner, N. (2018). Adjusting to Robots: Worker-Level Evidence. Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute Working Paper 13, and the authors' ECB/CEPR presentation of December 2018, which sets German displacement at roughly half the American magnitude and attributes the difference to firing costs, unions and works councils, and retrainability.

Workforce sentiment, 2026

12. ADP Research (2026). Today at Work, Issue 1, drawing on the ADP Global Workforce Survey of more than 39,000 workers across 36 markets. 22 per cent strongly agreed their job was safe from elimination; 18 per cent of individual contributors, 21 per cent of frontline managers, 23 per cent of middle managers, 31 per cent of upper managers, 35 per cent of C-suite executives.
13. Mercer (2026). Global Talent Trends 2026, survey of 12,000 workers and business leaders. 40 per cent fear AI will make their job obsolete, against 28 per cent in 2024.
14. ManpowerGroup (2026). Global Talent Barometer. Regular AI use up 13 points to 45 per cent; confidence in using technology down 18 per cent; 89 per cent confident in current-role skills; 43 per cent expect automation to replace

their job within two years; 64 per cent intend to stay with their employer; 56 per cent report no recent training; 57 per cent report no mentorship.

15. Qualtrics (2026). EX Trends Report. Frontline workers 39 per cent hopeful or excited about AI changing their work, against 53 per cent of knowledge workers; 53 per cent of frontline workers trust senior leaders to use AI responsibly.

Change management, justice and gain-sharing

16. Hughes, M. (2011). Do 70 per cent of all organizational change initiatives really fail? Journal of Change Management, 11(4). Reviews five published instances of the claim and finds no valid and reliable empirical evidence for the figure in any of them.

17. Beer, M. & Nohria, N. (2000). Cracking the code of change. Harvard Business Review; and Hammer, M. & Champy, J. (1993). Reengineering the Corporation. Cited as the traceable origins of the 70 per cent claim, not as support for it.

18. Cooper, C. L., Dyck, B. & Frohlich, N. (1992). Improving the effectiveness of gainsharing: The role of fairness and participation. Administrative Science Quarterly, 37, 471–490.

19. Kwon, S., Kim, M. S., Kang, S.-C. & Kim, M. U. (2008). Employee reactions to gainsharing under seniority pay systems: The mediating effect of distributive, procedural, and interactional justice. Human Resource Management, 47, 757–775.

20. Case meta-analysis of 33 gain-sharing implementations spanning fifty years, identifying employee involvement in plan design, formal involvement structures, and employee favorability as predictors of success, robust across organizational size, union status, technology and environment. Reported in the organizational-development and justice literatures; see also Welbourne, T. M. (1998), Untangling procedural and distributive justice, Group & Organization Management, 23(4).

21. Konovsky, M. A. (2000). Understanding procedural justice and its impact on business organizations. Journal of Management, 26(3).

Tacit knowledge and deployment mechanism

22. Polanyi, M. (1958). Personal Knowledge: Towards a Post-Critical Philosophy. University of Chicago Press. Source of the formulation that we know more than we can tell.

23. Trade and practitioner literature on tribal knowledge, hidden factories, and the automation of official rather than actual processes, used in §5 for mechanism description only. Magnitudes attributed to these sources are not used anywhere in this paper.

24. Practitioner literature on automation ROI and islands of automation, describing the capital-efficiency gap between theoretical machine capacity and realized throughput under OEE decomposition, and workforce resistance as a ceiling on realized capacity. Used for mechanism only.

Companion papers

25. Nussbaum, K. (2026). The Humanoid as Capital Asset: The Asset Thesis for Embodied AI. Applied Futures Research, REA·004.

26. Nussbaum, K. (2026). The Pilot Trap: Why Automation Deployments Stall Between Proof and Scale. Applied Futures Research, REA·005.

27. Nussbaum, K. (2026). Own, Lease, or Rent the Robot: Capital Structure for Assets With No Residual Value History. Applied Futures Research, REA·006.