THE INTERIOR ARCHITECTURE OF TRANSFORMATION

The Maybe Series

Ten articles on holding what you cannot yet resolve

JOSHUA T. FRASER, ED.D.

THE INTERIOR ARCHITECTURE OF TRANSFORMATION

The Maybe Series

CONTENTS

There is an old story, traced most reliably to the *Huainanzi* text of 139 BCE, about a farmer whose horse runs away. His neighbors gather at the fence. "What terrible luck," they say. The farmer looks up from his work. "Maybe," he says.

The next day the horse returns, bringing with it three wild horses. The neighbors gather again. "What wonderful luck," they say. The farmer wipes his hands on his trousers. "Maybe."

His son tries to ride one of the wild horses, falls, and breaks his leg. The neighbors arrive, as neighbors do. "What terrible luck." The farmer's answer does not change.

The following week, military officers come through the village conscripting young men for a war from which most will not return. They pass over the farmer's son because of his broken leg. The neighbors are back at the fence with their wonderful luck, and the farmer is still standing there with his one word, and two thousand years later we are still telling this story because we have not yet learned what it is trying to teach us.

That teaching is more urgent now than at any point in recent memory. And the urgency itself is part of the problem, because urgency is the enemy of the posture the farmer is modeling. Urgency wants conclusions. It wants us to sort, to decide, to pick a lane so we can start building in it. The farmer refuses. Not because he doesn't care, but because he has been alive long enough to watch the sorting go wrong more often than it goes right.

Alex Finn, writing in *Wired*, describes AI as the great democratizer, a technology that will put creative and analytical power into the hands of people who have never had access to it.^1^ The economists on the Work Futures panel at Davos describe the same technology as the most significant displacement event in the history of labor markets. They are both presenting evidence. They are both making predictions. They are both, if we are being honest, probably right about some things and wrong about others. And the most dangerous thing any of us could do right now is pick a side and stop paying attention to the other.

That last sentence sounds reasonable. It sounds like something a thoughtful person would say at a panel discussion before yielding to the next speaker. But I want to push on it, because living it is significantly harder than saying it. The discourse has sorted itself into camps with remarkable efficiency, and the camps have given themselves permission to stop listening. The enthusiasts are building while the critics are warning, and neither group is required to sit with the discomfort of the other's evidence, because each camp has constructed a narrative sufficient to explain away the inconvenient parts. Both camps have constructed explanations that protect their certainty, mirror images of each other, each selecting evidence that confirms what the camp already believes and treating the selection as comprehensiveness. The explanations are products of the camps, not of the evidence, and neither is honest about what it excludes.

Here is where it gets uncomfortable for me personally, and I think the discomfort matters enough to name. I build with these tools. I use Claude Code nearly every day. I have watched it do things that took me hours in a fraction of the time, and the feeling that produces is genuine, almost intoxicating. There are mornings when something that felt impossible yesterday is finished before breakfast, and in those moments the enthusiasts' narrative feels not just plausible but obvious. Of course this is good. Of course this changes everything. And then I catch myself, because "of course" is the sound a mind makes when it has stopped examining its own conclusions, and the farmer never says "of course." The farmer says "maybe," which is the harder word because it requires you to hold the good feeling and the unease at the same time without letting either one win.

The unease is real too. The same architecture producing that morning efficiency is producing something elsewhere that I have not fully mapped. Jobs are shifting. Creative work is being repriced. The question of who benefits and who absorbs the cost is not settled, and pretending it is settled, in either direction, is a move I recognize from other domains where I have watched leaders simplify their way into harm they could have anticipated if they had been willing to stay in the discomfort a little longer.

Carl Jung wrote about the shadow: the parts of ourselves we refuse to acknowledge because they contradict the story we need to believe.^2^ Every camp in the AI discourse has a shadow. The enthusiasts' shadow is the harm they are not examining because it would complicate their

narrative of progress. The critics' shadow is the genuine benefit they are not acknowledging because it would weaken their case for caution. The farmer has no shadow in this story. Or rather, his shadow is the cost of his posture: the loneliness of being the person who will not join the celebration or the mourning, the particular exhaustion of holding uncertainty when everyone around you has found the relief of a conclusion.

I want to be careful here, because the farmer's posture can be romanticized into wisdom from above or serene detachment or the pose of someone who has figured things out and is waiting for the rest of us to catch up. The farmer is none of these. The farmer is a person standing in a field, doing his work, watching events unfold, and refusing to let anyone else's certainty infect his own judgment. That refusal carries no comfort and little admiration in the moment. The farmer's "maybe" is what integration sounds like when feeling and thinking run as one event encountering a situation whose meaning has not finished arriving, and the word carries grief and understanding and uncertainty simultaneously because the channels are not taking turns. His neighbors have the warmth of shared conviction. They have the pleasure of being right together. The farmer has a field that needs planting whether the news is good or bad.

This is different from indifference, and the distinction matters. Indifference is a posture you adopt when the stakes are low. The stakes here are not low. The farmer is not disengaged. He is practicing the most rigorous form of engagement available, because his posture requires him to hold the full weight of complexity without setting any of it down. The neighbor with wonderful luck has set down the weight of caution. The neighbor with terrible luck has set down the weight of possibility. Both of them feel lighter. Both of them are less accurate than the farmer, who has not moved and has not simplified and has not allowed the pressure of a gathered crowd to compress reality into something more manageable than reality actually is.

I think about this when I read the latest round of predictions. Someone with credentials and a platform says AI will eliminate 40% of jobs within a decade, and the critics share it as confirmation, and the enthusiasts dismiss it as fearmongering, and both camps have processed the claim in under thirty seconds because they already know what they think. The farmer would read the same prediction and do something almost nobody does anymore: hold it. Not accept it. Not reject it. Hold it alongside the other evidence, alongside the morning when the tool did something remarkable, alongside the afternoon when it produced something careless, alongside the question of who is building these systems and for whom and at what cost. Holding all of that is heavy. Setting some of it down feels like clarity, but the farmer knows it is not clarity. It is relief, and relief is not the same thing.

Maybe.

That word, placed honestly, is the hardest single word in any conversation about the future. It refuses prediction. It refuses tribe. It refuses the comforting fiction that we already know enough to stop paying attention to evidence that contradicts our position. "Maybe" is not a punt. It is a discipline. It is the decision to keep all the evidence on the table when every social pressure you face is telling you to sort it into piles and

throw half of it away.

This is the posture these pieces are written from. Not above the debate. Not outside it. Inside it, using the tools, building the systems, doing the work, and refusing to let any of it become a sermon for either side, because the moment you start preaching from inside your experience you have stopped paying attention to it. These ten articles are an attempt to stay in the farmer's position for as long as the thinking requires: to examine AI through the lens of suspended judgment, to take the capabilities seriously without mythologizing them, to take the risks seriously without catastrophizing them, and to see what becomes visible when you refuse to let urgency collapse your field of view.

If that sounds like something worth exploring further, you are in the right place. If it sounds like indecision, I understand. The farmer heard that too, from the neighbors who always knew exactly how to feel.

He planted his field anyway.

^1^ Finn, A. (2024). "The Great Equalizer: How AI Tools Are Democratizing Creative Work." *Wired*.
^2^ Jung, C.G. (1951). *Aion: Researches into the Phenomenology of the Self*. Princeton University Press.

ARTICLE 2

The Tool and the Territory

Every significant technology creates a second landscape invisible to people focused on the first. The visible landscape is the tool itself: its features, its interface, what it can do and how fast it does it. The invisible landscape is what the tool exposes about the people using it, the organizations absorbing it, and the assumptions that were load-bearing only because nobody had tested them until now. Most conversations about artificial intelligence are still happening on the first landscape.

Alfred Korzybski gave us a sentence in 1931 that has outlived nearly everything else he wrote: "The map is not the territory."^1^ He meant that representations of reality are not reality itself, that the model is not the thing being modeled. Applied to AI, the insight sharpens: the tool is the map, and the territory is what it reveals about how you think, how you decide, how you organize, and what you have been avoiding. The people who will handle this moment well are not the ones who master the features. They are the ones who recognize what kind of ground they are standing on.

There are three territory shifts worth examining carefully.

The Expertise Shift

When a machine can produce a competent first draft of nearly anything, competence stops being the differentiator. You can ask an AI system to draft a strategy document, a communications plan, a research synthesis, a proposal outline, and the result will be structurally sound, appropriately sourced, and coherent enough to serve as a working starting point. Five years ago, the person who could produce that document efficiently was the expert in the room. Their expertise was the bottleneck, and the bottleneck conferred authority.

Now the bottleneck is gone. The drafting is handled. What remains is the part that was always more important but rarely valued on its own terms: the judgment to know whether the thing produced is right for this community, this moment, this particular confluence of needs that no training set has ever encountered.

Gary Klein's research on expert decision-making is instructive here. Klein spent decades studying how firefighters, nurses, and military commanders make high-stakes decisions under time pressure, and what he found was that expertise lives not in the ability to generate options but in the ability to recognize patterns and assess fit.^2^ The expert firefighter does not run through a decision tree. The expert firefighter walks into a burning building and knows, from sensory cues that resist explicit articulation, that the floor is about to give way. That recognition capacity is what AI cannot replicate, because it requires embodied experience in a specific context, accumulated over years, refined through failure. Judgment, not competence. The two have always been different; we just could not see the difference clearly because competence was scarce too, and competence was easier to measure.

Now competence is abundant. Judgment is exposed as the thing that actually matters. Most of our professional evaluation systems and credentialing structures are still measuring competence. The territory has shifted. The maps have not.

The Efficiency Trap

Everyone assumes that reclaimed time is automatically reclaimed well.

It is not. This is something you discover from the inside, not from the outside. In the early months of building serious automation, the time savings are real and they feel like victory. Hours freed from repetitive tasks, from formatting drudgery, from the mechanical labor of first-draft production. And then comes the honest accounting of what that freed time actually goes toward. The answer, more often than is comfortable to admit, is: more of the same category of work. Different tasks, but the same level of urgency and the same quality of attention. The automation makes it possible to run faster on a track that was already going the wrong

direction.

This is not a new pattern. Jevons identified it in 1865 when he observed that improvements in coal-burning efficiency did not reduce coal consumption; they increased it, because efficiency made coal-powered activity cheaper and therefore more attractive.^3^ The dynamic applies to cognitive work with uncomfortable precision. Save two hours on document production, and those two hours do not flow toward reflection, relationship-building, or strategic thinking. They flow toward the next two documents in the queue, because the environment you operate in is designed to absorb every freed minute into the next urgent demand.

Efficiency, it turns out, is not a direction. It is an accelerator. It makes you go faster, but it does not tell you where to go, and it does not stop you when you are pointed at the wrong thing.

The territory question is not "how do I save time?" It is "what structural changes do I need to make so that saved time flows toward the work that matters instead of being recaptured by the work that does not?" That is a systems question, a design question, a question about organizational architecture and incentive structures and the courage to leave the queue half-empty. No prompt will answer it.

The Authority Question

When AI can synthesize research faster than any individual, when it can pull from a broader corpus of knowledge than any human expert has read, what does it mean to be "the expert in the room"?

The question is uncomfortable, and it should be, because it lands differently depending on where your authority actually lived. If your authority was primarily informational, if being the person who had read the most and could retrieve it fastest was your primary value proposition, then the machine has something to say about your future. That is the discomforting truth at the center of this shift, and no amount of "AI is just a tool" reassurance dissolves it. For some roles, for some people, the ground has genuinely moved, and pretending otherwise is a kindness that helps no one.

But there is a second truth sitting next to the first, and it complicates the picture in ways that matter. If your authority was about something harder to name, something that lives in the space between information and action, between data and wisdom, between what the research says and what this particular situation actually needs, then the machine has not taken anything from you. It has clarified what was yours all along. Dreyfus and Dreyfus argued decades ago that expert knowledge is fundamentally non-propositional, that it resides not in rules and facts but in the practitioner's embodied capacity to perceive relevant features of a situation that novices cannot even see.^4^ AI operates entirely in the propositional domain. It can retrieve, synthesize, and articulate what has been written down. It cannot perceive what has not.

I suspect the honest answer for most people is both. Some portion of what you do is now replicable by machine, and some portion is not, and you probably do not have a clean accounting of which is which. That means the authority shift is going to be uncomfortable for a long time. The people who handle it well will be the ones who can sit in that discomfort without collapsing it into either "AI is making me irrelevant" or "AI cannot touch what I do." Both of those are maps. Neither is the territory.

Staying on the Ground

There is a strong temptation, when territory shifts like these appear, to theorize about them from above. To stand at a distance and produce frameworks, taxonomies, strategic roadmaps. The frameworks are not useless, but they share a limitation: they describe the territory from the perspective of someone who is not walking through it. The map gets refined. The walking does not happen.

Philip Jackson's ethnographic work in classrooms demonstrated that the most consequential dynamics in educational settings are not the ones visible from the administrator's office or the policy document; they are the ones visible only from inside the room, over time, through sustained presence.^5^ The same holds here. The expertise shift, the efficiency trap, the authority question: these are conditions to be navigated from within, and the navigation will teach you things the frameworks cannot.

I am not going to resolve any of these shifts here. I have named them, examined their contours, traced their implications as far as I can from where I sit. The territory is yours. Walk it.

^1^ Korzybski, A. (1933). *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. International Non-Aristotelian Library.

^2^ Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.

^3^ Jevons, W. S. (1865). *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*. Macmillan.

^4^ Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Free Press.

^5^ Jackson, P. W. (1968). *Life in Classrooms*. Holt, Rinehart and Winston.

ARTICLE 3

What If I'm Wrong About Both Things

What happens when your openness and your wariness are both telling the truth?

I have been sitting with that question for months, and it has not become easier. If anything, it has become heavier, because the evidence on both sides keeps accumulating with the kind of steady insistence that makes you suspect the universe is not planning to resolve this one for you.

Here is where I am. I use these tools. I build with them. I have watched something that would have taken me three full days compress into forty minutes of focused collaboration with a machine that does not get tired, does not lose the thread, does not need me to explain the same constraint twice. That compression is real. It changed what Tuesday looked like. It changed what was possible before the week started. And if I am being honest, it changed something about how I think about my own capacity, because the boundary between what I can do and what I can do with this is no longer where it used to be, and the new boundary is farther out, and that feels like something worth noticing.

It also feels like something worth being cautious about, because every technology that has ever expanded individual capacity has also, without exception, been pointed at people with less power and used to consolidate the advantage of people with more. That is not a prediction. That is a historical pattern so consistent it barely qualifies as observation. It is closer to gravity.

So I hold both. The compression is real and the pattern is real and they exist in the same moment, in the same tool, sometimes in the same keystroke. And the thing I cannot do, the thing that keeps me up on the nights when I am honest enough to stay with the discomfort instead of collapsing it into a position, is separate them. I cannot extract the benefit and leave the risk behind. I cannot point at the tool and say this is good or this is dangerous, because it is both, simultaneously, and the simultaneity is the part that nobody seems to want to talk about.

The discourse is sorted. I described this in the first piece, the neighbors at the fence with their wonderful luck and their terrible luck, and I meant it as a description of a problem, not a metaphor I planned to inhabit. But here I am, standing where the farmer stands, and it is lonelier than I expected. The enthusiasts have a community. The critics have a community. The people in the middle, the ones holding the full weight of the contradiction without setting either side down, mostly just have a headache and an inbox full of articles from both camps, each one so confident that it makes the uncertainty feel like a personal failure rather than an intellectual position.

I want to be clear: this is not performed uncertainty. I am not hedging so that I can claim prescience regardless of which direction things go. I genuinely do not know. When I watch what these tools can do for a first-generation college student who has never had a writing tutor, who has never had someone sit with their ideas and help shape them into language that carries the weight those ideas deserve, I feel something that I would be lying to call anything other than hope. When I watch the same architecture being used to surveil workers, to automate decisions about who gets resources and who does not, to accelerate the same deficit narratives that were already doing harm at human speed, I feel something I would be lying to call anything

The Efficiency Lie

In 2017, a McKinsey Global Institute report estimated that 60 percent of all occupations had at least 30 percent of their constituent activities that were "technically automatable."^1^ The figure was presented as opportunity. Productivity gains. Labor optimization. The word "efficiency" appeared 87 times in the report. The word "equity" appeared four times. The word "harm" did not appear at all.

That ratio tells a story. Not about McKinsey specifically, but about the conceptual architecture surrounding efficiency as a category. Efficiency is treated as a directional good, a force that moves organizations forward by definition. The assumption is so deeply embedded that questioning it sounds like arguing against progress itself. No one in a professional setting has ever been criticized for making something more efficient. The concept carries an automatic positive valence that insulates it from the scrutiny applied to every other organizational decision.

This insulation is the problem.

Efficiency is not a value. It is a multiplier. It amplifies whatever it is applied to, with complete indifference to what that thing is. If efficiency is applied to good work, the result is more good work done faster. If it is applied to surveillance, the result is surveillance at higher resolution and greater frequency. If it is applied to the same inequitable structures that were already producing harm, those structures begin operating with more speed and precision than their architects imagined possible. The multiplier does not evaluate direction. It has no capacity to. Direction is a human input, and the productivity narrative treats it as settled when it is not.

Shoshana Zuboff identified this pattern two decades before the current AI moment. In *The Age of Surveillance Capitalism*, she documented how efficiency logic, once decoupled from questions of purpose, becomes a mechanism for extracting value from human behavior without accountability to the humans whose behavior is being extracted.^2^ The tool works. The question of what it is working toward goes unasked, not because it is difficult but because the efficiency frame makes the question seem irrelevant. If something is faster, it is better. That is the axiom. And it is false.

Three patterns appear consistently when organizations adopt AI-driven tools at scale. These are not edge cases or cautionary outliers. They are the center of the adoption curve. They describe what happens most of the time, in most organizations, under normal conditions.

The first pattern involves data.

An organization implements AI-powered analytics capable of surfacing performance metrics in real time, disaggregated across every variable the system can identify. The tool is technically excellent. It does precisely what it was built to do. The problem is that the metrics it pulls are the same deficit-framed metrics the organization has been using for years: who is failing, who is behind, who is at risk, who is in the pipeline for consequences. The AI makes it possible to generate these reports hourly instead of monthly, which means conversations about people shift from "how is she doing this quarter" to "her numbers dropped since Tuesday."

The tool accelerated the frequency of deficit-focused surveillance without anyone pausing to ask whether deficit-focused surveillance was the practice that deserved acceleration.

Safiya Umoja Noble's work on algorithmic systems is instructive here. In *Algorithms of Oppression*, she demonstrates that automated systems do not introduce bias; they operationalize existing bias at scale, embedding it into infrastructure where it becomes harder to see and harder to challenge.^3^ The analytics dashboard does not create deficit thinking. It makes deficit thinking ambient, continuous, and architecturally invisible. The old quarterly report was a document someone had to produce, distribute, and explain. The real-time dashboard is just there, refreshing itself, generating a perpetual stream of deficit-coded information that shapes every conversation without anyone deciding it should.

Speed was never the problem. Direction was the problem. The direction was already wrong before the tool arrived. The tool made it faster.

The second pattern involves communication.

A network of teams adopts AI-generated messaging templates. The templates are well-crafted, consistently formatted, and available on demand. The problem is volume. Because communication is now cheap to produce, output triples. People who previously received two messages a week are now receiving six or seven. The messages, while polished, have lost something the handwritten note from a specific person used to carry: the signal that a human being was paying attention to a specific recipient and chose to spend time composing something for them alone.

Hartmut Rosa's sociological analysis of acceleration is relevant. Rosa argues that technological acceleration does not produce more time; it produces more tasks to fill the time that was freed.^4^ The efficiency gain in message production creates a relational loss in message reception, and the relational loss is invisible in every metric the organization tracks, because the metrics measure messages sent, open rates, response times. Not trust. Not the felt sense that someone on the other end of the communication actually knows who they are writing to. The metrics cannot capture what was lost because what was lost was never measured, and what is never measured does not exist inside an efficiency frame.

The volume increase is treated as evidence that the system is working. More messages, more engagement, more reach. By every available metric, communication has improved. By the experience of the people receiving it, something has been replaced by its simulation, and the simulation is convincing enough that the replacement goes unnoticed by everyone except the people on the receiving end. They notice. Their response is to stop reading.

The third pattern involves decision-making.

An organization uses AI to optimize scheduling and resource allocation across multiple sites. The algorithm reduces conflicts, balances workloads, and maximizes utilization. It is technically impressive. It also, because the algorithm optimizes for operational efficiency rather than equity, consistently routes the newest and least experienced personnel toward the highest-need situations. Those situations have the most flexibility in their parameters and are therefore easiest for the algorithm to fill. The algorithm did not intend to reproduce inequity. It was not designed with that goal. It optimized for the variable it was given, and the variable it was given reflected the same structural priorities that had been producing inequitable outcomes for years.

Virginia Eubanks documented this dynamic in *Automating Inequality*: systems designed to optimize for efficiency within existing structures do not challenge those structures; they reinforce them, encoding decades of inequitable practice into algorithmic logic that presents itself as neutral.^5^ The scheduling tool is neutral in the same way that a highway routed through a Black neighborhood is neutral. It solves a technical problem. The technical problem was defined by a set of priorities that were never examined, and the solution to the technical problem produces consequences that fall disproportionately on people who had no voice in defining the problem.

Faster.

Each pattern has the same architecture: a tool that functions exactly as designed, applied to a system that was already misaligned, producing more of what the system was already producing but at greater speed and resolution. The tool is not the failure. The direction is the failure. And direction is a human question that no technology answers, because the question is not technical. It is moral. It asks not "can this be done more efficiently" but "should this be done at all, and if so, for whom, and at whose expense."

The efficiency frame does not suppress this question. It does something worse. It makes the question feel unnecessary. If something is faster, cheaper, more scalable, the efficiency frame treats the case as closed. The burden of proof shifts to anyone who wants to slow down, and slowing down is coded as resistance, as Luddism, as failure to adapt. The person who asks "what is this actually for" is positioned as an obstacle to progress, which means the question gets asked less often as the tools get more powerful. The inverse relationship is precise: the more consequential the question becomes, the less likely it is to be raised.

Ruha Benjamin names this dynamic "the New Jim Code," the production of social inequality through the design of ostensibly neutral technologies.^6^ The neutrality is the mechanism. Because the tool is presented as value-free, the values embedded in its application become invisible. Efficiency is the master category that enables this invisibility, because efficiency, unlike justice or equity or care, does not require the practitioner to specify a beneficiary. Efficient for whom. Efficient toward what end. Efficient at whose expense. These questions are grammatically available but structurally discouraged. The efficiency frame is complete without them.

There is a version of this argument that ends with reassurance. A paragraph about how efficiency, properly directed, remains a powerful tool. A sentence about how the answer is not rejecting automation but ensuring it serves the right purposes. A closing metaphor, perhaps the farmer again, asking what is being planted before celebrating the speed of the planting.

That version is available. It is also, in its own way, another efficiency: the efficient resolution of discomfort, the rapid conversion of a disturbing observation into a manageable action item. The reader encounters a problem, receives a solution, and moves on. The discomfort lasted exactly as long as the argument required it to, and not one sentence longer.

The patterns described above are not problems awaiting solutions. They are outcomes of a logic that is functioning correctly. The efficiency frame is not broken. It is working. It is producing exactly what it was designed to produce: more of whatever already exists, faster, at greater scale, with less friction. The question of whether what already exists deserves to be amplified is outside the frame entirely, and the frame is the only thing most organizations know how to look through.

^1^ Manyika, J., Chui, M., Miremadi, M., et al. (2017). *A Future That Works: Automation, Employment, and Productivity*. McKinsey Global Institute.

^2^ Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.

^3^ Noble, S. U. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.

^4^ Rosa, H. (2013). *Social Acceleration: A New Theory of Modernity*. Columbia University Press.

^5^ Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press.

^6^ Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity Press.

ARTICLE 5

Both/And Is Harder Than It Sounds

Marcus tried. That is the part that matters, and the part that makes the rest of this harder to write.

He was a curriculum director in a mid-sized suburban district, the kind of leader who read widely and genuinely believed that holding complexity was a professional obligation. When his district began piloting AI tools for lesson planning, Marcus positioned himself as the person who could hold both sides. He told his team that AI was simultaneously powerful and dangerous, that the time savings were real and the costs were not yet visible. He said both/and at least twice in every meeting. He built a slide deck with "Both/And Thinking" in the header. He believed it.

And then his superintendent asked him to make a recommendation.

Not a both/and recommendation. A recommendation. Should the district expand the pilot or not? Should they invest in the platform or redirect the money to coaching? The superintendent did not want tension held beautifully. She wanted a decision by Thursday.

Marcus recommended expanding the pilot. He told himself the decision was both/and because it came with caveats: monitor for quality, protect planning time, sunset the tool if outcomes dropped. But the caveats were footnotes. The decision was binary. Expand or don't. And the moment he made it, the both/and collapsed into something that looked, from every angle that mattered, like a side.

His team noticed. The teachers who had concerns heard "expand" and stopped raising objections, because objections to an approved initiative feel like resistance rather than thought. The teachers who loved the tool heard "expand" and started evangelizing, because endorsement from leadership feels like permission to stop questioning. Within six weeks, the pilot had become a program, the caveats had become optional, and Marcus was defending a position instead of holding a tension.

He is not in the rest of this article. His story does not resolve.

The both/and is not a compromise. That is the most common misreading, and it does the most damage, because compromise feels sophisticated enough to pass for complexity.

Compromise lives in the middle. It takes two positions and locates the point equidistant from both, then calls the result balance. The both/and does not live in the middle. It lives in the tension between two things that are simultaneously true and seemingly incompatible, and it refuses to resolve that tension into something tidier. AI is democratizing creative capacity and commodifying human expertise. AI is reducing operational burden and intensifying surveillance culture. AI is accelerating transformative work and reproducing inequitable patterns at scale. Each pair is true at the same time, and if you resolve any by picking a side, you become a less accurate observer of what is happening.

The difficulty is not intellectual. Intellectually, both/and is easy. You can nod at it in a meeting, put it on a slide, use it as a discussion prompt. The difficulty is operational: the moment someone asks you to act, the both/and wants to collapse, because action moves in a direction, and direction feels like a side.

That collapse is where the slogan lives. And I am frustrated by how quickly people get there.

"Both/and" has become one of those phrases that signals thoughtfulness without requiring it, a verbal badge that says "I am holding complexity" while the person wearing it makes the same binary decisions everyone else makes, just with better framing language. I have watched it happen in rooms where I was the one facilitating, and I have watched it happen in my own thinking, which is worse. The phrase becomes a container for the feeling of sophistication, and the feeling substitutes for the practice, and by the time anyone notices the substitution, the decision has already been made and the both/and is decorative.

This is not a small problem. It is the central failure mode of every framework that asks people to hold tension rather than resolve it. The tension is genuinely uncomfortable, and the human nervous system is not built to sustain unpleasant states when resolution is available. Our neural architecture is a prediction machine designed to collapse ambiguity as rapidly as possible.^1^ For most of our evolutionary history, unresolved ambiguity signaled danger. The rustle in the grass is either a predator or the wind, and the humans who survived decided fast rather than standing there admiring the complexity while the thing in the grass made its own decision.

That architecture served us on the savanna. It serves us poorly now. The AI question is not predator-or-wind. It is both-at-the-same-time, and our brains want to collapse it because collapsed ambiguity feels like safety, even when the collapse produces a less accurate picture. The farmer's "maybe" is fighting this architecture directly. That is why the parable resonates across two thousand years: not because the farmer is wise, though he is, but because his posture requires a discipline that contradicts our default cognitive setting.

And most of us, when the superintendent calls, do not hold it.

Three practices help build the capacity. Not wanting the capacity, because wanting to hold both/and is like wanting to run a marathon; the wanting is irrelevant, the training is not.

The first is the pre-mortem. Before any significant AI decision, write two stories. One is the story of brilliant success: the problem solved, the time saved, the people you serve genuinely better off. The other is catastrophic failure: the tool reproduces what it was supposed to interrupt, the time savings evaporate into new demands, the thing that was automated turns out to have been holding something important together. Write both with equal conviction, then read them side by side and look for what they share. The shared element is almost always an assumption that both stories take for granted. That assumption is where the real decision lives.

This was tested when deciding whether to automate a complex data-gathering process. The success story assumed faster data would produce faster action. The failure story assumed automated gathering would bypass the sense-making conversations that give data its meaning. The shared assumption was that speed matters. When examined, speed was the wrong variable entirely. The value was in the conversations people had about the data together. Automating the collection was fine, but the timeline could not compress, because no algorithm compresses trust-building into a more efficient format.

The second practice is the shadow audit. It asks one question: What is my enthusiasm refusing to examine? Enthusiasm has a shadow. When the audit runs honestly, specific things surface: time saved and reinvested poorly, automation that worked technically but missed something relationally, systems built that were elegant and unnecessary because building felt more comfortable than sitting with the mess. Each is a place where enthusiasm functioned as defense against a harder truth.

The third is the Tuesday test. Whatever you believe about AI on Sunday night, after reading the latest piece and feeling the warm coherence of a narrative that explains everything, check whether it holds by Tuesday afternoon, in the actual work. Sunday night beliefs are conceptual; Tuesday afternoon beliefs are experiential, tested against actual friction and actual people and systems that do not behave the way the think piece predicted. If the Sunday belief survives Tuesday unchanged, you are not paying close enough attention. If it shatters completely, you are being too reactive. The both/and lives where the belief is modified but not abandoned.

Here is the discomfort I cannot resolve.

Every practice I just described can become another version of the collapse. The pre-mortem becomes a ritual that produces a predictable list and gets filed. The shadow audit becomes a performance of self-awareness that substitutes for change. The Tuesday test becomes a habit of mild revision that never touches the core commitment. Each practice is vulnerable to the same thing it was designed to interrupt,

because the human capacity for turning genuine discipline into comfortable routine is, as far as I can tell, infinite.

The people who hold the both/and most consistently are not the smartest or the most informed.^2^ They are the ones who have learned to tolerate the discomfort of not knowing while still showing up. They have a kind of earned steadiness that comes from having been wrong enough times that certainty lost its appeal and right enough times that despair lost its grip. They are not in the middle. They are in the tension, and they have stopped wishing it would resolve, and they have started building from inside it.

That is the posture. It is not natural. It is not comfortable. And it is the most honest place to stand when the technology moves faster than our capacity to understand it, and the only irresponsible position is the one that pretends to have it figured out.

^1^ Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
^2^ Tetlock, P.E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press.

ARTICLE 6

What AI Cannot Do (and Why That Matters More Than What It Can)

Artificial intelligence cannot sit with another person in distress and mean it.

That claim requires no hedging, no qualification, no "at least not yet" addendum that preserves the possibility of future disruption. The limitation is not technical. It is ontological. Co-regulation, the process by which one nervous system stabilizes another through proximity, breath, posture, and the quality of sustained attention, requires two embodied agents whose biological systems are literally entraining to each other in real time.^1^ A language model can generate the words "I'm here for you." It cannot be here for you, because "here" requires a body, and "for you" requires a self that has something at stake in the encounter. The sentence "AI cannot sit with another person in distress and mean it" is not a prediction about current limitations that future architectures will overcome. It is a statement about the kind of thing meaning is.

This distinction matters because the most consequential work in education, in organizational leadership, in any field that involves human development, is relational work that depends on embodied presence. And that work is precisely the work most organizations have buried under operational demands they never chose but have stopped questioning.

The Catalog of What Requires a Body

Consider what happens in a room when a facilitator holds silence during a difficult conversation. The silence is not empty. It is full of micro-signals: the shift of someone's weight in a chair, the quality of eye contact breaking and reforming, the barely perceptible change in a colleague's breathing that communicates "I am deciding whether to trust this process." Stephen Porges's polyvagal theory documents how the human autonomic nervous system reads these signals through a process he calls neuroception, a subconscious evaluation of safety that operates below conscious awareness and responds to cues no camera or microphone can fully capture.^2^ The facilitator who has learned, through years of practice, to trust a four-minute silence and let it do its work is not simply withholding speech. They are actively co-regulating with every person in the room through embodied signals that no natural language processing system can replicate, because the signals are not linguistic. They are biological.

Or consider the knowledge that accumulates over years of relationship with a specific person. A principal who has worked with a teacher for six years knows the difference between Tuesday's flatness and Thursday's flatness. She recognizes that the particular quality of "I'm fine" spoken in the hallway this morning carries a weight that the same words did not carry last week. This is not sentiment analysis. It is relational knowledge, knowledge that exists only inside the relationship that produced it, that cannot be extracted from that relationship and encoded in a training set, because the knowledge is not a pattern in data. It is a sensitivity built through accumulated mutual presence.^3^

AI cannot make the judgment call that contradicts the data because something about this particular moment requires a response the pattern would never predict. Gary Klein's research on naturalistic decision-making demonstrates that expert practitioners in high-stakes fields routinely make decisions that violate their own protocols, not from carelessness but from a form of recognition that integrates contextual information the protocol cannot hold.^4^ The experienced teacher who departs from the lesson plan because something in the room has shifted, the social worker who ignores the risk assessment score because she knows this family in ways the instrument does not, the principal who overrides the discipline matrix because the student in front of him is not the student the matrix was designed for: these are not errors. They are acts of contextual judgment that require a decision-maker who is present in the situation, not processing it from outside.

AI cannot look someone in the eye and say "this is causing harm" and carry the professional and personal risk of that honesty. Honesty that costs something is a fundamentally human act. It requires a body with something to lose, a career that could be damaged, a relationship that could fracture, a reputation that could be diminished. The weight of moral speech comes not from the content of the words but from the fact that a particular person, in a particular position, chose to say them knowing what they might cost. An AI system can generate the sentence "this practice is causing harm to students." It cannot risk anything by saying it, and

speech without risk is not courage. It is text.

The Discomfort of the Inventory

The catalog is extensive, and extending it further is not the point. The point is what the catalog reveals when you take it seriously. Each item on the list, co-regulation, relational knowledge, contextual judgment, moral speech, names a category of work that only humans can do. And in most organizations, that work receives the least structural protection. The people who should be doing the irreplaceable work are instead formatting documents for reports no one will read carefully, sitting in meetings that rehearse information everyone already has, processing administrative tasks that consume hours of attention that could have been directed toward the humans they serve.

This is where the argument becomes uncomfortable, and the discomfort is the point. If you accept that AI cannot do the relational, embodied, morally consequential work, then you also accept an obligation to audit how much of your organizational time is currently consumed by work AI could do. And that audit will reveal something most leaders do not want to see: that the organization has been spending its most expensive resource, human attention and presence, on its least consequential tasks, not because anyone decided this was a good idea but because operational defaults accumulated over years until the defaults became invisible.

The educational philosopher Nel Noddings argued that caring is not a feeling but a practice, one that requires what she called "engrossment," a full turning of attention toward the other person that cannot be faked or automated.^5^ If Noddings is right, then the current allocation of professional time in most schools and districts is a structural betrayal of the caring relationship that education is supposed to be built on. Teachers who spend two hours per evening on documentation are not failing to care. They are being structurally prevented from caring by systems that consume the attention caring requires.

The Genuine Contribution

This leads to an argument about AI's contribution that the marketing materials consistently miss. AI's most important function is not what it produces. It is what it displaces. Every minute of administrative processing that a well-designed system handles is a minute returned to the work that requires embodied human presence. The draft generated in twenty minutes instead of three hours, the communication template that eliminates forty-five minutes of formatting, the data summary that replaces two hours of spreadsheet manipulation: these are not impressive feats of technology. They are acts of structural liberation, small and unglamorous, that free human beings to do what only human beings can do.

But the liberation only works if you know what you are freeing people for. And this is where most organizations fail, not in adopting AI but in failing to do the prior work of identifying the irreplaceable. Michael Polanyi's concept of tacit knowledge, knowledge that cannot be fully articulated or transferred through explicit instruction, maps closely onto the catalog above.^6^ The relational sensitivity, the contextual judgment, the embodied co-regulation: these are forms of tacit knowledge that practitioners develop through years of practice and that no system, however sophisticated, can replicate. But Polanyi also observed that tacit knowledge is invisible to the systems that depend on it. Organizations benefit from it without recognizing it, protect it accidentally rather than deliberately, and discover its absence only after the people who carried it are gone.

The inventory, then, is not optional. It is the precondition for any responsible use of AI in human-centered work. The inventory asks: What do we currently ask humans to do that a system could handle? And what do we currently ask humans to do that only a human can do? And are we structurally protecting the second category, or are we allowing the first category to consume the time, attention, and presence that the second category requires?

Hubert Dreyfus spent decades arguing that artificial intelligence would never replicate human expertise because expertise is not rule-following but a form of embodied, situational responsiveness that develops through immersion in practice.^7^ His argument was widely dismissed during the early decades of AI research, and the recent success of large language models has been treated as a definitive refutation. It is not. What large language models have demonstrated is that pattern completion across vast corpora of text can produce outputs that are indistinguishable from human writing in many contexts. What they have not demonstrated, and what Dreyfus would have predicted they could not demonstrate, is embodied judgment: the capacity to be in a situation, to feel its weight, and to respond from within it rather than from above it.

The Center of the Argument

The most important use of artificial intelligence is protecting the space for natural intelligence to do what only it can do. That sentence carries the weight of everything this article has argued, and it only holds if the people responsible for implementation have done the honest, uncomfortable work of knowing what "only it can do" means in their specific context.^8^ Not in general. Not as an abstraction. In this building, with these people, on this particular afternoon when a student needs a human being who is free enough, present enough, and skilled enough to sit down across from them and mean it.

The farmer would recognize this argument. You cannot automate the part of the work that matters. You can only clear the ground so that the person doing the work has room to do it well. Whether they will do it well is a question no system can answer in advance.

^1^ Porges, S. W. (2011). *The Polyvagal Theory: Neurophysiological Foundations of Emotions, Attachment, Communication, and Self-Regulation*. W. W. Norton.

^2^ Porges, S. W. (2004). "Neuroception: A Subconscious System for Detecting Threats and Safety." *Zero to Three*, 24(5), 19-24.

^3^ Hollingsworth, S. (1989). "Prior Beliefs and Cognitive Change in Learning to Teach." *American Educational Research Journal*, 26(2), 160-189. See also Schon, D. (1983). *The Reflective Practitioner*. Basic Books. On knowledge embedded in relationship rather than extractable from it.

^4^ Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.

^5^ Noddings, N. (1984). *Caring: A Feminine Approach to Ethics and Moral Education*. University of California Press.

^6^ Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press.

^7^ Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press. Originally published 1972 as *What Computers Can't Do*.

^8^ This framing draws on Vallor, S. (2016). *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. Oxford University Press. Vallor argues that technology's ethical significance lies not in what it does but in the kind of human character it cultivates or degrades.

ARTICLE 7

The Question I Ask Before Every AI Decision

I have a question I ask before every AI decision, and I am not sure it is the right one. I want to say that up front because the original version of this article presented the question like a gift, something clean and portable, something you could unwrap on Monday morning and start using immediately. That version was easier to write. This one is harder because I have been living with the question for a while now, and living with a question is different from presenting one.

The question is: What is my evidence that this will serve the people I am responsible for better than what it replaces?

I still ask it. I ask it in meetings, in planning sessions, in the quiet moments when I am deciding whether to adopt a tool or keep doing something by hand. But I have started noticing things about the question that make me less certain it works the way I thought it did.

Let me start with what is genuinely useful about it. The question contains three requirements that most AI adoption decisions skip entirely, and walking through them still matters.

The first requirement is in "what is my evidence."

Not my intuition. Not my excitement. Not the testimonial from a vendor or the case study from an organization operating in a completely different context. My evidence. Data I have collected, observations I have made, outcomes I have measured. When you ask "what is my evidence," you force yourself to distinguish between information and interpretation, between what you know and what you have concluded, between data

and the story you are telling about the data. Most AI adoption decisions are driven by interpretation masquerading as evidence. "This will save time" is an interpretation. The evidence would sound more like this: "I tracked my time on this task for three weeks, the average was two hours, the AI alternative completed the task in fifteen minutes with comparable quality as measured by these specific criteria." The interpretation sounds like evidence, but it is not, and the gap between them is where bad decisions live.

I believe that. I have seen it play out. But here is where the uncertainty starts: I have also watched myself construct evidence that confirms what I already wanted to do. The question asks for evidence, and I provide it, and the evidence is real, but the selection of which evidence to foreground is shaped by a preference I formed before I started looking. The question does not protect against that. It asks what my evidence is. It does not ask whether I went looking for evidence that would support the conclusion I had already reached.

The second requirement is in "the people I am responsible for."

Not yourself. Not your convenience or your workload or your sense of professional accomplishment. The people the work is supposed to serve. This needs to be said plainly because most AI adoption decisions serve the person making them first and the people they are responsible for second. The honest version of the sentence is usually "this will make my professional life more manageable" rather than "this will genuinely improve outcomes for the people I serve." Personal ease is not wrong. It is legitimate. But it should be named for what it is, because when we dress up convenience as service, we lose the ability to ask whether they are actually aligned.

That part I still trust. But I have noticed something about how I apply it. When I want to adopt a tool, I define "serve" broadly: it serves them because it frees my time, which I then spend on higher-value work, which eventually reaches them. When I do not want to adopt a tool, I define "serve" narrowly: it does not directly improve their experience, so it does not count. The question lets me adjust the definition of service to match whatever I have already decided, and I am not always aware that I am doing it.

The third requirement is in "better than what it replaces."

This is the part I still find most genuinely useful, because it requires examining the current state rather than assuming it is broken. Not every manual process is wasted effort dressed up as tradition. Some of the things done slowly are done slowly for a reason, and the slowness is the value. A message drafted personally takes longer than an AI-generated version, and the time it takes is part of what makes it meaningful, because the recipient can feel the difference between a communication that cost someone time and a communication that was produced at scale. "Better than what it replaces" forces you to look honestly at what is currently happening before replacing it with something faster.

But here, too, I have caught myself in something uncomfortable. Last fall I was evaluating whether to automate a data tracking process. I asked the question. The evidence was clear: the manual process consumed

significant time, produced a measurable error rate, and created delays between gathering and use. The AI alternative eliminated entry errors, reduced processing time dramatically, and made information available same-day. The decision was yes, and I felt good about it.

What I did not examine, not until weeks later, was what happened to the time I reclaimed. I had told myself the reclaimed time would go to more important work. Some of it did. Some of it disappeared into things I cannot account for. The question helped me make the adoption decision well. It did not help me make the downstream decisions, the ones about what to do with the capacity that opened up.

There was another decision where the question worked more cleanly. A proposal came forward to use AI to generate personalized outreach. The templates were polished and the efficiency gain was real. I asked the question. I observed the current process carefully. What was happening was not just message generation. It was thinking about specific people, remembering specific context, recognizing specific needs. The AI could match the format of that work but not the relational intelligence embedded in the process. The time spent was not waste. It was relationship maintenance, and the relationship was the channel through which every subsequent communication would travel. The answer was no. AI handled formatting assistance, but the thinking about each person stayed human, because the thinking was the value.

I am more confident about that decision. But I notice that I am most confident about the decisions where the question gave me the answer I already preferred. The no felt right before I asked the question. The question gave me a structured way to arrive at a conclusion I had already reached intuitively. Maybe the question's job is to slow you down enough to check your intuition against something more rigorous. But maybe the question is sometimes a formalization of preferences rather than a genuine interrogation of them, and I cannot always tell the difference.

A third decision landed somewhere in between. An automated analysis protocol was in development. The data gathering and pattern recognition could be handled algorithmically, faster and more thoroughly than manual analysis. But this particular process was a community process. The findings needed to be examined by the people they described. The patterns needed to be interpreted by the people living inside them. The conclusions needed to be co-constructed, not delivered. The answer was not yet. The mechanical work of data preparation was automated. The human work of sense-making together was protected.

Three decisions. One yes, one no, one not yet. In the original version of this article, I presented these as evidence that the question works. Now I present them as evidence that the question helps, which is different from works, and the difference matters more than I used to think.

The question does not produce clean answers. It produces better questions disguised as answers, and sometimes I mistake the disguise for the real thing. It slows me down, and slowing down is valuable. It forces me to name who I am serving, and naming is valuable. It insists on evidence, and evidence is valuable. But it

does not protect me from my own capacity to construct a careful, well-reasoned justification for something I wanted to do anyway.

I still ask it. Every time. I do not have a better question, and the absence of a better one is not a reason to stop asking this one. But I hold it differently now. Less like a tool that produces reliable outputs and more like a practice that requires its own scrutiny, a question that needs to be questioned.

I keep asking. I keep watching what the asking does and does not do. The farmer would probably say that is enough.

Maybe.

ARTICLE 8

Three Practitioners, One Word

Nora

The literacy coach had been piloting an AI reading assessment tool for six weeks when the principal asked her to present results. She had the data. The tool had flagged seventeen students whose decoding patterns suggested compensatory gaps the classroom assessments missed, and fourteen of those seventeen turned out to need exactly the support the tool predicted. The accuracy rate was striking. The time savings were real.

But three students the tool got wrong kept pulling at her. One was processing a family separation that had reorganized his entire nervous system. One was an emergent bilingual whose Hmong literacy was sophisticated in ways the English-only tool could not see. One, she suspected, was performing struggle because the small-group pullout was the only part of his day where an adult sat next to him and paid full attention, and that was worth more to a nine-year-old than any assessment score.

The tool could not know any of this. It was not designed to know it, because this kind of knowing lives in the accumulated hours of a particular relationship between a particular adult and a particular child, and that is not a data problem. It is a presence problem.

She walked into the meeting and presented both results. The fourteen and the three. She recommended continuing the pilot with the relational caveat built into the protocol as the first line, not a footnote. The principal wanted a cleaner recommendation. Nora understood why. Clean recommendations move through

systems faster, and systems reward speed.

Maybe.

David

The assistant superintendent had spent two months building a policy analysis protocol with AI assistance. The district was reviewing its family engagement framework, and the volume of documents, three years of community input, survey data from nine schools, focus group transcripts, board minutes, was genuinely more than any team could synthesize by hand in the timeline the board had set. The AI tool surfaced themes across documents, identified contradictions between what the district said it valued and what the budget actually funded. That finding alone was worth the process, because the contradiction had been sitting in plain sight for two years and nobody had assembled the pieces, not from carelessness but because the volume made the pattern invisible to any single reader.

He did not present the synthesis directly. He took it to the community advisory council first, not as findings but as questions. He said: this is what the tool surfaced. You are the people described in these documents, and the tool does not know what it is like to live inside the patterns it identified. I need you to tell me what it got right and what it missed and what it found that is technically accurate but contextually wrong in ways that matter.

The council revised a third of the themes. Not because the AI was inaccurate, but because accuracy and meaning are not the same thing, and the distance between them is where community trust lives or dies. A finding that is statistically defensible and experientially wrong does more harm than a finding that is approximate and honestly held, because the defensibility becomes a weapon against the people who know from living inside it that something is off.

David presented the revised analysis to the board. It took longer. It was messier. A board member asked why he had not simply presented the AI analysis, which was more coherent.

Maybe.

Rosa

The high school history department chair had been asked to build a PD sequence on AI integration for her content team. The question of what AI means for history instruction is not simple, because history is not a content delivery problem. The discipline is built on holding multiple, contradictory accounts of the same event and constructing an interpretation that accounts for the evidence without collapsing the contradiction. If that

sounds familiar, it should.

Rosa built three sessions. The first gave everyone unstructured time with the tools. No framework, no rubric. She wanted her team to develop their own experience before anyone told them what to think about it. The second asked each teacher to bring one example where the tool had done something genuinely useful and one where it had produced something that looked competent but was wrong in a way that mattered for the discipline. Recognizing disciplinary wrongness requires disciplinary expertise, and the whole point was to demonstrate that the expertise is the thing the tool cannot replace: not knowing the content, which the tool can access, but knowing what matters about the content, which it cannot.

The third session was where it got interesting. Rosa asked the team to design one unit where AI played a meaningful role and one where it was deliberately excluded, and to articulate why. The conversation that followed was the best professional dialogue she had facilitated in four years, not because the topic was AI but because the structure required people to make specific claims about what they valued in their practice and defend those claims with evidence. That kind of specificity is rare in professional development because it is uncomfortable.

Her principal observed the third session and asked whether the model could scale to other departments.

Maybe.

The Word Returns

Three practitioners, three contexts, three different decisions. None arrived at a clean answer. The clean answer is the one that has set something important down, and they chose not to set it down.

The farmer's neighbors wanted him to sort the events into good luck and bad luck because sorting is a relief. Nora, David, and Rosa are doing what the farmer did: refusing the sort. Not because they do not care about the outcome, but because they have been in the work long enough to know that premature sorting is where the most consequential errors live, the ones you do not see until the consequences have already reached the people you were trying to serve.

William Perry, studying how college students develop intellectually, found that the most sophisticated thinkers are not the ones who hold the most knowledge.^1^ They are the ones who have learned to make commitments in the face of uncertainty, to act without requiring the comfort of certainty first. That is what these three are doing. They are not waiting for the answer. They are acting inside the question.

This is what "maybe" sounds like in practice. It sounds like a literacy coach presenting both the fourteen and the three. It sounds like an assistant superintendent taking an AI synthesis to the people it describes before

calling it findings. It sounds like a department chair trusting her colleagues to think before she tells them what to think.

Maybe.

That word, returning here for the first time since the opening article, is not softer than it was. It has been tested by everything between here and there: the territory distinction, the anxiety of honest use, the efficiency critique, the both/and discipline, the limits of computation, the question that precedes every decision. The word has traveled through all of it and arrived unchanged, which is the point. The farmer's answer does not change because it was never about the specific event. It was about the posture, and the posture holds whether the news is wonderful or terrible or, as it usually is, both at once.

The farmer prepares the field. These three are preparing theirs.

^1^ Perry, W.G. (1970). *Forms of Intellectual and Ethical Development in the College Years*. Holt, Rinehart and Winston.

◆

The farmer has been in these pages before. He appeared in the first article, standing at a fence while his neighbors cycled through their certainties, and he offered the same word each time they arrived with their verdicts. Maybe. The parable was doing introductory work then, establishing a posture, laying the ground. Nine articles later, the farmer returns, and the story is no longer an introduction. It is a reckoning with whether any of what followed has earned the posture the farmer was modeling before we had done anything to deserve it.

Because here is what has happened between then and now, across these articles and across the months they represent. We examined the tools and the territory they reshape. We sat with the efficiency lie long enough to feel it pulling at us from the inside, the seductive compression of complex human work into metrics that feel like progress. We tried to hold both/and and watched it buckle under the weight of real decisions where both sides cannot actually coexist without someone absorbing the cost. We asked what AI cannot do. We asked the question before every decision. We watched practitioners attempt the farmer's posture in their own settings, with their own pressures, with varying degrees of success and failure. All of that has been the work of the series. And now the farmer is back, not to validate the work but to complicate it one more time.

The original parable, traced to the *Huainanzi* text, is about a man whose responses never change regardless of the event. The horse leaves; maybe. The horse returns with companions; maybe. The son breaks his leg; maybe. The son is spared from conscription; maybe. The neighbors, who represent the rest of us, keep arriving with their interpretations, their categories of luck and misfortune, their confident readings of what each event means. The farmer keeps planting. The parable is usually told as a story about wisdom, about the farmer's superior perspective, his ability to see what the neighbors cannot. I have come to think that reading is too generous, and the too-generous reading is worth examining because it protects us from something the parable is

actually trying to surface.

The farmer is not wise. The farmer is practiced. That distinction matters more than it might first appear, because wisdom suggests arrival, suggests a vantage point from which events can be properly assessed, and the entire point of the farmer's posture is that no such vantage point exists. What the farmer has is not superior knowledge. What the farmer has is a discipline built through seasons of watching predictions fail, through harvests that contradicted every forecast, through years of learning that the relationship between an event and its consequences is longer and stranger than any single moment of interpretation can capture. The farmer's "maybe" is not the calm of someone who understands. It is the restraint of someone who has been wrong enough times to distrust the feeling of being right.

That is a less comfortable version of the story than the one usually told, and the discomfort is precisely where its usefulness lives for the question of how organizations prepare for what is coming.

Because the preparation question is not a prediction question, and we keep confusing the two. The prediction question sounds like this: What will AI do to our field? Which roles will it eliminate? How fast should we move? When is the right time to adopt? These are the questions filling keynotes and retreats and board study sessions, and they are, with genuine respect for the people asking them, the wrong questions. Not because they are unimportant, but because they are unanswerable in the way they are being asked. Nobody knows what these systems will do to any field over the next decade. The people who claim to know are selling either their enthusiasm or their fear, and both are products, not predictions.

The preparation question is different. It does not ask what will happen. It asks what conditions need to exist so that whatever happens, the organization can respond with something other than panic or performance. What relational trust needs to be in place before the next capability shift so that when it arrives, people are making choices rather than reacting to pressures they do not understand? What evaluative disciplines need to exist so that when someone proposes a new tool, there is a shared, operational definition of "better" that everyone can evaluate against, rather than a mission statement nobody has read since the strategic plan was adopted?

Wendell Berry, writing about agricultural communities decades before anyone was thinking about artificial intelligence, observed that the health of a farm cannot be separated from the health of the community that tends it.^1^ The farmer who prepares the field is not performing a technical task. The farmer is sustaining a relationship between soil and season and labor and community that has never been reducible to any single input. Berry's point translates with uncomfortable precision: the organizations preparing well for AI are not the ones purchasing the best tools. They are the ones tending the conditions that determine whether any tool lands as support or surveillance, as resource or weapon.

Four elements of that tending have become visible over two years of watching organizations do this well and poorly, and I want to name them without pretending they are a framework. They are observations, not a program.

The first is building the human infrastructure before the technical infrastructure. If people in an organization do not trust the leadership making adoption decisions, adding AI tools will amplify the distrust. The tools will feel like monitoring, because in a low-trust environment every new system feels like monitoring. People have learned from experience that new systems are often deployed to surveil rather than support, and that learning is not irrational. It is pattern recognition earned through exposure. The relational foundation that transforms information from judgment into insight does not exist in these settings, and no tool can create it. Technology adoption built on a relational deficit reproduces the deficit at higher resolution. The sequence matters.

The second is establishing what "better" means before adopting anything. This sounds straightforward until you try it. Most organizations do not have a shared, operational definition of what improvement looks like for the people they serve. They have mission statements and strategic plans and value propositions, but they do not have the specific, agreed-upon criteria that would allow anyone to look at an adoption decision six months later and say, with evidence, this served people better than the alternative. Building that shared definition takes longer than the adoption itself. It requires argument and revision and the discomfort of naming specifically what you are trying to produce, which means naming what has been failing to be produced. That naming is where most organizations stop, because the conversation becomes uncomfortable in ways that purchasing software never does.

The third is protecting the irreplaceable. Identify the human work that must not be automated and build structural protections around it so that efficiency gains flow toward that work rather than away from it. Without those structural protections, freed capacity does not reach the irreplaceable work. It gets absorbed by operational demands that are always hungry for more, and the work that only a human being paying full attention can do continues to starve. The inventory is not complicated: What work in this context can only be done by a person who is fully present? What work is currently done by people that a system could handle? Once those two lists exist, the structural question is how to ensure that time recovered from the second list actually reaches the first. Without structural answers, default behavior wins, and default behavior in any organization is to absorb freed capacity into more operational tasks. You have to design against the default, which means you have to name the default, which means you have to admit that the default is already winning.

The fourth is holding space for the both/and at the organizational level. The person who built an AI-powered workflow and the person who refuses to adopt anything new are both essential voices, because the builder sees possibilities the resister misses and the resister sees risks the builder ignores, and the organization needs both perspectives to do this well. This requires a culture that can hold contradiction without resolving it

by picking a direction and marginalizing the perspective that lost. Most organizational cultures cannot do this. They resolve tension through decision, which sounds healthy until you notice that "decision" in these contexts usually means one group's evidence gets elevated and another group's evidence gets dismissed, and the dismissed group stops contributing honestly, and the organization loses the very perspective it most needed to hear.

These four elements are not a checklist. They are not something you complete before moving on. They are ongoing, and they are harder than the technology itself, and that is exactly why most organizations skip them. The technical work is visible, purchasable, demonstrable. The human work is slow, relational, and difficult to show in a slide deck. The gap between what organizations invest in and what actually determines whether the investment lands well is the gap between the farmer and the neighbors. The neighbors are buying weather forecasts. The farmer is tending soil.

I want to be honest about my own uncertainty here. I do not know if these four elements are sufficient. I do not know if I have named them correctly, or if I have missed the fifth element that turns out to matter more than the other four combined. I have watched enough organizations attempt this work to believe these observations are grounded in something real, but I have also watched enough of my own confident observations collapse under the weight of contexts I did not anticipate. The farmer's posture applies to the person describing the farmer's posture. Maybe these elements are right. Maybe they are partial. Maybe the next two years will surface something none of us can see from here.

The farmer prepares the field. He does not know what the season will bring. He does not know if the rain will come or if the drought will hold or if the horses will stay or leave or return with company. He tends the soil because the soil is what he can tend, and because he has learned, through more seasons than his neighbors have watched, that preparation is not the same as prediction, and that the people who confuse the two are usually the ones standing at the fence.

The field still needs planting. That much has not changed.

^1^ Berry, W. (1977). *The Unsettling of America: Culture and Agriculture*. Sierra Club Books.

ARTICLE 10

Maybe (Reprise)

This series began with a parable about suspended judgment. It is ending with an argument about the same thing, because the parable has done its work and repeating it would be a way of avoiding what comes next.

What comes next is an accounting.

The "maybe" posture was held for nine articles. It was applied to AI adoption, to institutional decision-making, to the interior lives of people who build systems that affect other people. It was tested against enthusiasm and against criticism and against the particular pressure that comes from audiences who want you to pick a side so they can decide whether to listen. The posture held. Whether it should have is a different question, and it has not been answered yet.

Here is what the posture protected against.

It protected against premature certainty. There were stretches during this writing when the trajectory of AI in human work felt settled, fully legible, pointing somewhere specific. Those stretches followed mornings when something elegant got built, when time collapsed in a way that felt like progress. The posture did not allow staying in those mornings. It sent the thinking back to the afternoon, where the system had failed in a way nobody anticipated, where people with legitimate objections had not been consulted because the morning had felt so clear that consulting seemed unnecessary. The posture kept things located in practice rather than floating above it.

It protected against tribal alignment. Both camps made offers. The AI enthusiasts wanted this work as evidence that the tools deliver. The AI critics wanted it as a cautionary account of what gets lost. Both offers came with something genuine: the comfort of shared conviction, the relief of a settled identity, the warmth of a group that agrees with your conclusions and does not require you to hold contradictions. The posture declined both. The loneliness of that was real. The accuracy it produced was worth the loneliness, though saying so does not make the loneliness retroactively comfortable.

It protected against narrative closure. Writers want stories that resolve. Here is the problem, here is what it means, here is the recommendation. The posture does not allow that sequence to complete. It insists the problem is still shifting, the meaning still contingent, the recommendation premature. This is uncomfortable for anyone who produces written work, because writing pulls toward conclusion the way water pulls toward level. Refusing the conclusion requires sustained effort, and the effort is invisible to the reader, which means it looks like indecision when it is actually the most expensive form of intellectual honesty available.

Here is what the posture cost.

It cost speed. People who have chosen a side move faster because their frame tells them which information to process and which to discard. They act with confidence that was, at times, enviable. The posture slows everything down. It requires looking at the information your frame would have filtered out, and looking takes

time, and sometimes the slow response arrives after the decision has already been made by someone faster.

It cost surface authority. Not the deeper confidence that comes from doing honest work over time, which remained, but the surface kind, where you walk into a room and say what should happen without qualification. Everything got qualified. "And also" appeared more than any audience wanted, and the room sometimes responded with a mixture of respect and confusion that does not convert easily into followership.

It cost belonging. In a polarized discourse, platforms accumulate around the clearest positions. The strongest advocates and the strongest critics both gather audiences, because clarity is compelling even when it is partial. The posture is useful to a smaller group, the group that has tried certainty and found it insufficient, and that group exists but it is not large.

Here is what the posture revealed.

It revealed that the most consequential transformation is interior. The technology shifted rapidly. New capabilities arrived at intervals close enough to prevent adjustment to the previous ones. But the significant changes happened inside the practice, in the way it became possible to think about tools, about power, about whose judgment gets centered and whose gets automated. The posture forced regular examination of motive. Is this being built because it serves people, or because building produces a feeling of competence that is easier to access than the relational work it replaces? Is this being written because something needs saying, or because the topic draws attention and attention is its own reward? The questions are not comfortable. The answers are not always flattering. Asking them consistently produced better work than any capability adopted this year, and that sentence is true and it is also the kind of sentence a person writes when they want credit for their own discomfort, which is itself worth examining.

It revealed that technology is a mirror before it is a tool. Every system built showed something about the builder before it showed anything about the capability. The automation revealed what had been avoided. The first thing you do with a powerful new instrument tells you more about your priorities than any mission statement, and what got built first, across the field, was not always what was most needed.

It revealed that suspended judgment is not a strategy for one domain. It functions for technology decisions, for policy decisions, for organizational change, for the way a person relates to their own certainty in any context where the world is more complex than the narrative constructed to contain it.

The series that grew from this practice is not about AI. It was never about AI. It is about the interior architecture that determines whether any powerful instrument serves liberation or reproduces the patterns it was supposed to interrupt. The question was never whether the technology is good or bad. The question was always about the person holding it.

Nine articles traced that question through different angles. None of them resolved it. The parable at the center of the series is about a person who watches events arrive and declines to sort them, and the power of that story is not in its wisdom but in its cost: you never get to stop. There is no point at which enough evidence has accumulated to permit a conclusion. The posture is permanent or it is nothing, and permanence is not a thing most people want from their intellectual commitments.

That is the discomfort the series has been building toward, and it lands here without resolution because resolution would betray everything the preceding nine articles argued. If the posture is honest, it applies to itself. If suspended judgment is the discipline, then the judgment of the discipline is also suspended. The series cannot conclude that "maybe" is the right word, because concluding that would be a form of certainty, and certainty is what the posture exists to resist.

So the series ends where it began, with the same word, carrying the same weight it carried at the start, except that now the weight includes everything examined across nine articles and the cost of carrying it and the fact that carrying it did not produce arrival and was never going to.

Maybe.

What do you do with a word that refuses to become an answer?

The Maybe Series

The Interior Architecture of Transformation

The Interior Architecture of Transformation

Joshua T. Fraser, Ed.D.