THE HUMAN QUESTION
with Sharon Higgins
EDITION NO. 4 · SEPTEMBER 27, 2026

When AI Begins to Shape Culture

Who decides the reality we inherit?

We keep arguing about AI one incident at a time: Did a student use it to cheat? Did a hiring algorithm show bias? Did an artist lose control of their work? Did a factory floor lose its second shift?

These arguments aren't wrong. They're just not the argument.

I've spent enough years running workshops in economic and business development, sitting with small-business owners trying to get a loan approved, and placing interns with clients to notice a pattern: while we debate the branches, something is growing at the root. AI is not simply entering our lives as a tool. It is entering the institutions that decide, on society's behalf, what counts as knowledge, who counts as qualified, what counts as valuable work, and which version of events counts as true. That is not only a technology question; it is cultural, and it is the one we keep skating past.

What Our Institutions Permit

We tend to think of culture as the soft stuff: language, tradition, or shared belief. But culture is also manufactured, quietly and continuously, by institutions: schools decide what's worth knowing; universities decide who's credentialed; employers decide which skills have market value; banks decide who gets the capital to build something; governments decide which decisions carry public force.

Culture is not just what we celebrate; it's what our institutions permit.

Every one of those institutions is now being mediated by AI, which means AI isn't just entering a market. It is stepping into the machinery through which a culture reproduces itself, generation after generation. I've watched a few "transformative" technologies move through institutions this way before. Each one arrived promising efficiency, and each one left having steadily redrawn who gets to decide what matters. This one appears to be moving faster than any of the others did.

That changes the question. It is less "Will AI take my job?" and more this: once AI starts filtering who gets hired, who gets funded, who gets trained, and whose version of history surfaces first, who authorized that, and on what grounds?

Some institutions are beginning to answer that question out loud. New York City's public school system, the largest in the country, has barred AI from grading and from other consequential decisions about students. Teachers may still use it for planning and routine tasks, but professional judgment stays with the educator.

The questions beneath that policy sound like implementation details: Who retains judgment? What counts as effective instruction? How should bias and equity be assessed? Which decisions must remain human? They are not implementation details. They are statements of institutional value. A school system that decides an algorithm may help a teacher but may not decide a child's grade, placement, promotion, or graduation is doing more than writing an AI policy.

It is drawing a cultural boundary around human judgment.

What makes this moment unusual is its timing. The city's Technology in Schools Coalition, which brings together students, educators, parent leaders, union partners, and outside experts, is not due to deliver its policy recommendations until April 2027. The rules are being written while the transformation is already underway. We are not observing this cultural shift after the fact. We are participating in it.

In my last edition, I argued that capability is not authority. The question that follows is harder: Who decides where capability is granted authority? The answer does not depend on whether AI exists, or even on whether people use it. It depends on where our institutions permit AI-generated capability to become institutional authority.

It is worth being precise about what "AI" means here. Generative AI, the chatbots and drafting tools most people mean when they say the word, is the newest and most visible part of a longer movement toward algorithmic decision-making. The cultural argument isn't limited to systems that generate language. It includes the older systems that rank, score, and predict, the ones already embedded in lending, hiring, and admissions long before anyone had heard the term "large language model." Some of the clearest evidence for this argument comes from those older systems.

Education and a Shrinking Pipeline

Education is the clearest early case for generative AI specifically, and we've already settled into a comfortable phrase for it: "AI will personalize education." That phrase is usually where hard questions go to die. The real questions are less marketable. If a model can explain almost anything instantly, what's left for a curriculum to insist a student still remember? The risk was never mainly that students would cheat. It is that institutions, chasing efficiency, will redesign learning around whatever AI happens to measure well

and discard what it can't, without ever voting on it: curiosity, dissent, and the slow parts of growing up.

But the version of this I've watched most closely isn't in a school system. It's in my own business development practice, where for years I ran workshops and placed interns with small businesses. The interns worked in marketing, client communication, and the clerical backbone of running a business while they earned their degrees. The task itself (draft the social post, build the flyer, organize the client files) was never really the point. The point was a 20-year-old learning what it feels like to sit with a business owner, take a vague ask, and turn it into something usable, with someone correcting their mistakes along the way.

In the small-business environments I know, those opportunities have begun to shrink, and AI is accelerating the change. It now drafts the marketing copy, the flyer, and the first-pass content calendar faster, and often better, than a student still learning the craft. The business no longer needs the intern to produce the deliverable. That means an entire generation's first real exposure to how a business operates, not the theory but the lived experience of it, may be disappearing as quietly as the deliverables it used to require.

The Bottom Rung

The same pattern shows up one level higher, in professional work, where I've seen employers scale back junior roles. The reasoning has usually been explicit: AI now produces the research memo, the first-draft analysis, the entry-level output those roles existed to do. In retrospect, the decision makes sense, but entry-level work was never only about its output. It is where professional judgment gets built: years of routine tasks that teach a person to recognize patterns, to notice when something looks wrong, and to earn trust before anyone hands them the real decisions.

The outcome was easy to notice; the development happening underneath it was not. Automate that layer, and an organization looks leaner for a few years, then discovers it has no one ready to run the place. The experts of twenty years from now are supposed to be learning today. If we remove the environments in which novices become experts, where, exactly, will they learn?

We may be automating the bottom rung and only later notice we've also removed the way anyone reaches the top.

For the students in my internship programs, the loss comes even earlier. The rung is disappearing before some of them ever get a foot on the ladder at all.

The Deeper Fracture

Beneath the questions of work and opportunity sits something harder to name: the question of what we still hold in common. When schools, search systems, and public institutions each deliver an individualized version of knowledge, personalization stops being merely a service feature. It begins to affect the common ground a culture depends on.

A shared culture requires at least some shared reality: a common set of facts we can disagree about, a common history we can argue over, a common vocabulary for naming what matters. Personalization is useful precisely because it narrows what each of us sees. Across an entire society, that same narrowing can leave us with fewer and fewer things we have all encountered together.

AI does not need to lie to erode that ground.

It only needs to consistently choose what to include, what to rank first, and whose language sounds authoritative. That is curation functioning as culture inside the very institutions this edition is about, and it happens without anyone formally deciding it should. If each of us is shown a slightly different world, what will we still be able to talk about together?

Consequence, Not Sophistication

So what, exactly, should we regulate? Not the technology itself. That is too blunt an instrument, as likely to strangle a useful tool as a harmful one. The sharper principle, and the central argument of this edition, is this:

The greater the consequence of an AI-assisted decision, the greater the obligation surrounding how it gets deployed.

An AI drafting a meeting agenda is not in the same category as one recommending who gets hired, who gets a loan, or who gets denied a benefit. Consequence, not sophistication, is the right axis.

What would that look like in practice? When an institution uses AI in a decision that shapes someone's life, the people affected should know it is being used. Its likely effects should be weighed before it is deployed, not discovered afterward. A human being should remain answerable for the outcome, because "the algorithm decided" cannot become an institutional alibi. There should be a real way to appeal. And permission to use it should be revisited over time, not granted forever.

One nuance is worth insisting on: deployment is consequential even when a human technically makes the final call, if the AI's recommendation is what did the deciding.

Otherwise, "a loan officer signed off" becomes a loophole wide enough to drive every hard case through.

There is also a prior layer this framework doesn't cover on its own. Some risks arise in how a system was built, not how it's used: unlawful training data, concentrated ownership of the infrastructure, bias built into the model before any institution ever deploys it. The fuller version, then, is that we should regulate *development* where the system itself creates systemic risk, and regulate *deployment* where its use creates consequential human outcomes. Different problems call for different points of intervention.

When Prediction Becomes Prescription

Here is the part that should unsettle us more than it currently does: a system built to *predict* who's likely to repay doesn't stay neutral for long once an institution starts acting on it.

I've watched this from the other side of the table, working with small-business owners trying to access capital. Modern underwriting is increasingly data-driven and, in many settings, algorithmically assisted: models may weigh variables such as credit history, time in business, cash flow, collateral, and other indicators of risk. What those indicators can fail to capture is the full picture: the owner who took a personal credit hit covering payroll rather than laying anyone off, the decade-long relationship with the same suppliers, the business that survived something that closed three competitors. What registers instead is a file that resembles businesses that defaulted in the past, and the loan is declined on that basis.

The trouble is that "who defaulted historically" is not a neutral category. It reflects who got fair terms in the first place, who had a banker willing to look past a thin file, and who had the capital to weather a bad quarter without ever showing up as risk. A model trained on that history doesn't correct the pattern. It automates it and hands the lender's decision back with the authority of a number instead of a person's judgment.

The system didn't just predict who would struggle to access capital. Left alone, it helps decide who does.

That is a harder problem than "Is the model accurate?" An accurate model built on an unjust pattern doesn't correct the injustice. It can formalize it and give it the appearance of mathematical inevitability. Consequential deployment needs scrutiny before scale, not a bias audit after a decade of declined applications.

Accuracy Is Not Legitimacy

A school's grade, a bank's underwriting decision, an employer's evaluation, a government's ruling: each carries weight not purely because it is correct, but because people accept the institution's authority to make that call. That authority is trust, built over time, and it is not the same thing as statistical performance.

Ask any small-business owner who has been declined by an algorithm what the experience feels like. There is no loan officer to make the case to, no chance to explain the missing context, often not even a clear reason beyond a score. If those judgments increasingly come from systems the affected person cannot examine, challenge, or understand, accuracy is no longer the only relevant question. A decision can be perfectly correct on its own terms and still be illegitimate because the person subject to it never had a meaningful way to contest it.

Perhaps the quieter danger isn't that AI will get things wrong. It is that it will get enough things right, often enough, that institutions stop feeling obligated to explain themselves at all.

Legitimacy erodes long before anyone can point to an error.

The New Divide

The divide that matters most may not be who has access to AI. It may be between people who direct it and people who are directed by it. In the intern programs I've run, it shows up as the difference between organizations that redesign around the technology on purpose and those that let a program quietly disappear.

Some businesses restructure the internship instead of eliminating it. The student directs the AI, checks its output, catches what it gets wrong, and becomes the judgment layer instead of the labor layer. That is arguably a better apprenticeship than the one it replaced. But it requires deliberately choosing to preserve the pipeline, not just noticing that the labor cost dropped. That kind of redesign also takes resources and mentorship time, which a well-funded firm can absorb far more easily than a smaller practice can. And the small businesses flattened into a bad credit score by a lending algorithm are, more often than not, the same ones that can't spare the time to reinvent an internship program. It is easier to simply not offer one this year.

That is a harder inequality to see than the old digital divide, because it doesn't show up as a lack of access. The gap is in who has the standing, and the slack, to redesign the system on purpose, and who just quietly loses the program, the loan, or both.

The Uncomfortable Middle

None of this is an argument against AI. It is an argument for governed experimentation: evidence before scale, real accountability, and a willingness to shut a system down when it fails the people living under it. AI has also expanded access and opened doors once guarded by cost or geography. The goal was never to protect the existing culture from AI; our culture has plenty worth correcting on its own. The goal is to use the tool to fix what's unjust without letting it automate those same injustices under a new, harder-to-question authority.

The Human Question

AI is not going to independently decide who we become as a culture. It doesn't have that kind of agency, and it is a mistake to talk about it as though it does. What will decide it is thousands of smaller institutional choices: what a school board approves, what a bank's underwriting model is trained to weigh, what a business development practice decides an internship is still for, what a company stops paying humans to do.

That is the more hopeful way to see it, because it means the outcome isn't fixed by the technology. It is still being decided, in ordinary meetings, by people who could ask better questions than most of us currently do.

If AI is going to help determine who gets educated, who gets trained, who gets hired, who gets capital, and what we accept as real, shouldn't we decide deliberately how much authority we're willing to hand it before we've already handed it?

The future we build with AI will not arrive as a single event we get to vote on. It will be authorized, quietly, one declined loan and one lost internship at a time, and we will inherit whatever we failed to examine.

A Constructive Answer

Can we let AI into our institutions without surrendering the judgment that makes them worth trusting? I believe we can. Doing so requires us to decide deliberately and in the open, where AI may assist and where a human must remain answerable. AI can inform the decisions that shape our lives. It should not quietly become the one making them. Where that line is drawn must remain ours to decide.

Sharon Higgins, MPA, MSEd

SOURCES & FURTHER READING

New York City Public Schools, "Guidance on Artificial Intelligence and Screen Time," 2026–27 school year. Schools.nyc.gov

New York City Public Schools, "New York City Public Schools Announces Release of AI Guidance for Educators and School Leaders," March 24, 2026. Schools.nyc.gov

Office of the Mayor, City of New York, "Mayor Mamdani and Chancellor Samuels Put Students First with Nation's Broadest Generative AI Moratorium in Schools," September 2026. nyc.gov

Government Technology, "NYC Schools Prohibit AI for Grading, Discipline, IEPs," March 25, 2026. govtech.com

Andreas Fuster, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, "Predictably Unequal? The Effects of Machine Learning on Credit Markets," Journal of Finance 77, no. 1, 2022.

Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace, "Consumer-Lending Discrimination in the FinTech Era," Journal of Financial Economics 143, no. 1, 2022.

Solon Barocas and Andrew D. Selbst, "Big Data's Disparate Impact," California Law Review 104, 2016.

Cathy O'Neil, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy, Crown, 2016.

Virginia Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor, St. Martin's Press, 2018.

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023. nist.gov

European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III (high-risk uses in education, employment, and creditworthiness), 2024. eur-lex.europa.eu