

FIELDWORK NO. 033 · AUTOMATION & LABOUR FORMATION

Can We Predict What New Work Automation Creates?

History suggests that some of the work created by automation can be seen before it exists. The strongest clue is not what the machine replaces. It is what the machine suddenly needs.

AUGUST 2026 · THIRTEEN HISTORICAL TRANSITIONS, 1950-2011

01

The Answer

Yes, partially.

We took thirteen automation transitions old enough to have settled labour markets, set a cutoff date for each one shortly after the technology became commercially real, wrote down what should have been visible at that moment, and only then looked at what happened. The framework worked in one category and performed badly in three others. The failures were consistent enough to be useful.

The most forecastable new work appears where automation requires an input that nobody currently produces in the form the machine needs. Everything else is much harder to see coming, and almost impossible to size or time.

That single mechanism accounted for the strongest results across the set. Address hygiene for optical mail sorting. Continuously maintained navigation databases for flight management systems. Specimen collection around automated clinical testing. Structured documentation support around electronic health records. In each case the work was legible from the machine's requirements before the labour market reacted to it, and in several cases the industry acted on it years before anyone had a name for the job.

8 of 13

CASES WHERE NEW WORK FORMED AROUND AN INPUT THE MACHINE REQUIRED AND NOBODY SUPPLIED

55 → 1

USPS REMOTE ENCODING CENTRES AT PEAK, AND TODAY

-97.5%

US ONE-HOUR PHOTOFINISHING EMPLOYMENT, 2000 TO 2024, WITH NO REPLACEMENT LAYER

The Most Predictable New Work

In 1983 the US Postal Service was deploying optical character readers that could sort machine-readable mail cheaply and could not read much else. The bottleneck moved instantly from sorting to input quality, and the input was written by the sender, who had no reason to care.

This was not a subtle inference available only in retrospect. A government audit dated 6 January 1983 was already recommending that the Postal Service improve its readability guidelines, gather data on how mailers responded to requests that they clean up their mail, broaden the service that converted existing customer lists, and keep the national address directory updated. Read that as a contemporaneous document and it is the whole forecast. The machine consumes an input produced badly by an outside party, somebody has to convert the existing stock, and somebody has to maintain the authority against which addresses are checked.

What followed was a presort and address-hygiene industry, paid through postage discounts rather than wages, and a certification regime for the software that does the cleaning. The work did not appear on the automating organisation's payroll. It appeared in its customers' supply chain.

Read the machine's input requirements. They often tell you where the next labour market will form.

The aviation case is sharper still, because the industry moved first. Flight management systems needed a current, standardised, worldwide navigation database that could not be allowed to drift. In 1973 the airline electronics committee convened a standards effort. In 1975 it published the specification that still governs how navigation data is prepared and transmitted. That is years before the two-pilot glass cockpit made the whole thing indispensable, and roughly a decade before the certified data-production industry matured.

An industry does not convene a standards committee for an input it expects to take care of itself. The constraint was seen, acted on, and institutionalised before the occupation existed.

Automated clinical chemistry followed the same shape from the other direction. Once a single analyser could run twelve tests on one sample, the scarce thing became the sample. Getting blood out of a patient is manual, and it separated out into its own certified occupation. Automation created a hands-on job upstream of itself.

Seeing a Problem Is Not Enough

This is the part most people get wrong, and the historical record is unusually clean about it.

Lotus 1-2-3 shipped in January 1983. Within nineteen months *Business Week* was running a feature on how personal computers trip up executives, cataloguing eight-figure forecasting errors caused by a single missed formula. The *Wall Street Journal* covered a Dallas energy company's spreadsheet loss in August 1984 and asked in print by 1986 whether software firms could be held responsible when a program makes a costly error. Commercial spreadsheet auditing tools were selling in the City of London by the mid-eighties.

The problem was real, recognised immediately, and tooled. Forty years later, spreadsheet auditing is a niche service and not an occupation. There is one environment where it became routine paid work: UK tax auditors sampling the spreadsheets businesses use to calculate VAT. That is the one environment where somebody outside the firm was entitled to demand proof.

Digital photography makes the point even more directly. By 2000 the organisation and retrieval problem was obvious to everyone looking at it. Kodak was presenting an electronic shoebox browser at an academic workshop, working from the assumption that a family collection would run to tens of thousands of images. A contemporaneous MIT paper named the problem and stated, in the same breath, that the aim was to solve it "using software agents rather than human labor, as much as possible."

They were right. One-hour photofinishing employment fell 97.5 percent between 2000 and 2024 and essentially nothing formed to replace it. The scarcity was correctly identified, and the party feeling it was a household with no budget line for photo librarians.

A new scarcity creates demand for a solution. It creates human work only when the party feeling the scarcity has a reason and a budget to buy labour.

The accountability regularity

Some new work does form around externalities and legal boundaries. Scanner mispricing produced a permanent state inspection regime that still runs. Self-checkout moved scanning to an unsupervised customer and generated a supervision and loss-analytics layer. Technology-assisted review in litigation produced defensibility work the moment a court took a view on it. Regulated laboratory quality control attaches to every run, forever.

The pattern underneath is not that hard-to-verify systems create reviewers. It is narrower.

Verification work scales when someone outside the automating organisation is entitled to demand proof. A regulator, a court, a payer, a platform, an insurer, or a customer with contractual recourse. Where nobody has that standing, the work stays small no matter how consequential the errors are.

Treat this as a structural regularity rather than a forecasting rule. When we went back through contemporaneous sources looking for evidence that anyone saw the accountability layer coming, we found none in any case examined. Aviation recognised the missing data input in 1973 and did not build a formal acceptance regime for data suppliers until 2005. The regularity is real. It is not something people reliably anticipate.

0 4

What History Does Not Let You Predict

Four things, and the fourth is the one that costs money.

Scale. The framework got magnitude wrong in both directions, repeatedly. Nothing in the 1976 diagnosis of supermarket scanning suggested that third-party analysis of the resulting data would become an industry. Contemporaries did slightly better and still badly underrated it: the 1970 business case that persuaded the grocery industry to adopt the barcode was built on checkout labour and repricing savings, with the information benefits deliberately downplayed as soft and hard to quantify. That soft benefit became scan-data analysis and category management. You can identify the right new scarcity and still misjudge its value by an order of magnitude.

Timing. Input work arrives with adoption. Everything else waits for a trigger, and the trigger is often unrelated to the technology. Sealed containers made cargo unobservable in 1956. The large security-screening apparatus arrived in 2002, for reasons that had nothing to do with shipping.

Form. Automated clinical panels produced a substantial amount of new human work in following up abnormal results in people who were never previously tested. That work landed on physicians. It never got a title and never appeared in any occupational classification. A function can matter economically and remain invisible to labour statistics permanently.

Durability. Automation reliably creates a human layer for the cases it cannot handle, and that layer is reliably temporary. The Postal Service built 55 remote encoding centres, staffed by people reading handwriting the machines could not. When they opened, the equipment could read about two percent of handwritten addresses. Within fifteen years it read about ninety-five percent. One centre remains.

Exception work is easy to predict and dangerous to treat as durable. If the job exists because the machine has a residual error rate, the machine is coming for it next.

There is also a null result worth holding onto. Run the diagnostic on the automatic elevator in 1950 and every question comes back empty. No input the passenger doesn't already supply free, no cost landing on an outside party, no boundary crossed, nothing scarce that anyone could buy. A maintenance and inspection layer formed and nothing else did, and elevator operator remains the only occupation in the 1950 US Census that automation is generally agreed to have eliminated outright. A diagnostic that cannot return "nothing forms here" is not measuring anything.

History is much better at telling you where new work is likely to form than how many people it will employ.

0 5

Five Questions to Ask Before You Automate

0 1

What input does this system require that nobody currently produces in the form it needs?

Who owns that input today. Who will have to create it, clean it, structure it, version it, or vouch for it. If the answer names a party outside your organisation, that is where the work forms, and you will probably pay for it through a discount or a vendor rather than a wage.

HIGHEST CONFIDENCE QUESTION IN THE SET

0 2

What happens when this action occurs a hundred times more often?

Which failure that was previously too small to matter becomes economically important. Then ask who eats it. If the answer is your own organisation, expect much less new work than intuition suggests. If it is an outside party with leverage, expect a durable function and expect it late.

0 3

Who is entitled to demand proof?

Does the output now enter the jurisdiction of a regulator, a court, a payer, an insurer, a platform, another company, or a customer with recourse. If nobody outside has standing to ask, do not budget for a verification workforce. The spreadsheet is forty years of evidence on that point.

0 4

What becomes scarce next, and who feels it?

Then the second half, which is the half people skip. Does that party have a budget for human labour. Scarcity felt by consumers gets absorbed by software. Scarcity felt by organisations turns into roles.

0 5

Does this new work exist only because the machine is currently weak?

If yes, plan it as a project rather than a department. The technology improving is the same event as that work disappearing.

0 6

What This Means for AI

The following are forecasts produced by applying the historical framework, not findings. They should be read at that weight.

Machine-readable organisational context

This is the strongest forecast available, because it rests on the strongest historical mechanism. Agents need authoritative, current, structured knowledge of internal rules, permissions, exceptions, contract terms, customer states, workflows, entitlements, and approved actions. Almost no company maintains any of this in the form an agent can consume.

Expect work to form around producing that context, keeping it current, versioning it, defining who is allowed to change it, and vouching for it when agents act across organisational lines. That is the address-hygiene pattern and the navigation-database pattern, one abstraction level up. Forecast the function and resist naming the job.

Accountability and authorisation

Where agents act on an organisation's behalf against outside parties, those parties will want to know what the agent was permitted to do and who is answerable when it goes wrong. Expect work around authorisation, attestation, defensibility, audit trails, and liability allocation. Timing is not forecastable, and history says it will arrive suddenly, after a court decision or a large loss, rather than gradually.

Exception handling

Expect it early and expect it to be temporary. Treat headcount here as a bridge.

Verification

Do not forecast a general army of AI output reviewers. The historical evidence argues against it, and the spreadsheet case argues against it in the most direct way available: the problem was recognised immediately, tools were built, and the occupation never came. Expect durable verification work concentrated in the places that already have external claimants, which is regulated finance, healthcare, legal, and safety-critical systems.

07

Where the Forecast Can Break

The framework rests on an assumption that held in all thirteen historical cases and may not hold now.

Every automation studied here was narrow. It took one step of a workflow, and the adjacent steps stayed beyond it, so preparing the input, checking the output, handling the exception, and connecting the system to everything else had to be done by something other than the machine. The only available something was a person. The framework is, underneath, a method for locating the steps a narrow machine cannot reach.

General-purpose AI may perform input preparation, verification, integration, and exception handling itself. If it does, several of the layers the framework locates get attacked by the same system almost immediately, and the forecasts above about context production and exception work degrade badly.

What survives that assumption failing are the forecasts grounded in accountability rather than capability. Somebody has to be answerable to an outside party. That is a legal and social fact about who can be sued, fined, licensed, or fired, and it does not dissolve because a system got better.

The more general the automation becomes, the less safe it is to assume that today's complementary work becomes tomorrow's occupation.

Which leaves the durable version of the idea, and it is a smaller one than we expected to find. To forecast the work that automation creates, do not start with job titles. Start with what the new system consumes, what breaks when it scales, who is entitled to demand proof, and what becomes scarce once execution gets cheap.

Thirteen automation transitions between 1950 and 2011, each with a fixed historical cutoff shortly after commercial adoption. Structural forecasts were recorded before outcomes were examined, then scored against what occurred. A separate validation pass tested six cases against sources published at or before their cutoff, to distinguish constraints that were visible at the time from constraints that are only visible now.

LIMITATIONS

The framework identifies where new work forms and does not estimate how large it becomes. Five magnitude errors across thirteen cases, in both directions. Cases where automation produced nothing are systematically harder to find than cases where it produced something, so the hit rate reported here is inflated by an unknown amount. Several of the strongest results sit in domains with heavy public-sector involvement, which may reflect a real mechanism or may reflect better documentation. Where a regulatory or business-model change lands within a few years of an automation, the method over-attributes to the technology and has no internal way to detect it.

PRINCIPAL SOURCES

US Government Accountability Office (GGD-83-24, January 1983); ARINC Specification 424 (1975) and FAA Advisory Circular 20-153; US Bureau of Labor Statistics industry employment series (September 2025); ECR Retail Loss research on self-checkout; contemporaneous business press, 1974-2001; US Census occupational data. Full case files and evidence grading held internally.