

# From Chaos to Intelligence:

## Democratizing AI for the Connected Jobsite

*Why construction productivity has been flat for 70 years — and how letting humans speak, write, photograph, and upload in their natural way (while AI does the structuring) can finally close the data gap.*

Arcudi Alessio, Pernice Gabriella

---

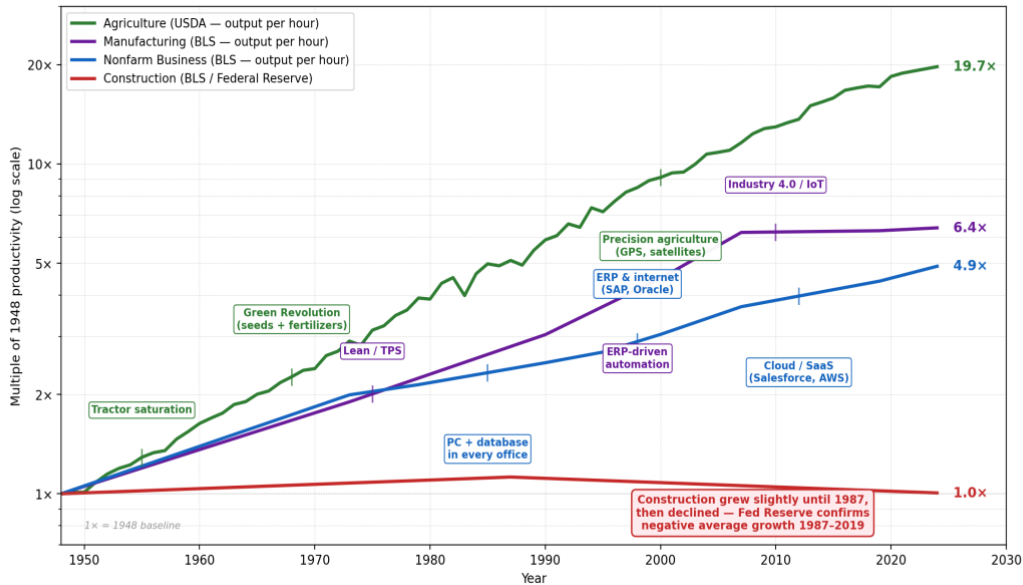
### 1. The Productivity Anomaly

Over the last 70 years, every major sector of the economy has had its breakpoint. Agriculture had several: tractors replaced animal power, then the Green Revolution introduced high-yield seed varieties and synthetic fertilizers, then GPS-guided precision agriculture turned every field into a sensor network. According to the U.S. Department of Agriculture, agricultural output per hour worked grew **by a factor of 17 between 1948 and 2017** (USDA ERS, 2022). Manufacturing absorbed Toyota's lean production system, then automation, then Industry 4.0; U.S. manufacturing productivity grew by roughly **6.4x** over the same period. Services and the broader nonfarm business sector digitized internal operations through ERP and CRM platforms (SAP, Oracle, Salesforce), then moved to the cloud, then layered analytics and AI on top of centralized data — productivity grew about **5x** (BLS, 2026).

Construction had none of this. In a recent working paper from the Federal Reserve Board, Garcia & Molloy (2025) state the finding bluntly: *“Of all major industries, construction is the only one to have registered negative average productivity growth since 1987.”* The authors tested every plausible source of measurement bias — unobserved quality improvements, deflator mismeasurement, structure quality changes — and concluded that even under generous assumptions, construction productivity growth has indeed been weak. Indexed to 1948, U.S. construction labor productivity grew slightly through the 1970s and 1980s, peaked around 1987, and has declined since. Net result today: approximately **1.0x** — the industry is right back where it started.

**Figure 1** shows the divergence. Every other sector has a story; construction has a flat line.

### U.S. Labor Productivity by Sector, 1948-2024



Sources: USDA Economic Research Service (Agricultural Productivity in the U.S., 1948-2021); U.S. Bureau of Labor Statistics (Productivity & Costs); Federal Reserve Board, Garcia & Molloy (2025, FEDS Working Paper 2023-052r1).

Figure 1. U.S. labor productivity by sector, 1948–2024 (multiples of 1948 productivity, log scale). Sources: USDA Economic Research Service (Agricultural Productivity in the U.S., 1948–2021); U.S. Bureau of Labor Statistics (Productivity & Costs); Federal Reserve Board, Garcia & Molloy (FEDS WP 2023-052r1).

This is not a controversial claim and it is not new. What is new is the magnitude of the underlying waste it implies. Construction is one of the largest sectors in the global economy, employing about 7 percent of the world’s working-age population and generating around \$10 trillion annually (MGI, 2017). McKinsey estimates that if construction productivity caught up with the rest of the economy, the industry’s value added would rise by \$1.6 trillion per year. The gap is also growing more expensive over time. Between 2000 and 2022, construction productivity improved only 10 percent in total — roughly 0.4 percent per year — versus a 50 percent gain in the total economy (McKinsey, 2024).

Construction is the second-least-digitized sector in the world on the MGI Industry Digitization Index, ahead only of agriculture and hunting (MGI, 2015) — and even agriculture has long since moved on. What makes this anomaly remarkable is that construction is not technology-averse at the output layer. Buildings today use BIM models, prefabricated modules, drones, and ground-penetrating radar. The Federal Reserve paper notes this directly: nail guns, prefabricated inputs, and information technology have all reached the industry, yet the productivity needle has not moved. The cutting edge lives on the drawings and in the factory. The problem lives somewhere else entirely.

## 2. The Office Door Problem

The most advanced digital tools in construction stop at the office door. On site, reality is still captured through calls, WhatsApp messages, photos stored in personal phone galleries, and handwritten notes that may or may not make it back to a spreadsheet. This is not a marginal observation: it is the operating reality for the vast majority of projects, including those run by large general contractors.

A useful diagnostic comes from rework. Studies put rework at anywhere between 5 and 20 percent of total project cost, with most research clustering in the 4–10 percent range (CII, 2011; Rhumbix, 2024). Recent peer-reviewed work by Love and colleagues at ASCE reached a more striking finding: **the rework costs that contractors typically record are only about one-quarter of the rework costs that are actually incurred** (Love, ASCE, 2026). When the same projects are tracked with rigorous discipline — every correction, every redo, every quietly-absorbed mistake — the real number is roughly four times higher than what gets reported. Of the rework that *is* recorded, 26 percent is caused by miscommunication and 14–22 percent by inaccurate information (PlanRadar, 2025).

*In other words: the single biggest source of construction waste is information that does not flow — and the industry literally does not know how much it is losing because the data never gets captured.*

This is consistent with what the adoption-barrier literature has been saying for years. Fieldwire’s analysis identifies five forces holding back construction technology, none of which are fundamentally about technology itself: cultural drag, broken upstream workflows, fear of disruption, lack of internal champions, and software that punishes the field user rather than helping them (Nitchen, 2025). HP calls the typical pattern “organ rejection” — a new tool fails to deliver visible benefits, workers notice, adoption collapses (HP, 2024). Zepth’s review converges on the same root causes: fragmented workflows that new tools cannot capture, mismatch between software design and site realities, unclear ROI for the end user, reversion to legacy tools (Zepth, 2026).

There is also a demographic dimension. The median age in U.S. construction is 42 and rising; fewer than 10 percent of workers are under 25 (Vizz Technologies, 2024). On many sites, a meaningful share of the workforce does not have the local language as a first language, and some workers have limited literacy in any language. Asking these people to “fill out the form in the app” is asking them to do three new things at once — operate an interface, work in a second language, and translate physical reality into a tabular schema — while standing on a noisy site under time pressure.

### **3. Interaction Cost is the Real Bottleneck**

From a Human–Computer Interaction perspective, the construction tech adoption problem is not mysterious. It is a textbook case of high interaction cost colliding with high cognitive load at the point of work.

Cognitive Load Theory (Sweller, 1988) predicts that when extraneous load — the load imposed by *how* information is presented or captured rather than by the task itself — exceeds working memory capacity, users disengage or revert to simpler tools. The Technology Acceptance Model (Davis, 1989) shows that perceived ease of use and perceived usefulness are the two strongest predictors of adoption; the Unified Theory of Acceptance and Use of Technology (Venkatesh et al., 2003) adds effort expectancy and facilitating conditions as critical moderators.

Most construction software fails on every one of these dimensions for the field user. It raises **input friction** (more taps, more fields, more menus than the original phone message it replaces), **cognitive friction** (the user has to remember which category, which form, which project, which template), and **coordination friction** (the captured data sits in a silo disconnected from the channels where decisions actually happen).

Distributed Cognition (Hutchins, 1995) offers the alternative framing: knowledge in a complex socio-technical system is not held by individuals but is distributed across people, artifacts, and environments. The job of a good tool is not to make every worker into a data-entry clerk; it is to capture knowledge *where it emerges*, in the form in which it naturally emerges, and consolidate it into a shared record.

This is the gap our software was designed for.

#### 4. The MELA Work Hypothesis: Let Humans Be Human, Let AI Structure

We have been training people to speak software for forty years — to remember which menu, which field, which dropdown, which category. The contrarian move, now newly possible, is to invert the labor: humans do what humans do naturally, and AI does the structuring.

The principle is simple: **the worker should never have to structure data for the system. The worker speaks, writes, photographs, uploads documents — whatever modality is natural in that moment. The LLM translates this unstructured input into structured, organized, analyzable records.** This is the entire architecture of MELA Work.

Concretely, a foreman opens the app the same way they open WhatsApp. They might:

- **Speak** — “we poured the third floor slab this morning, two delivery notes from the supplier, the rebar layout had a problem on the east side”;
- **Write** — a quick chat message, in their own language, using their own phrasing;
- **Photograph** — the site, the delivery note, the invoice, the damaged material, the completed work;
- **Upload** — a PDF from the supplier, a screenshot of a drawing, a voice memo recorded an hour earlier.

Behind the scenes, AI does the work that until now has been delegated to humans:

- **Natural Language Processing** parses informal updates — progress, blockers, requests for information — and turns them into structured project entries with the right categories, the right project assignment, the right date, the right scope code.
- **Computer Vision** extracts data from photos (site documentation, delivery notes, invoices) and converts them into searchable, auditable records and real-time cost intelligence.
- **Document understanding** ingests PDFs, scans, and screenshots and extracts the relevant information — vendor, amounts, line items, references — without manual transcription.

- **Multilingual handling** (30+ languages) means a crew that is Polish, Romanian, Albanian, Italian, and Arabic on the same morning can each report in their own language, with everything normalized into the same structured record.

The system never asks the user to remember categories or schemas. It infers them. It never requires typing on a jobsite where hands are dirty and cognitive bandwidth is going into not-falling-off-the-scaffolding. It works on the device every worker already has, in whichever modality is convenient at that moment.

What this changes, in HCI terms, is the entire cost structure of capturing field information. **Input friction drops** because the worker uses their natural modality instead of a form. **Cognitive friction drops** because the worker does not have to maintain a mental model of where things belong in some abstract schema. **Coordination friction drops** because the structured record is generated automatically and pushed into the systems where decisions get made.

This is a fundamentally different bargain than every previous generation of construction software. Earlier tools asked the human to do the structuring work because the machine could not. That constraint no longer holds. With modern multimodal LLMs, the structuring work is the part the machine is best at.

## 5. Data Collection is the Breakpoint

Here is the claim the whole paper has been building toward.

The reason agriculture had a productivity revolution between 1950 and 2000 was not just that someone invented the tractor. It was that mechanization made yield, cost per acre, soil chemistry, and weather *measurable at scale*. Precision agriculture went further: every square meter of a field became a data point, and once the data existed, optimization followed almost automatically. The breakpoint was not the machine. It was the data the machine generated.

The reason nonfarm business productivity jumped after 1995 was not the personal computer in itself. It was that ERP and CRM systems created what consultants call a “single source of truth” — a centralized, structured record of every transaction, customer interaction, and inventory movement. According to MGI, the most digitized sectors are also the most data-rich, and they consistently outperform on productivity. ERP-using companies report 20 percent reductions in operational costs (Digitrix, 2025); manufacturing operations adopting data-driven lean methodologies have achieved up to 100 percent improvements in operational excellence with a typical 30 percent cost reduction (Hexagon AI, 2025). The breakpoint was centralized data management.

The reason construction has not had its breakpoint is that construction has never collected the data. The information exists — every site generates an enormous amount of it every day — but it is dispersed across personal phones, voice calls, paper notes, and chat threads, never reaching a central record in a structured form.

Friction in the interface is therefore not just a user-experience problem. It is the *upstream cause* of the productivity stagnation. If the cost of capturing a piece of field information is higher than the perceived value of that information to the person capturing it, the information will not be captured, and the data-driven optimization loop never starts. Every other sector closed this loop in its breakpoint decade. Construction has not — because, until very recently, removing the human structuring cost was technically impossible.

## 6. What Doubling Data Collection Would Mean

Our deployments with contractors show that an AI-mediated approach — where humans interact naturally and the LLM does the structuring — can roughly **double** the volume of structured field data that reaches the central project record, because it removes the per-report cost almost entirely. Workers who would have made zero structured entries on a given day make several; reports that would have been one-line WhatsApp messages become structured records with photos, location, timestamps, and trade context attached.

This matters because, in the sectors that have already had their breakpoint, the relationship between data density and productivity is well-documented. The MGI Digitization Index shows a clear correlation between digital usage intensity and productivity growth at the sector level (MGI, 2015). Manufacturing studies tie data-driven decision-making to measurable improvements in defect rates, cycle times, and cost (Bloom et al., U.S. Census, 2016). The pattern is consistent: more granular, more frequent, more reliable data leads to better decisions, faster cycles, and lower waste.

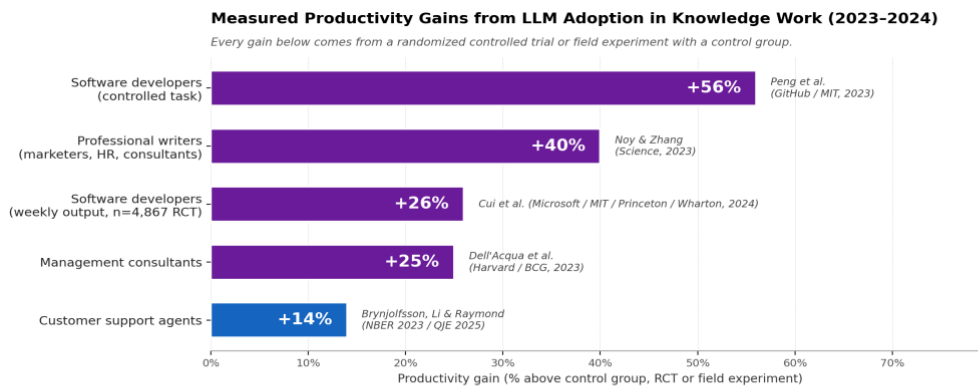
In construction, the same mechanism would attack the largest known source of waste — the 5–20 percent of project value that goes to rework, of which roughly a quarter is caused by miscommunication and another fifth by bad data (PlanRadar, 2025). If the data exists, miscommunication is detectable and correctable in real time. If the data exists, the post-mortem on a delay or cost overrun is no longer an exercise in archaeology. If the data exists, cost intelligence becomes proactive rather than retrospective. And the four-to-one under-reporting of rework documented by Love and colleagues — where actual rework is roughly four times what gets recorded — becomes a measurable, addressable problem rather than an invisible drain.

The reason this has not happened in construction before now is that until very recently it was technically impossible to collect this data without imposing prohibitive interaction costs on the field. Forms required typing. Categories required training. Languages required translation. Photos required tagging. Documents required transcription. Each of these costs was small in isolation and crushing in aggregate, and the field rationally rejected the tools. Modern multimodal AI removes every one of those costs: the human contributes raw input — speech, text, photos, documents — and the machine handles the rest.

## 7. The Proof Is Already In: What LLMs Have Done in Other Sectors

We do not have to speculate about what happens when LLMs remove the structuring cost from a workflow. It has already happened in several sectors, and the gains have been measured under rigorous experimental conditions.

**Figure 2** summarizes the results. Every number below comes from a randomized controlled trial or field experiment with a control group — not vendor claims, not surveys, not pilots.



■ Customer support: a "data-heavy" job — the bottleneck is producing structured records (tickets, notes, dispositions) under time pressure.  
 ■ Other knowledge work (writing, coding, consulting): the bottleneck is producing structured output (text, code, analysis).

Sources: Brynjolfsson, Li & Raymond (NBER WP 31161, 2023; QJE 2025); Peng et al. (arXiv:2302.06590, 2023); Cui et al. (Microsoft/MIT/Princeton/Wharton, 2024); Noy & Zhang (Science 381, 2023); Dell'Acqua et al. (HBS/BCG, 2023).

Figure 2. Measured productivity gains from LLM adoption in knowledge work, 2023–2024. Sources: Brynjolfsson, Li & Raymond (NBER 2023; QJE 2025); Peng et al. (2023); Cui et al. (Microsoft/MIT/Princeton/Wharton, 2024); Noy & Zhang (Science, 2023); Dell'Acqua et al. (Harvard/BCG, 2023).

- **Customer support agents (+14%):** Brynjolfsson, Li & Raymond studied 5,000 agents at a Fortune 500 firm using an LLM that monitored conversations and suggested responses. Productivity rose 14% on average, with the lowest-skilled workers gaining 35%. Customer sentiment improved and requests to escalate dropped 25%. The paper was published in the *Quarterly Journal of Economics* (2025).
- **Software developers (+26% to +56%):** A controlled experiment by Peng et al. (MIT/GitHub) found developers using Copilot completed an HTTP-server task **55.8% faster** than the control group. A larger RCT by Cui et al. across Microsoft, Accenture, and a Fortune 100 manufacturer (n=4,867 developers) measured **+26% weekly output**.
- **Management consultants (+25%):** Dell'Acqua et al. at Harvard/BCG ran a field experiment with 758 BCG consultants. Those with GPT-4 completed tasks **25% faster** and produced **40% higher quality** output. Lowest-performing consultants improved by 43%.
- **Professional writers (+40%):** Noy & Zhang, in *Science* (2023), ran a randomized trial on 444 college-educated professionals — marketers, HR, grant writers, consultants, data analysts. ChatGPT users completed writing tasks **40% faster** with 18% higher quality.

The common thread across all of these settings is the same one we have been describing for construction. In each case, the bottleneck was not “do the work” — it was **produce a structured output from informal human inputs**: a customer’s question turned into a ticket disposition; an idea turned into code; a meeting note turned into a strategic memo; a brief turned into a finished document. The skilled work — judgment, expertise, decision-making — remained with the human. The LLM absorbed the structuring labor.

Customer support is the closest analogue to construction. Both involve workers under time pressure who must produce a structured record from messy human-generated input, in environments where the tool's overhead has historically been the binding constraint on data quality. When the LLM took over the structuring work, customer support agents did not just work faster — they produced records that were more complete, more consistent, and more useful to the organization downstream. The data that was always supposed to be captured finally got captured.

The construction site has every condition for the same dynamic, only more so. The structuring cost is higher (more variables, more contexts, multilingual workforce, physical environment). The pre-existing data loss is larger (most field information never reaches a structured system at all). And the downstream value of the structured data is greater (rework, cost overruns, and schedule risk are all data-deficient problems waiting for the data to exist).

## 8. Conclusion: A Refinement of the Process, Not a New Tool

Construction is the last major sector that has not had its data-collection breakpoint. The reason is not that construction workers are technology-averse, that the industry is uniquely fragmented, or that ROI is unclear in the abstract. The reason is that, until now, every digital tool offered to the field has had a higher interaction cost than the analog alternative. People are rational; they have stuck with WhatsApp, calls, and paper because, for the specific act of capturing a piece of site reality, those tools have been better products.

The Federal Reserve paper on construction productivity made a quiet but important observation: the industry has had many genuine innovations over the last forty years — nail guns, prefabrication, BIM, GPS-guided equipment — and none of them moved the productivity needle. The lesson is not that innovations do not work in construction. The lesson is that **the bottleneck was never in the tools. It was in the process**, specifically in the information process that connects the field to the rest of the organization.

This is why the shift we are describing is not “another platform.” It is a refinement of the process — *who does the structuring work*. For four decades, software has required the human to do the structuring: pick categories, fill fields, translate physical reality into a schema. The modern bargain is the opposite: the human contributes raw input in whatever form is natural, and the AI does the structuring. The result is a simple causal chain:

***LLM-mediated input → frictionless data collection → centralized, structured data → optimization, productivity, and digital management at scale.***

Each link in this chain has been independently validated. The LLM-mediated input layer is now measurable across customer support, software, consulting, and professional writing (Section 7). Frictionless data collection is what doubles capture volume in our own deployments (Section 6). Centralized structured data is what gave agriculture, manufacturing, and services their productivity

breakpoints (Sections 1 and 5). And the chain only fails to close in construction because, until very recently, the first link was technically impossible.

**Sometimes technology is not a new tool. Sometimes it is a refinement of the process — a quiet change in which actor in the system does the cognitive labor.** Tractors did this to ploughing. ERP did this to operational reporting. Multimodal LLMs are now doing it to data capture in the field. The pattern is the same; the breakpoint just keeps moving forward, sector by sector.

We call our approach *scaffolding* deliberately. The AI is not the work; the work is still done by humans. The AI holds up the work so the human can do their actual job — building things — while contributing, in passing, to a data record that makes the whole system smarter. When that scaffolding is in place at the scale of thousands of workers across thousands of sites, the same dynamic that lifted agriculture, manufacturing, and services off their plateaus becomes available to construction.

***The breakpoint will not look like a tractor. It will look like a foreman talking into a phone, photographing a delivery note, forwarding a document — and a machine in the background turning all of it into the data the industry has needed for seventy years.***

---

## References

- Bloom, N., Brynjolfsson, E., Foster, L., Jarmin, R., Patnaik, M., Saporta-Eksten, I., & Van Reenen, J. (2016). Data in Action: Data-Driven Decision Making in U.S. Manufacturing. U.S. Census Bureau, CES Working Paper 16-06.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at Work. *Quarterly Journal of Economics*, 140(2), 889–942. (Working paper: NBER WP 31161, 2023.)
- Construction Industry Institute. (2011). A Guide to Construction Rework Reduction. Research Team 203.
- Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., & Salz, T. (2024). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. Microsoft / MIT / Princeton / Wharton.
- Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, 13(3), 319–340.
- Dell’Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Harvard Business School Working Paper 24-013.
- Digitrix. (2025). CRM & ERP Digital Transformation Roadmap for 2025.
- Fieldwire by Hilti / Nitchen, K. (2025). The 5 Barriers to Tech Adoption in Construction.
- Garcia, D., & Molloy, R. (2025). Reexamining Lackluster Productivity Growth in Construction. Finance and Economics Discussion Series 2023-052r1, Federal Reserve Board.
- Hexagon ALI Resources. (2025). How to Use Manufacturing Data to Make Better Decisions.
- HP Industrial Print. (2024). Mythbusting Digital Transformation in Construction.

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Love, P. E. D. et al. (2026). How Much Does Field Rework in Construction Actually Cost? ASCE Civil Engineering Source.
- McKinsey Global Institute. (2015). *Digital America: A Tale of the Haves and Have-Mores (Industry Digitization Index)*.
- McKinsey Global Institute. (2017). *Reinventing Construction: A Route to Higher Productivity*.
- McKinsey & Company. (2024). *Delivering on Construction Productivity Is No Longer Optional*.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. arXiv:2302.06590.
- PlanRadar. (2025). *Cost of Rework in Construction: Causes, Data & Prevention*.
- Rhumbix. (2024). *The True Cost of Construction Rework and How to Prevent It*.
- Sweller, J. (1988). Cognitive Load During Problem Solving: Effects on Learning. *Cognitive Science*, 12(2), 257–285.
- U.S. Bureau of Labor Statistics. (2026). *Productivity and Costs — long-term labor productivity by sector for selected periods*. <https://www.bls.gov/productivity/>
- U.S. Department of Agriculture, Economic Research Service. (2022). *Agricultural Productivity in the U.S. and “Increases in Labor Quality Contributed to Growth in U.S. Agricultural Output,” Amber Waves*.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425–478.
- Vizz Technologies. (2024). *What Is Holding Back Technology Adoption in the Construction Industry*.
- Zepth. (2026). *Why Construction Tech Always Fails at Adoption*.