

# Anthropic Just Published Honest Data Nobody in AI Wants to See

By Marco Kotrotsos

Medium

Mar 10, 2026

The company behind Claude used its own usage data to measure how AI is actually affecting jobs. The findings are uncomfortable for everyone, including Anthropic.

Most AI jobs research is speculation dressed up as analysis. Consulting firms estimate which tasks could theoretically be automated, map those tasks to occupations, and produce alarming charts. The methodology is sound enough. The problem is that theoretical capability and actual adoption are completely different things, and nobody had the data to measure the gap.

Until now. Anthropic researchers Maxim Massenkoff and Peter McCrory just published a paper that does something no AI company has done before: it combines what AI can theoretically do with what people are actually using it for. The distinction matters more than you'd think.

## The gap between “could” and “does”

Previous research from Eloundou et al. estimated that AI could handle 94% of tasks in computer and math occupations. Anthropic's observed data shows that only 33% of those tasks are actually being performed by Claude in practice. That's a 61-percentage-point gap between what's possible and what's happening.

Office and administrative work show a similar pattern. Theoretical capacity sits at 90%. Actual usage is a fraction of that.

This gap is currently the thing keeping most knowledge workers employed. Not their irreplaceability, not the quality of their work, but the fact that organizations haven't figured out how to deploy AI at scale yet. Legal constraints, technical limitations, integration costs, human review requirements, all of these create friction that slows adoption. The researchers expect every one of these barriers to diminish over time.

Read that again. The moat protecting your job isn't your skill. It's organizational inertia. And inertia, by definition, is temporary.

## Who's actually exposed

This is where the data gets uncomfortable.

Anthropic ranked occupations by “observed exposure,” meaning the overlap between what Claude can do and what people are actually using it for in practice. The top ten:

1. Computer programmers, 75%
2. Customer service representatives, 70%
3. Data entry keyers, 67%
4. Medical record specialists, 67%
5. Market research analysts and marketing specialists, 65%
6. Sales representatives, 63%
7. Financial and investment analysts, 57%
8. Software quality assurance analysts, 52%
9. Information security analysts, 49%
10. Computer user support specialists, 47%

The occupations with zero exposure? Cooks. Motorcycle mechanics. Lifeguards. Bartenders. Dishwashers. Dressing room attendants. If your job requires you to physically be somewhere and do something with your hands, you're fine. If your job involves processing information, analyzing data, writing code, or communicating with humans through text, you're on this list.

The demographic profile of exposed workers makes the pattern even clearer. Workers in the most exposed occupations earn 47% more than workers with zero exposure. They are 16 percentage points more likely to be female. They are nearly four times more likely to hold a graduate degree (17.4% vs 4.5%).

AI isn't automating low-wage manual labor. It's automating educated, well-paid knowledge work. The people who spent the most on their education and built the most specialized careers are the ones most exposed.

### **The employment signal that matters**

The headline finding is measured and careful: “no systematic increase in unemployment for highly exposed workers since late 2022.” People in AI-exposed jobs are not losing their jobs at elevated rates.

Not yet?

But the researchers found something else. For workers aged 22–25, hiring into AI-exposed occupations has dropped approximately 14% compared to 2022 levels. Less-exposed occupations maintain a steady 2% monthly job finding rate. The most exposed occupations have seen that rate decline by about 0.5 percentage points. No similar decline exists for workers over 25. This is the pattern that should concern everyone paying attention.

The jobs aren't disappearing for the people who already have them.

They're disappearing for the people trying to break in. If you're a senior software engineer, your job is safe for now. If you're a recent graduate trying to land your first programming job, you're competing against tools that didn't exist when you chose your major. The report frames this carefully: young unemployed workers may be staying in existing roles, pursuing alternatives, or returning to school. But the signal is there. The entry-level funnel is narrowing.

## **The scenario nobody wants to model**

Buried in the paper is a thought experiment. The researchers describe what would happen if AI capabilities expanded to fully automate the most exposed occupations. Not all at once, but a sustained displacement similar to what happened during the 2007–2009 financial crisis.

During the Great Recession, US unemployment doubled from 5% to 10%. The researchers note that a comparable shock is “absolutely possible” in the white-collar labor market if AI adoption catches up to AI capability.

Right now, 33% of theoretical capacity is being utilized. If that number moves to 66%, the disruption hits a different scale entirely. And there's nothing in the data suggesting the gap will stay this wide. Every barrier the researchers identify, legal constraints, integration costs, human review requirements, is a barrier that companies are actively working to remove.

Anthropic CEO Dario Amodei has said publicly that AI could disrupt half of entry-level white-collar work. Microsoft AI Chief Mustafa Suleyman said most professional work could be replaceable within 12–18 months. These aren't critics or skeptics making these claims. These are the people building the technology.

## **What 97% means**

Here's a number that didn't get enough attention: 97% of the tasks people actually use Claude for fall into categories that are theoretically feasible for the model. Of those, 68% involve tasks where Claude can fully handle the work without human assistance. Only 3% of observed usage involves tasks that AI fundamentally can't do.

People are already self-selecting toward tasks AI can handle. They're not using Claude for the hard stuff and failing. They're using it for the stuff it's good at and succeeding. That means the 33% observed-vs-theoretical gap isn't about capability limitations. It's about deployment friction. The AI can already do the work. Organizations just haven't set up the pipelines to let it.

When those pipelines exist, the 33% becomes 50%, then 70%, then the

scenario nobody wants to model.

### **What the report doesn't say**

Anthropic deserves credit for publishing data that undermines its own growth narrative. A 14% decline in entry-level hiring is not the kind of finding that makes enterprise customers eager to expand their Claude deployment. The fact that AI disproportionately affects educated, higher-earning workers contradicts the comfortable story that AI mainly eliminates tedious, low value tasks.

But the report also has limitations it acknowledges openly. The usage data comes from Claude specifically, not from the entire AI ecosystem. GPT, Gemini, Copilot, and dozens of specialized tools handle tasks that don't show up in Anthropic's data. The real observed exposure across all AI tools is higher than what this report captures.

The unemployment analysis can detect differential increases of about 1 percentage point. Smaller shifts, the kind that accumulate slowly before anyone notices, would be invisible in this framework. By the time the signal is clear in the data, it may be too late for the workers it affects.

And the report measures displacement at the occupation level, not the task level. A software engineer who spends 30% less time writing boilerplate code doesn't show up as "displaced" in unemployment statistics. They show up as "more productive." Until their employer realizes they need three engineers instead of five, at which point two of them show up in the data, but only after the decision was already made.

### **The practical takeaway**

The researchers explicitly frame their work as "a first step toward cataloging the impact of AI on the labor market."

First step. They built an early warning system, and the early warning is: the disruption is not theoretical, it's measurable, it's concentrated in educated knowledge work, and it's hitting young workers first.

If you're in one of the exposed occupations, here's what the data actually tells you.

Your current job is probably safe in the short term. Existing workers in exposed roles show no elevated unemployment. The organizational friction that slows AI adoption is real, and it buys time.

**Your next job? That might not exist anymore.**

The entry-level pipeline is already narrowing. If you're planning a career transition or expecting to move laterally into a similar role, the assumption that equivalent roles will be available in 2–3 years is not supported by the trend.

The value of your expertise is shifting. The gap between theoretical capability and observed usage exists because someone still needs to direct, verify, and integrate AI output. The people who understand both the domain and the tools are the ones who will bridge that gap. Pure domain knowledge without AI fluency, and pure AI fluency without domain knowledge, both lose.

Organizations are slower than you think, but faster than you hope. The 33% utilization rate will not stay at 33%. Every quarter, the friction decreases. Every new integration, every new API, every new deployment pattern makes it easier to close the gap. The question isn't whether your industry will be affected. It's whether you'll be ready when the adoption curve catches up to the capability curve.

Anthropic built a thermometer and stuck it in the labor market. The reading says the fever is low but rising. The responsible thing to do is not to ignore the thermometer because the number isn't scary yet. It's to start building immunity while you still have time.

The full report: [anthropic.com/research/labor-market-impacts](https://anthropic.com/research/labor-market-impacts)

*Marco Kotrotsos, specializing in practical AI implementation for organizations ready to close the gap between AI hype and AI value. With 30 years of IT experience now focused purely on AI deployment, he works hands-on with companies to turn AI potential into measurable business outcomes.*