

Economics from the Top Down

new ideas in economics and the social sciences

The Laws of Useful Automation

Blair Fix

October 03, 2026



Given our current landscape of AI hype and doom, everyone seems to have an opinion about whether chatbots are ‘good’ or ‘bad’. Fewer people have thought about how chatbots fit into the long-term history of automation. And even fewer folks have reflected on the conceptual requirements that make automation useful. Here are my thoughts on this latter topic.

In my view, automation is useful when it meets two criteria:

1. The *product* is more important than the *process* that creates it.
2. The product can be validated *without* auditing the creation process.

If we look at successful forms of machine automation, they tend to meet both criteria. For example, think of a pencil factory that automates the production of pencils. Here, the pencil is the ‘product’, and whatever happens inside the factory is the ‘process’.

Now for the end user, the process of pencil production is largely irrelevant. It doesn’t matter if the pencil is hand crafted, built in an assembly line, or conjured by a Star Trek [replicator](#). As long as the pencil works in the hands of the user, its creation process is unimportant.

That brings me to the second requirement for useful automation: the product must be easily validated. To validate that a pencil works, you don’t need to audit the process that created it. You just give the thing a go. Yup, this pencil works. Nope, that one’s a dud.

Thinking more broadly, the history of automation has been dominated by this sort of process where a machine replaces human labor in the production of some sort of physical commodity. The automation works because (1) the

commodity is valued more than the process that creates it, and (2) it's easy to verify the quality of the commodity without auditing the entire chain of production.

Intellectual automation

Looking to more recent technological progress, it turns out that successful automation need not be physical. Intellectual automation can also be useful. Just look at *computers*, which owe their name to the desire to automate computation.

Many forms of computation pass our automation requirements: the product is more important than the process, and the process is easy to validate. For example, when I ask a computer to calculate $\sqrt{200}$, I don't care about the algorithm it uses, or about the details of its chipset. I just want the answer. Likewise, once I have the result, I don't have to audit the computer's innards to validate the answer. I just check that the number squares to give 200.

In more general terms, any intellectual task that meets our two requirements is game for computer automation. But for years, some forms of intellectual work seemed beyond the reach of computers. Writing code is a good example. Historically, programming was a tedious task that took years of specialized training. But with the rise of neural nets and their associated chatbot interfaces, these barriers are being torn down. Chatbots can now code more quickly than any human. But is this ability useful?

Well, the answer depends on whether the application passes our two requirements. (Is the product more important than the process? And can the product be verified without auditing the process that created it?) Clearly, many forms of coding pass this test. Web design is an obvious example. When a blogger asks for a nice website, they usually don't care how it's accomplished. Likewise, the blogger can easily tell if the final design is what they want. (They just browse the website.)

In short, chatbots are useful for automating intellectual tasks like web design (and any other easily verified piece of software). Yes, the bots will put some folks out of work. Yes, they'll raise questions about the skills required in the workforce. And yes, they'll be used in ways that [undermine labor power and harm workers' health](#). But these issues are nothing new. They're a historical

feature of all forms of automation. What interests me more (at least in this essay) is the ways in which chatbot automation might be fundamentally useless, or even downright harmful.

On the useless front, I'm skeptical that chatbots can be used to fully automate scientific analysis. Here's why. When a scientist analyzes their data, they might think that the product of their inquiry is the 'results' section in their published paper. And in some sense, that's true. But the (big) caveat is that the usefulness of this result depends on whether the analysis pipeline is *correct*.

Now, suppose that a scientist automated their work by getting a chatbot to code the entirety of their analytic pipeline. How does the scientist know that their results are correct?¹ Well, if the analysis is complicated, the only sound way to assess its correctness is to audit the underlying code. That requires significant time and skill. Meanwhile, this time-skill investment largely defeats the purpose of automation. Of course, chatbots can no doubt help scientist solve specific coding problems. (These bots lower the barrier to successful programming.) But the notion that chatbots will fully automate scientific analysis is, frankly, laughable.

True, some academics will surely try this fully automated approach. In fact, I expect that the scientific literature will become increasingly polluted with bot junk. But I'd argue that we can't blame chatbots (solely) for this pollution.

¹Of course, the truth is that even the best-trained scientists make mistakes, which is why good science requires replication. And regarding code, there's an old saying that you shouldn't reinvent the wheel ... don't recode an algorithm that someone has already solved. Which is why scientific code is typically full of libraries and functions which the scientist in question did not write.

For example, when I get R to calculate a matrix inverse, I'm actually calling ancient Fortran code for doing linear algebra. To me, this code is a black box — I have no idea how it works. So how do I know that this code works correctly? Honestly, it's a matter of trust. These libraries are free and open source, and have been used for ages. If they had a gaping problem, scientists would not use them.

Now, this game of trust comes on a continuum. I greatly trust R's matrix inverse functions. I put less trust in code from a random Github repository. And I put even less trust in the code delivered by a chatbot. Sure, the chatbot code may be 99% good. But that 1% bad stuff can be a killer. Imagine a world in which all scientific analysis and all scientific libraries were coded by chatbots with no supervision from scientists. Each time a bot calls a Python or R library, there's a 1% chance of error. As the code expands, the error compounds, to the point that virtually everything the bots spit out is wrong.

Still, chatbots shine in the domain where automation has always been useful — when the results are easily verifiable. Refactoring code is a good example. If I do an analysis in R and someone else wants to refactor the code into Python, a chatbot could make short work of the task. Sure, the chatbot might make mistakes along the way, but the user could tell by comparing the R output to the Python output.

The root problem is the incentive structure in universities — a structure that values the production of academic papers far more than the process that creates them. But when it comes to good research, its social value lies entirely in the scientific *process*.

Process dependence

Like science, many areas of human life have a process dependence, in which the usefulness of a ‘product’ hinges solely on the process of doing it. Schooling is the most ubiquitous example. When a teacher asks students to solve a math problem, the product of this task is the correct answer. But the usefulness of this activity lies mostly in the process of doing it. Sure, a calculator will tell you the sum of $21 + 54$. But if the goal is to learn basic arithmetic, the use of a calculator is not ‘automation’. It’s *cheating*.

The obvious conclusion is that skill acquisition cannot be automated. If the goal is to learn basic arithmetic, a calculator is self-defeating. If the goal is learn to the principles of English spelling, a spell checker is unhelpful. If the goal is to learn to read, a text-to-voice processor is educational sabotage. And if the goal is to learn to write, well, chatbots are your mortal enemy.²

In this light, we can think of formal education as a teaching method that forces students to re-experience, in a curated and abbreviated form, problems that long stumped our ancestors. For example, it took thousands of years of doing arithmetic before humans automated the job with calculators. By then, mathematics was a mature field. When today’s students learn math, they replay this history over the course of a few years. At first, ‘math’ consists solely of raw computation. Later on, students learn more symbolic logic, and the number crunching gets delegated to machines. In short, when it comes to intellectual automation, everything hinges on the order of operations. First, you learn a difficult skill; then you discover that you can automate it.

Do not automate

Thinking further about process dependence, it seems likely that some intellectual tasks should *never* be automated. Writing is the most obvious example.

²In schools systems, teachers often speak in the language of ‘accommodations’ — as in text-to-voice is an ‘accommodation’ for dyslexia. While this use of technology is well intentioned, it’s also tragic. The truth is that if a teenager cannot read effectively, *all* available resources should go into solving this [highly solvable problem](#). In my view, it’s unethical to forge ahead with other curriculum in the face of gaping illiteracy.

To understand why automated writing is bad, we need to first deal with the overloaded nature of the English language. In English, the word ‘write’ has a misleading double meaning. In one sense, to ‘write’ means to ‘scribe’ — to put an already existing sentence onto paper. This form of ‘writing’ takes skill, but is grounds for useful automation. Once upon a time, video captions were transcribed by a human listener. But today, the captioning can be generated by natural language processors.

The problem with automated ‘writing’ comes with the second meaning of the word. To ‘write’ is not just to scribe; it’s also to craft a set of coherent arguments that other humans can follow and understand. To automate this activity is an oxymoron, because the product (a rational argument) can’t be separated from the process itself. Or as my mentor Jonathan Nitzan once told me, “You don’t really know what you think until you write it down.”

Here’s what he means. ‘Writing’, in this sense of the word, is essentially codified thinking. When an idea is written down, reading this idea leads to all kinds of interesting consequences. Often, the writer realizes that the idea is vague or incomplete. And so they revise it until things make sense. Once the idea is coherent, rereading it prompts new ideas and new connections. Many are dead ends. Some are fruitful.

To be frank, this iterative process can be torturous. When I write blog posts about my research, the final essay usually conceals an iceberg of revised or discarded thought. Sure, I’d like to avoid this torture and still have the final well-argued essay. But to avoid the torture of ‘writing’ is to avoid the discomfort of *rational thought*. Automating this task does not ‘save time’; it saves us from *thinking*.

Automation for whom?

For automation to be ‘useful’, the product must be more important than the process by which it is created. But there is a sticky question that I’ve so far avoided: more important *for whom*?

For the user of a pencil, the way that this commodity was manufactured is largely irrelevant. But for the folks who live beside the pencil factory, the manufacturing process is often more salient than the product itself, particularly if pollution is involved. Likewise, chatbots might be great for the Silicon Valley programmer, but they’re a Faustian bargain for the utility planners who have to power local data centers.

When it comes to the big picture of automation, the process by which it occurs is incredibly consequential ... often far more so than the product being automated. It's one thing to have a Star-Trek-like computer powered by nuclear fusion; it's quite another to have a sycophantic chatbot powered by fossil fuels.

Unfortunately, we humans are notoriously bad at assessing the big-picture consequences of our automation schemes. As a rule, we build first and ask questions later. Today, we seem to be automating tasks that should not be automated (using fuels that are steadily spoiling the planet). The internet is increasingly littered with chatbot slop, to the point that search engines can feel pointless. If a search query returns page after page of bot slop, it's obviously more sensible to pose the question directly to a chatbot. But if the chatbot is trained on bot-slop, how can you trust its answer?

Now, I'm personally skeptical of claims about AI-driven doom, but mostly because they go in the wrong direction. If there's a risk that AI will kill industrial civilization, it's not because the machines will take over. Far more likely, in my opinion, is that we get a future that looks like an *anti-singularity* — a future in which our technology becomes so powerful (and so polluting) that it gradually undermines its own existence. As humans subcontract thinking to fossil-fuel-fed chatbots, the information environment (as well as the natural environment) becomes polluted with slop, and the slop-trained bots themselves grow increasingly senile. “Water the crops with [Brawndo](#),” the bots say. Meanwhile, AI-driven education has left humans gullible enough to listen.³

In short, automation can be useful if it frees us from drudgery (in a way that doesn't destroy the earth) and leaves more time for creative thought. But if automated chatbots [gobble fossil fuels](#) in order to liberate us from thinking, well, civilization had a nice run.

³If there are long-term benefits to chatbots, they will come by strategically *withholding* chatbot use while students learn difficult and uncomfortable skills. Actually using a chatbot takes about as much skill as using a calculator. Which is funny, because no one would propose a ‘calculator-driven education’. But many folks will claim that AI is going to revolutionize schooling. Well, I work in high schools and can tell you that so far, what's been ‘revolutionary’ is that chatbots have killed the take-home essay. Learning to craft long-form thought used to be a standard feature of high-school education. Now it's not.

Support this blog

Hi folks, Blair Fix here. I'm a crowdfunded scientist who shares all of my (painstaking) research for free. If you think my work has value, consider becoming a supporter. You'll help me continue to share data-driven science with a world that needs less opinion and more facts.



Further reading

Doctorow, C. (2026). *The reverse centaur’s guide to life after AI: How to think about artificial intelligence before it’s too late*. Verso Books.