

Artificial intelligence isn't a good argument for basic income

A major study backed by OpenAI's Sam Altman shows that cash could be a great tool against poverty. No need for fears of runaway AI.

By Oshan Jarow

11 min. read · [View original](#)

We're [flooded](#) by guaranteed income pilot experiments that offer some promising results, but [don't seem](#) to be moving us any closer to actual federal policy. Yet findings published today from the largest randomized basic income experiment in the US to date, backed by Sam Altman and OpenAI, should get your notice.

The study, held from November 2020 through October 2023, gave 1,000

recipients \$1,000 per month, no strings attached. It's one of the biggest and longest trials ever run on direct cash giving. Many other [basic income pilots](#) have given people \$500 or less, and rarely for more than a year or two.

While the study was run by a group of academics, it was [set in motion](#) by OpenAI CEO Sam Altman. For years, Altman has been publicly worrying that basic income will become necessary as AI eliminates traditional jobs while creating huge stockpiles of wealth held by a few. "If public policy doesn't adapt accordingly, most people will end up worse off than they are today," [he wrote](#) in 2021. (Disclosure: Vox Media is one of several publishers that have signed partnership agreements with OpenAI. Our reporting remains editorially independent.)

Altman isn't alone. Many major figures in the tech world, from [Elon Musk](#) to the "godfather of AI," [Geoffrey Hinton](#), believe AI will usher a wave of technological unemployment, and basic income will become necessary to keep us all afloat.

On this, I'm conflicted. In general, unconditional cash can be [an effective anti-poverty policy](#), but the extensive list of small experiments and accumulating evidence doesn't seem to be furthering the case of basic income to becoming a reality. So maybe any press is good press, and all the attention AI leaders are bringing to basic income will help it along. Spreading the fear of AI stealing all our jobs seems to garner much more enthusiasm for basic income than continually [pointing at the evidence](#). (Just look at Andrew Yang's [2020 presidential campaign](#), which centered on universal basic income as a response to automation.)

But hitching [the case for basic income](#) to fears of rapid AI progress makes it far more vulnerable than it needs to be. If there's no great wave of AI-driven unemployment, if the AI [bubble bursts](#) and turns out to be [hardly any different](#) than cycles of innovation and technological unemployment in the past, support for basic income would fall, too.

Yet the arguments for some kind of basic income are strong, no matter what

happens with AI. There could be absolutely no further progress, AI could come to a complete and eternal standstill, and the case for basic income would remain as strong as ever: unconditional cash can be a simple and flexible way to style income-support policies that actually reach everyone in need (at the trade-off of higher taxes).

The basic income movement might be better off severing ties with speculations about AI altogether. Then, the conversation could focus on what basic income can actually be: an effective anti-poverty tool that would neither stave off dystopia nor usher in a leisurely paradise, but instead, just [a world with less poverty](#).

News from the largest randomized study of basic income in the US to date

Many cash transfer advocates feel that the movement is [overdue to graduate](#) from research to actual policy. But the “Unconditional Income Study,” as this three-year endeavor by the nonprofit [OpenResearch](#) has been dubbed, added some welcome substance to the otherwise slim pile of research on big

cash transfers over long periods of time. (OpenResearch is not affiliated directly with OpenAI, but it grew out of the tech accelerator Y Combinator when Altman ran it, and has received a combined \$24 million from Altman and OpenAI's nonprofit organization.)

The study gave out monthly checks to people between the ages of 21 and 40 living in Texas and Illinois. To qualify, their 2019 household income had to be less than 300 percent of the federal poverty line: that would mean \$77,250 for a [family of four](#), or \$37,470 for an individual. The average participant's household income in 2019 was about \$30,000. One thousand people were randomized into the treatment group and received the full \$1,000 per month, while another 2,000 were part of a control group that got \$50 per month.

The results are coming out in batches, with [two papers](#) out today — one on [employment effects](#) and another on [health outcomes](#) — and another on consumption still in the works. So far, the results are kind of scrambled, refusing to

neatly fit into any particular view on basic income.

Recipients spent an average of \$310 more per month, mostly on housing, food, and car expenses. Overall, however, their incomes fell by about \$125 per month, excluding the transfers. The drop in earnings was largely driven by people choosing to work a little less (since the transfers meant they still came out ahead). In total, labor market participation declined by 2 percent, cashing out as working about 1.3 hours less per week, or roughly eight fewer days of work over the course of a year.

At first glance, this finding is somewhat at odds with [other basic income experiments](#) in the US, and challenges [the narrative advocates](#) have been pushing that unconditional cash won't make people choose less work. Sometimes, maybe it will. But eight fewer days of work per year doesn't exactly look like the doomsday scenario of everyone choosing to live out their remaining days on welfare, [surfing](#), or playing video games on the taxpayer's dime.

The cash's impact on health was a little more straightforward: it didn't seem to do much. Despite detailed data, including blood samples from some participants and nutritional intake, "We find essentially no evidence of improvements in physical health due to the transfers," the paper states, "and can rule out even small improvements."

There were short-lived improvements in mental health, mostly through reductions in stress. But they faded by the end of the first year. Recipients also showed a major increase in food security during year one, which also faded by the second year. If there was any silver lining in the paper on health, it was that recipients went to hospitals, doctors, and dentists more often, and spent about \$20 more per month on medical care as a result of the transfer, which the authors note could translate into indirect benefits in the long term.

"The appeal of cash transfers lies in the freedom that they give beneficiaries to make their own choices about what type of consumption to prioritize," the authors conclude. "However, the nature of that

freedom means that cash transfers are a blunt instrument for improving health and reducing health disparities specifically.”

The results suggest that policymakers interested in improving particular health outcomes might be better off pursuing more targeted interventions, like expanding Medicaid eligibility, reducing prescription drug costs, or making it easier to make primary care appointments. Targeting health directly, in other words, rather than simple cash.

As the researchers argue, “if policymakers seek to reduce poverty and improve the well-being of low-income groups more generally, cash transfers could still represent an important option, even if they do not improve health outcomes in the near term.”

Basic income restored a little bit of capitalism's forgotten promise: the freedom to choose more leisure time

One of the most controversial and politically charged parts of the basic income debate is how it would affect work. Among critical economists and

policy works, the [concern](#) isn't that we'll need basic income because there won't be any jobs left after AI, but that basic income could itself be the thing that [wipes away](#) our desire to work.

Whether these results support that fear depends on how deeply you read into the details. The results clearly show that people do work a little less. But going a bit deeper, the employment effects were concentrated among young, single parents. Does the option for single parents to choose an extra eight days off per year constitute a policy failure — or success?

The project also captured pretty detailed time-use data, letting researchers figure out where recipients were re-allocating their newly chosen time away from work. Though 1,000 people obviously made a variety of choices over the course of three years, the dominant trade was clearly toward more leisure time.

I've [previously written](#) about how, until recently, Americans across the political spectrum used to see the promise of capitalism as exactly that: delivering people the choice of cashing in

economic progress for more leisure time. Getting an economic boost and subsequently choosing to work a little less was not seen as a policy failure, but [the whole point](#).

“But for decades, workers haven’t even been getting that choice because, for the most part, productivity growth has ended up as higher profits and more inequality,” sociologist Aaron Benanav [told me last year](#). “It takes a political movement to cash that out in terms of a reduced workweek.”

Apparently, a generous basic income can at least help nudge that choice back into workers’ lives. After getting an extra \$1,000, recipients effectively traded in \$125 to get five-and-a-half more hours per month outside of work.

Interviews with participants paint a much brighter picture than the numbers

In addition to the survey data, OpenResearch held voluntary interviews with participants after the transfers ended to hear more about their experiences. Though only nine were recorded, I kept hearing variations on a

theme that wasn't captured in the statistics. (All names were changed by OpenResearch to protect the identities of participants.)

As Alina, a single mom living in Chicago, put it: "My whole mind kinda began to shift ... I had to stop myself so that I could have the mental mind shift, the paradigm shift to say, ok, we can actually engage in system planning."

Or Celene, who lived with her ex-wife and their kids in a friend's basement before the transfers helped them get their own home: "It helped my mind clear a little bit because I didn't feel like as much of a failure."

Or Dominic, who lost his job during the pandemic: "The biggest thing it did was allow me a time of peace of mind so I could explore what I really wanted out of life."

Everyone seemed to mention how getting unconditional cash was literally changing their minds for the better.

The mechanism doesn't have to be all that mysterious. Everyone faces their own

challenges, and people with low incomes especially so. But woven through all the differences is at least one shared consequence: stress. And we already know that stress can literally warp the mind, [especially around scarcity](#). Giving people cash comes with the flexibility to address whatever sources of stress are loudest in their lives, freeing up the mind.

And yet, according to the numbers, the sharp reductions in stress that recipients had in the first year didn't last, even though they shared these powerful stories after the third year. Which raises an interesting tension: Why is there such a big difference between the extremely positive stories that people are telling about their experiences, and the story that the numbers tell, which basically amounts to a shrug?

Part of it could be sampling bias: since the interviews were voluntary, maybe only people who had powerful and positive experiences opted in to the voluntary interviews to talk more about them. More broadly, though, there are parts of the human experience — and more

specifically, the experience of being a participant in a basic income experiment — that statistics either fail to capture, or even conceal.

When I spoke with Elizabeth Rhodes, who has a joint doctorate in social work and political science, and directed the study for OpenResearch, she explained that all the variation across how people used the money, or what a “good outcome” actually is based on someone’s particular circumstances, made them difficult to capture in top-line statistics.

One person might wind up finding [better work and taking on more hours](#) thanks to the cash, while another might scale back their hours to spend more time with their kids. Both cases can be good outcomes to the individuals involved, but “as an average treatment effect, where one person worked a lot more, while another worked less, it all kind of averages out.” she said.

We should insulate the case for basic income from any potential AI bubbles

Despite receiving funding from much of the tech world where talk of AI is

everywhere, none of the papers published from the Unconditional Income Study mentioned AI at all. That's a good thing, because as I mentioned earlier, we should insulate the case for basic income from the volatility and unknowns of AI.

Against the backdrop of all the “changing the world as we know it” talk, a number of analysts are beginning to think the AI [bubble](#) is already [beginning](#) to [burst](#). Maybe it will, and then it'll inflate again. And maybe one of these go-rounds will actually lead to changes that go beyond the cycles of innovation and technological unemployment [we've weathered for centuries](#). Or maybe they won't. Who knows!

The strongest arguments for unconditional income don't need to rest on speculation. We have a [lot of evidence](#). We have real-world experience from [a ton of pilot experiments](#) and [lapsed national policies](#). And we already have [an urgent need](#) for better anti-poverty policies, even if there wasn't an iota more of AI development.

Breaking the link between basic income and automation would also force us to consider what — if not UBI — actually would be a good policy response to major technological unemployment. Because if we really do all lose our jobs to AI, \$12,000 per year isn't going to cushion the blow much. This is a persistent confusion in debate around basic income: if you're talking about unconditionally giving people enough to live on indefinitely, you're either imagining a society of monks, that everyone lives in rural West Virginia, or you aren't actually talking about basic income at all, at least as it exists in any viable form.

The average living wage for a family of four in the US is about \$104,000 per year, according to the MIT [Living Wage calculator](#). Even if a guaranteed income on the scale of this study gave the full amount to children (which most proposals do not), the same family would only get \$48,000 per year, less than half the living wage. There is no realistic guaranteed income proposal that could serve as a substitute for wages altogether.

Altman himself [has argued](#) for something more nuanced in the past: taxes on AI companies and land values that feed into a [social wealth fund](#) where every citizen holds an equal share, receiving a dividend that rises or falls with the fund value.

That's basically a version of UBI that's styled away from guaranteeing an end to poverty, and toward being able to adapt to [runaway economic growth](#). In addition, maybe we'll need to consider [job guarantees](#). Or maybe compulsory employment is an iron cage we'll finally break by more deeply [democratizing](#) the economy, and whatever diminishing pools of work remain.

Either way, basic income advocates are left with a choice. You can welcome fears around AI as an incredibly effective marketing tool to spread awareness for basic income, and then try to jiu-jitsu that attention toward sturdier motivations, like [more effectively reducing poverty](#) than a patchwork of means-tested programs. Or, like me, you can try shouting into the void that we shouldn't associate basic income and

automation in the first place, since it sets the case for unconditional cash on shaky foundations, while there's [an empirically sturdy basis](#) already established.

With the OpenResearch studies, that basis continues to grow. Whether or not that'll lead to any policy breakthroughs before AI flops, fades into being just another new technology that shifts things around for a bit, or really does transform the economy, who knows?