



The End of Required Work: Universal Basic Income and AI-Driven Prosperity

How a tax on AI work might let everyone share the imminent bounty from AI automation.

Publication note

Originally published in [Towards Data Science](#). Author's copy on [Substack](#).

A previous version of this article was originally published in [Towards Data Science](#).

While the timeline remains somewhat uncertain, the overwhelming consensus is that machines will eventually take over most work. Rather than getting bogged down in a debate about exactly when this will happen, in two years or in twenty, it's more productive to ask what we can do to adapt and thrive in a world where work by artificial intelligence (AI) systems is growing increasingly prevalent, often at the cost of human jobs.

As a society, we are faced with a stark choice of what will happen when the machines eventually take over most of our jobs: will people be left destitute, or will they enjoy leisure time and be empowered to live life on their own terms? The latter scenario isn't just about not working. It's about having the freedom to pursue what brings us joy, without the burden of financial stress. It's about having the time and resources to explore our passions, travel, learn, and grow, without the constant pressure of making ends meet. In this scenario, people would still be free to work, but only because they want to, not because they have to. They could also pursue their own entrepreneurial ventures, whether a side hustle or a full-fledged startup, without the fear of falling into poverty hanging over them.

The fundamental challenge we face is figuring out how we can take the huge bounty that will eventually be produced by automated AI systems and share it in a reasonable way. In the US, our current tax structure is based primarily on taxing the work of individuals through income tax, but as more people inevitably lose work, that source of revenue is going to shrink while at the same time the need for public support will grow. *Instead of taxing a shrinking pool of human work, one possibility is that perhaps we should tax the growing work of AI systems instead.*

The Need for Universal Basic Income

Universal Basic Income (UBI) has gained attention in recent years as a potential solution to the challenges posed by automation. At its core, UBI is a government provision of a set amount of income for *every* individual, regardless of their background, employment status, or financial situation. This unconditional income would provide a financial foundation for everyone, untethered from any job or employer. It would not be contingent on

employment, savings, or income levels. It would be *universal*, meaning that everyone who is part of the system, from the most impoverished to the wealthiest, would receive the same amount of money.

Another key part of UBI is that it only provides a *basic* level of income. A person living off of only their UBI should be able to afford a safe and secure place to live, healthy food, and other basic necessities. If a person wants designer goods, exotic vacations, or other luxuries, then it is up to them to find a way to make additional money to get those things.

The biggest glaring problem with implementing UBI is figuring out how to fund it. In the US, individual income taxes are the primary source of government revenue. However, this situation would create a circular problem where taxes come from income, and income comes from taxes. Attempting to adjust the taxation system to cover UBI for both workers and non-workers would put an unsustainable burden on those still employed, who may feel understandably resentful. Moreover, as AI systems increasingly replace human jobs, the shrinking workforce would struggle to bear the growing weight of funding the entire system.

This issue of a shrinking base of income tax payers will be a pressing concern, irrespective of UBI. As jobs are taken over by AI, income tax revenue will shrink, while simultaneously the number of people needing assistance will grow. Continually increasing taxes on those who would still have jobs will not be sustainable, with or without UBI.

Leaving the majority of the population to starve on the street isn't a reasonable option, even for the cold-hearted. Not only would it be grossly unfair, inhumane, and reprehensible, it would also not be politically stable. If those people who end up directly benefiting from AI automation cavalierly tell those who don't to "go eat cake," then they probably should expect a violent backlash similar to the one that has become linked to that phrase.

If the alternative is infeasible, then we must find some way to support those without jobs. The current unemployment and welfare systems are inadequate, plagued by huge bureaucracies and irrational penalties for working. Even if these support systems were effective, they would still rely on a shrinking tax base, which is unsustainable and would only perpetuate the problem. Therefore, we need to explore alternative solutions that can support people without jobs and how to fund them from sustainable sources.

Challenges of Funding UBI

Some have already suggested a national sales tax as one possible solution to address the shrinking income tax base. While sales taxes can be regressive, meaning that people with less money tend to pay a larger fraction of their income relative to those with more money, a well-designed tax can mitigate this effect. For instance, exemptions for essential items like groceries, medical expenses, and housing could ensure that a sales tax is more progressive. Ideally, profitable corporations would bear the brunt of the tax, as they would be the direct beneficiaries of AI-driven efficiency. Some might expect that corporations would simply pass the cost on to customers, but the benefits of AI-driven productivity should offset those costs.

Building on this last point, we could make the taxation of AI-related savings explicit by designing a tax system that specifically targets the benefits of automation. By mildly taxing AI-related cost savings, corporations would still reap the benefits of automation while contributing to the broader social good. Importantly, as jobs are taken over

by AI systems, leaving more people in need of assistance, tax revenues would grow instead of shrink.

This core concept of requiring that the benefits of AI be shared with society is a fundamental goal. Under the current system, people work to produce value, and their salary is essentially a slice of that value minus the costs of production. The employer also takes a slice, as does the government. This leads to a zero-sum game where there is a limit to the total value created by a pool of workers, and trying to share the value broadly with those who don't have jobs invariably leads to shrinking the existing slices. With AI work, the costs of production are less, and there is no employee share because AI is not an employee and does not need a salary.

Whatever the details, one thing seems inescapable and essential: **If AI systems are going to take over human work and therefore human salary earning, then they must also replace human tax paying.**

Also note that while some work by AI systems will take jobs away from humans, the total amount of work done by AI is likely to be much greater than the number of human jobs displaced. One reason for this difference is that AI is capable of doing things that would be too difficult or dangerous, or conversely too trivial or boring for people to do. Another reason is that AI systems can be scaled up more easily than a human workforce. This greater base means that AI work could be taxed at a much lower rate than we currently tax human work.

Measuring AI Work

Imagine a world where the work done by AI systems can be measured and quantified. One possible approach is to use a concept called "Human Equivalent Effort Time" (HEET), which represents the amount of time a typical human would need to complete the same task. With HEET, we could track AI automation not only as it replaces human labor, but also as it creates new work opportunities for AI systems. A measurement like this would allow us to tax the work done by AI systems in a way that's fair, efficient, and sustainable.

By taxing AI work in HEET units, we can create a new revenue stream that supplements and eventually replaces traditional income taxes. This approach would enable continued revenue growth as AI systems become faster, more efficient, and easier to scale. Unlike human labor, AI systems are relatively inexpensive to maintain and can be replicated easily, allowing growth without all the difficulties typically involved when trying to hire more people.

A side point to note is that the AI systems discussed here are simply software designed to perform specific tasks, without any consciousness or self-awareness. Using a sentient AI for most work appears neither necessary nor desirable. An amusing illustration of this point is from the TV show Rick and Morty, where a brilliant scientist creates a sentient robot just to pass butter at the breakfast table. The self-aware robot experiences an existential crisis and sadness when it realizes its purpose is so trivial. While humorous, it nevertheless demonstrates potential pitfalls of creating sentient AI just for work tasks.

Like the butter-passing robot, humans experience boredom and depression when faced with tasks that are monotonous or of little value. Unnecessarily replicating that experience with our machines seems pointless and cruel. The vast majority of tasks can be accomplished without self-awareness, and there's no need to create complex artificial life to perform them. (Which is a good thing because we don't currently know how to make self-aware software.) Machines with no sentience do not need a slice of the value they create beyond the costs of

operating them.

I don't think measuring HEET would be clear and simple, but it doesn't seem impossible. We know a lot about human workers and what they are capable of, and that knowledge could be applied to assessing any task performed by an AI system and estimating how much time a human would take to do it or something similar. Making that determination for the millions of tasks that AI will be used for sounds tedious, but it's also something that could be automated by AI systems. One could imagine that instead of reporting how much a company paid in payroll and taxing that payroll, companies would instead report how much HEET was done by their software and pay a tax on that.

One could further imagine that an AI system managed by the IRS that would take descriptions of work done, assign an amount of HEET, and assess a tax. As with today's tax systems, this would require some honesty and also some penalties for lying. People would not enjoy the process, and there would likely be frequent appeals of the assessment. People would intentionally or by error fail to report work, and everyone would probably continue to hate the IRS. In other words, it would not be that different from today's taxation. We would be trading one frustrating bureaucratic tax system for another, with the key difference being that this new one would tax rapidly growing AI-driven productivity rather than dwindling personal income.

Some Details

One expected objection to taxing AI work is that it would be burdensome on companies and that it would stifle innovation and growth. However, that objection doesn't really stand up to scrutiny. Today, a large company wanting work done must hire someone, which requires paying the person and paying taxes. If AI systems are doing the work and being taxed, then there is no salary to pay, and the amount of taxation could be much smaller than what they would otherwise pay in payroll and payroll taxes. For example, instead of paying an employee \$50K, plus \$10K in payroll taxes (and hopefully some amount for benefits) to get one human unit of work done, a company might instead deploy ten human-equivalent AI systems, paying \$1K in taxes for each, thus getting 10x the work done for 1/6th the total cost while the government still collects the same amount in taxes.

The idea that everyone comes out ahead sounds nonsensical from our normal zero-sum perspective. How can everyone come out ahead? The critical change will be a vast pool of AI workers that do work yet need nearly nothing in return.

Individuals and small companies might look to AI automation to enable activities where labor is needed but they can't afford to pay for human staffing. For these small entities, paying a tax on the use of AI systems might make their business infeasible. In this respect, taxing HEET could stifle innovation and create barriers to new competition. However, the system could be designed to avoid this problem.

For example, both personal use and business use with less than \$1 million in gross annual revenue could be exempted from HEET tax. To avoid a sudden cost shock, the tax could ramp up between \$1 million and \$2 million gross annual revenue, allowing these businesses to adjust to the new tax gradually without being overwhelmed by the burden. These numbers are just examples, and one could imagine a carefully thought-out system with other appropriate exemptions.

The Future of Work and Human Wellbeing

The eventual future of work is going to very clearly be one of massive automation by AI systems, and those AI systems are going to be very productive. AI systems won't take sick days. An AI-run factory will not need to have space devoted to safe places for humans to stand, and it can run 24/7. An AI accounting service won't need to spend money on office space, HR, or training. If you have an enterprise staffed by AI and you want to expand, you just need to buy some more computers or maybe some more robots. The majority of the multitude of things that limit human productivity do not apply to AI systems. This bounty of savings is what will make something like UBI possible, but only if we find a way to share that bounty beyond the specific owners of the AI systems.

Today, being unable to find a job causes significant stress and insecurity, with unemployment being linked to a higher risk of depression and anxiety. Moreover, the financial and societal pressure to find a job can be overwhelming and debilitating. However, people without jobs who are financially secure often find that they're able to pursue their passions without the stress of needing to earn a living. For instance, a well-planned and well-funded retirement provides the financial freedom to devote oneself to rewarding activities that bring joy and fulfillment, without the limitation of needing to support oneself financially.

If we find a way to share the bounty produced by AI, we could create a world where no one needs to struggle with poverty or homelessness. People would have the freedom to pursue their passions and interests, and spend their days doing what brings them joy. They might devote themselves to creative pursuits, adventures, or community service, and have the time and resources to nurture meaningful relationships. In this world, work would be a choice, not a necessity, and people would not spend most of their waking hours working for someone else.

This isn't the first time someone has predicted "the end of work," but this time is fundamentally different. In the past, technological advances made human work more efficient. Instead of spending hundreds of person-hours to dig a hole, one person with a steam-shovel could dig the same hole in a couple of hours. Rather than painstakingly building something one at a time by hand, the same number of workers in a factory could churn out hundreds every hour. However, you still needed people. Now consider a factory with no people, or construction equipment that drives and operates itself. Production with zero people is only possible because AI systems can make operational decisions and handle exceptions on their own.

Final Thoughts

My suggestion of taxing the work done by AI systems is admittedly somewhat inchoate and vague. It probably has a hundred problems that would need answers and thousands of details that need figuring out. Also, there may be much better ideas that I've overlooked, both for funding and for distributing support. A national sales tax or a tax on HEET are just two possibilities. Maybe something as simple as taxing power spent on computing would suffice, or maybe we will need complex new mathematical theories relating to the value of a computation.

I also suggested some exceptions and thresholds that would be intended to keep the tax burden on the corporations that can afford it and that are benefiting from replacing human employees with AI systems. Determining what those exceptions and thresholds should be will require careful thought and analysis to make the numbers work. A badly designed system could be worse than doing nothing. This problem is a place for math and economics, not political squabbling and catering to special interests.

In conclusion, my proposal for taxing AI work is just one possible solution for sharing the fruits of AI automation. While still in its early stages, I believe this is an important conversation to have. As we move forward, it's crucial that we prioritize careful consideration and analysis to ensure that any solution we choose is fair, effective, and sustainable. Rather than getting bogged down in partisan politics, let's focus on finding a solution that benefits all people, not just a select few.

Regardless of what the actual solution ends up being, if we don't find a way to share the bounty produced by AI systems, then I fear that most people in the world are going to be left behind and suffer greatly. We have a choice in front of us: dystopian inequity or a bright world with ample resources for all. Unfortunately, the default if we do nothing is the dystopian one. I very much hope that we can put aside our differences and fears, and focus on realizing the brighter possibility.

About Me: James F. O'Brien is a Professor of Computer Science at the University of California, Berkeley. His research interests include computer graphics, computer animation, artificial intelligence, simulations of physical systems, human perception, rendering, image synthesis, machine learning, virtual reality, digital privacy, and the forensic analysis of images and video.

If you would like to read my other articles, consider subscribing. You can also find me on LinkedIn, Instagram, and at UC Berkeley.

Disclaimer: Any opinions expressed in this article are only those of the author as a private individual. Nothing in this article should be interpreted as a statement made in relation to the author's professional position with any institution.

This article and all embedded images are Copyright 2024 by the author. This article was written by a human, and both an LLM and other humans may have been used for proofreading and editorial suggestions.