First Quarter 2024

Economic Insights

Volume 9, Issue 1

Generative AI: A Turning Point for Labor's Share?

It's Not Just Who You Know, It's How You Know Them

Making Sense of Decentralized Finance

Question & Answer | Research Update | Data in Focus

FEDERAL RESERVE BANK PHILADELPHIA

Economic Insights

A publication of the Research Department of the Federal Reserve Bank of Philadelphia

Economic Insights features nontechnical articles on monetary policy, banking, and national, regional, and international economics, all written for a wide audience.

The Federal Reserve Bank of Philadelphia helps formulate and implement monetary policy, supervises banks and bank and savings and loan holding companies, and provides financial services to depository institutions and the federal government. It is one of 12 regional Reserve Banks that, together with the U.S. Federal Reserve Board of Governors, make up the Federal Reserve System. The Philadelphia Fed serves eastern and central Pennsylvania, southern New Jersey, and Delaware.

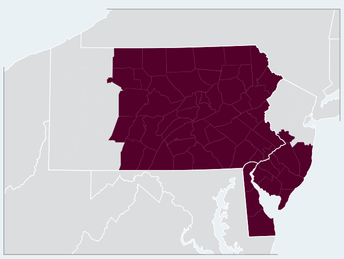

ISSN 0007–7011

Connect with Us

We welcome your comments at:
PHIL.EI.Comments@phil.frb.org

E-mail notifications:
www.philadelphiafed.org/notifications

Previous articles:
www.philadelphiafed.org/economicinsights

X (formerly known as Twitter):
@PhilFedResearch

Facebook:
www.facebook.com/philadelphiafed/

LinkedIn:
https://www.linkedin.com/company/philadelphiafed/

Contents

First Quarter 2024 Volume 9, Issue 1

Q&A...

with Joseph Abadi, an Economist here at the Philadelphia Fed.

Joseph Abadi

Despite having earned his doctorate in economics less than three years ago, Joseph Abadi has already produced an impressive body of work about the intersection of finance and technology. He also helped launch, in conjunction with the University of Pennsylvania, our ***Digital Currency Center***. For this issue, he examined decentralized finance, which uses blockchain technology to match buyers and sellers of assets. In this Q&A, he explains what led him to study economics, and what he hopes to accomplish through this work.

Where did you grow up?
Mostly in New York, but I also spent some of my childhood in Argentina. My dad was from Argentina, and at one point he had a business opportunity there. So, we moved there for a few years.

What did your parents do for a living?
My father was an investment banker. My mother was in finance too, but she stopped working for a few years when I was born.

Did their work shape your interest in finance, or did that develop organically later?
Because my parents worked in finance, it was a frequent topic of conversation around the dinner table. Also, because my dad was from Argentina, I was familiar with things like its currency crisis and inflation problems. That gave me an interest in economics and eventually in becoming an economist.

What led you to study math and physics in college?
I've always enjoyed thinking about and solving problems, and I really enjoyed studying math and physics in high school. But I maintained my interest in economics while studying them. I wanted my work to have an impact or at least be comprehensible to more than a handful of people. If you do a math PhD, there are probably only about four people in the world who will understand what you're studying. So, I got a PhD in economics.

After you earned your PhD, you could have pursued a career in finance, just like your parents. But instead, you went to work for the Philadelphia Fed. What led you to make that decision?
I'd had a couple of experiences in college interning at financial firms. It was interesting enough, but I wanted to focus more on research, which is what I enjoyed doing most. Writing my dissertation reinforced that impression. I really enjoyed doing my own research, getting to work on whatever topic I found interesting for as long as I needed to. Whereas, if you're working in industry, the questions are still interesting but you're often on a compressed schedule. Once you get far enough, it doesn't make sense to go further. You move on to the next thing because finance is about making money.

What lessons did you learn from physics and math that you apply to your research in economics?
Math and physics helped me think about and simplify abstract problems, and studying them gave me an intuition for how to solve a big problem before starting work on it. Just today I was working on a problem, trying to understand under what conditions a currency will circulate globally. For instance, why does the U.S. dollar circulate globally whereas the Argentine peso doesn't? A trick that I used to solve this problem is to think of extreme scenarios. For instance, what if one country has high inflation and the other doesn't? Then what will happen? What if one country is very small relative to international markets and one country is big? That's what I took from math and physics: making simplifications to get some idea of where to go with a problem.

What are your goals for the work you're doing here?
To write papers that are going to change the way people think about certain issues. For example, I cowrote a paper about the design of blockchains from an economic perspective.[1] Our goal was to help people understand the costs you incur when you design a blockchain recordkeeping system. We don't disparage blockchain recordkeeping systems. But if you were to read the paper, a natural thought is that these systems incur a lot of unnecessary costs. So, one change that might come from my research is making blockchain recordkeeping systems cheaper.

If you teach people about physics, you're not going to change the physical properties of the universe. But if you teach people about economics, it can lead to changes in the economy.
That's true. That's why economics is more difficult than physics. The economic laws of the universe aren't stable in the same way.

1 Joseph Abadi and Markus Brunnermeier, "Blockchain Economics," Federal Reserve Bank of Philadelphia Working paper 22-15 (2022), https://doi.org/10.21799/frbp.wp.2022.15.

Generative AI: A Turning Point for Labor's Share?

Thanks to artificial intelligence, labor's share of national income may no longer hold steady.

Lukasz A. Drozd
Economic Advisor and Economist
FEDERAL RESERVE BANK OF PHILADELPHIA

Marina M. Tavares
Economist
THE INTERNATIONAL MONETARY FUND

After years of slow and steady development, generative artificial intelligence (AI) technologies have exploded in popularity, and many experts believe that we are entering a new, AI-driven phase of the Industrial Revolution.[1] The advent of AI as the new engine of growth raises questions about the future of labor. Some have expressed concerns that, in the short run and the medium run, AI may lead to employment losses brought about by task automation and the skill obsolescence of the current labor force. But an additional risk of AI is in the long run: Unlike previous technologies, AI may undermine labor's share of national income, and technological innovation could, for the first time, permanently reduce the importance of labor in the economy, even if full employment is maintained.[2] The labor share—or the share of national income that labor receives as wages, salaries, and other compensation—has remained remarkably stable throughout much of the Industrial Revolution. Its permanent or even persistent decline would thus constitute a major break from past trends (Figure 1). This is troubling because, historically, automation, when combined with slower labor income growth relative to capital income, has been associated with rising income inequality and social unrest.[3]

In this article, we discuss how technology affects labor's share of income, why the labor share has been stable for so long, and why AI may threaten that stability.[4] The balanced nature of technological progress—which sometimes leads to the automation of tasks previously done by humans, but at other times improves the productivity of already-automated tasks—might have been the key to sustaining labor's stable share of income. If this is so, AI's unprecedented potential for job automation in the coming decades may disturb this long-standing equilibrium and cause the labor share to decline.[5] A declining labor share does not mean that labor income growth will stall, but to fuel labor income growth, AI will have to prove sufficiently productive in jobs it displaces us from.

AI's Pivotal Moment

Economic historians attribute much industrial growth to a handful of transformative innovations known as general purpose technologies (GPTs). A GPT has four key traits. First, it's identifiable as a generic technology or organizational system. Second, it's widely adopted across the economy. Third, it has multiple distinct applications. And fourth, it generates numerous spillovers that catalyze secondary innovations or even spawn new GPTs.[6] According to one analysis, as of the mid-2000s there had been at least 24 GPTs, ranging from language and the wheel to the steam engine and the computer.[7]

Each of these groundbreaking advances set the stage for numerous secondary innovations and drove economic growth for decades. For instance, the Industrial Revolution was kicked off by the steam engine, which revolutionized industries such as textiles, cotton agriculture, material science, and transportation. Then, in the last quarter of the 19th century, electrification and the internal combustion engine brought mass production, improved mobility, and telecommunication, leading to the Second Industrial Revolution. And beginning in the mid-20th century, the transistor gave rise to digital technologies, including computers and computer software, sparking the Third Industrial Revolution. According to many experts, AI, thanks to major advances in generative AI in the last decade, will soon join the pantheon of GPTs. Some experts even expect AI to become the leading technology of the Fourth Industrial Revolution, and for decades to come.[8]

See ***What Is Generative AI?***

The first three industrial revolutions featured very different technologies, but these revolutions' aggregate growth patterns look remarkably similar. This consistency led economist Nicholas Kaldor to formulate six stylized facts of industrial growth, known as Kaldor's facts, and it led to the development of a unifying (neoclassical) growth theory that makes no distinction between the different technological eras that drove growth.

According to one of Kaldor's facts, the incomes of capital and labor grew at similar rates even though most productivity gains could be attributed to the introduction or improvement of capital in production.[9] Economists recognized this feature as early as the 1940s, and it means that, in the long run, capital-productivity-augmenting technological progress equitably benefited both labor and the owners of capital despite ongoing automation.[10] In other words, labor somehow held its own against capital in the battle over which side would get the bigger piece of the economic pie, even though most economic growth was attributed to new or improved machinery (that is, to capital-productivity-augmenting technological progress).

If history is any guide, then, the emergence of a new GPT that drives growth need not affect growth patterns—in particular, labor's share of income.

So, what could make AI different from, say, the spinning jenny, which revolutionized the textile industry in the 19th century, or Ford's assembly line? The concerning aspect of AI, as we see it, is that it is a major GPT with the potential to broadly and persistently tilt the incoming flow of new capital-productivity-augmenting innovations toward those that automate tasks, rather than augment the productivity of capital in previously automated tasks. According to our theory, the balanced nature of technological progress in the economy as a whole—which sometimes leads to the automation of tasks previously done by humans but at other times improves the productivity of already-automated tasks—was the key to sustaining labor's stable share of income. If that's the case, AI may disturb this long-standing equilibrium and cause the labor share to decline over the coming decades.

See ***AI's Automation Potential***

Forces Affecting the Labor Share

To see how AI could affect the labor share, we must first understand how any new technology that makes machines more productive affects labor's share of income in an economy. Let's imagine the economy as a simplified system that produces a generic output—say, "a unit of gross domestic product (GDP)"—by performing a list of basic tasks.[11] The tasks are so basic, they can be executed by either machines or people. Each task needs to be completed only once, and no task can be skipped. (We discuss below how these simplifications affect the analysis.)

Because each production task can be executed by either capital (a machine) or labor (a worker), a new technology that increases capital productivity can influence a task in only one of two ways: If the task is already automated (that is, already done by a machine), it makes the machine that does that task more productive and hence cheaper to employ; but if the task is not yet automated (that is, done by a human) and the new technology makes it cost-effective to automate that task, the machine will displace a worker from that task. The impact on labor's share of income depends on which kind of technology we're talking about, and hence which kind of innovation is brought about by technological progress, which, as we shall see, turns out to be a crucial distinction for understanding how AI enters the picture.[12]

Productivity Effect of New Technologies

Consider the first case, where a technology boosts the productivity of capital in a task that is already automated. Let's assume that producing a unit of output costs a firm $100, and, due to competition, the unit of output also sells for $100. Now let's assume that a single task costs a firm $2 when performed by a machine, but, thanks to the new technology, the firm uses a next-generation machine to accomplish the same task for less, say $1.

Initially, the firm reduces its costs by $1 per unit of output, which immediately boosts the firm's profits. Since profits contribute to capital income but not wages, labor's income remains unchanged. Capital owners lose $1 in income payments for the machine, but they gain $1 in profits, and so there is no change in capital income either. Capital owners can pocket the gain because nothing else has changed, and we presume that firms pay workers just enough to keep them employed.

But if that technology becomes broadly available, other firms will adopt it as well, and all firms will seek to expand production because they can make (excess) profits. Competition between firms to expand their market share will erode profits, driving down the price of goods to a new break-even point of $99, and capital income payments will decrease by the saved dollar on the new machine. As that happens, labor can reap *all* the benefits from the new technology, for two reasons.

First, as the price of goods in the economy falls, both workers and the owners of capital pay less for the goods they purchase, leaving them more money to purchase more goods, which is to say that their real income—income adjusted for purchasing power—rises. Can labor's *share* increase because of the falling prices of goods implied by declining profit margins? Yes, and how much it increases depends on labor's (and capital's) initial share of income.

For example, if labor's initial share is two-thirds, as historically has been roughly the case, the increased purchasing power gives labor a 1 percentage point gain on two-thirds of each dollar generated via production and captured as labor income, which amounts to a total *real* income gain of two-thirds of a dollar in constant purchasing power (that is, in terms of goods one can purchase). These gains are earned on each unit of output produced. And labor's share of income is equal to the number of units of output it can purchase using the income it earned on each unit of output it helped produce. Thus, labor's increased purchasing power implies that labor's share of income has increased by two-thirds of a percentage point.[13] Simply put, to purchase goods, labor must effectively pay for its own input into production and for capital (including profits). Since capital costs less, labor can purchase more goods for the income it earns, and so its share rises at the expense of capital.

Remarkably, the benefits that accrue to labor don't end there. Since the new technology leads to a broad-based decline in prices, machines, which are also produced goods, should cost less. If the price of a machine, like the price of the good it helps produce, also falls by 1 percent, the additional cost savings associated with this decline will further drive down the price of goods, continuing the cycle of declining prices of goods and machines. This adjustment process won't stop until labor captures *all* the benefits from the new technology, with labor's share rising by as much as 1 percentage point and capital's share falling by a full percentage point.

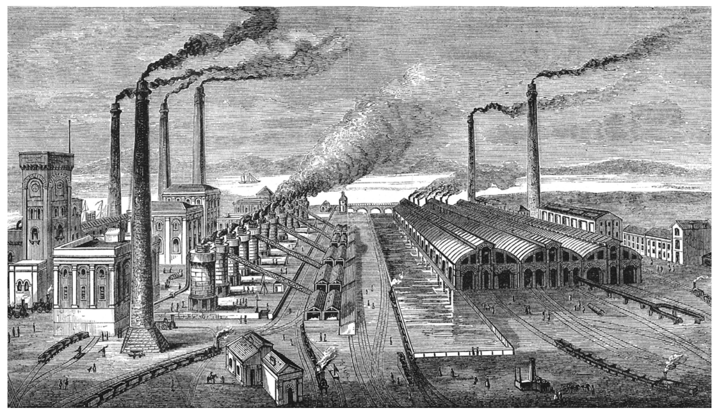

This *productivity effect of capital deepening* increases labor's share of income; it applies when innovations make machines more productive or cheaper to employ in tasks that are already automated. Although this is a simplified example, the productivity effect arises even when we allow for limited substitutability between tasks. Of course, if one task can effectively and perfectly substitute for some other tasks, capital will displace labor from these other tasks via task substitution, and this will dampen the productivity effect. However, broad substitutability across tasks is implausible. Technological processes for producing goods and services are fairly rigid and comprise a precisely defined sequence of myriad distinct tasks. For example, just because an automotive plant can replace a human welder with a robot doesn't mean it can do away with the workers who install windshields.[14]

Displacement Effect of New Technologies

But consider a very different scenario: For a given task, a new technology makes a machine only marginally more productive than labor. Let's assume that labor earns $2 per unit of output for this task, and the new machine performs the same task at nearly the same cost—say, $1.99. That is, the arrival of a new technology (and thus a new machine) leads to the automation of the task in question, but the difference of 1 cent between the cost of labor and capital implies that it does *not* save the adopt-

FIGURE 1

Labor's Share of Income Held Steady

This was true even during an era of widespread automation.

Labor's share of national income for the UK, the U.S., and France, 1855–2017

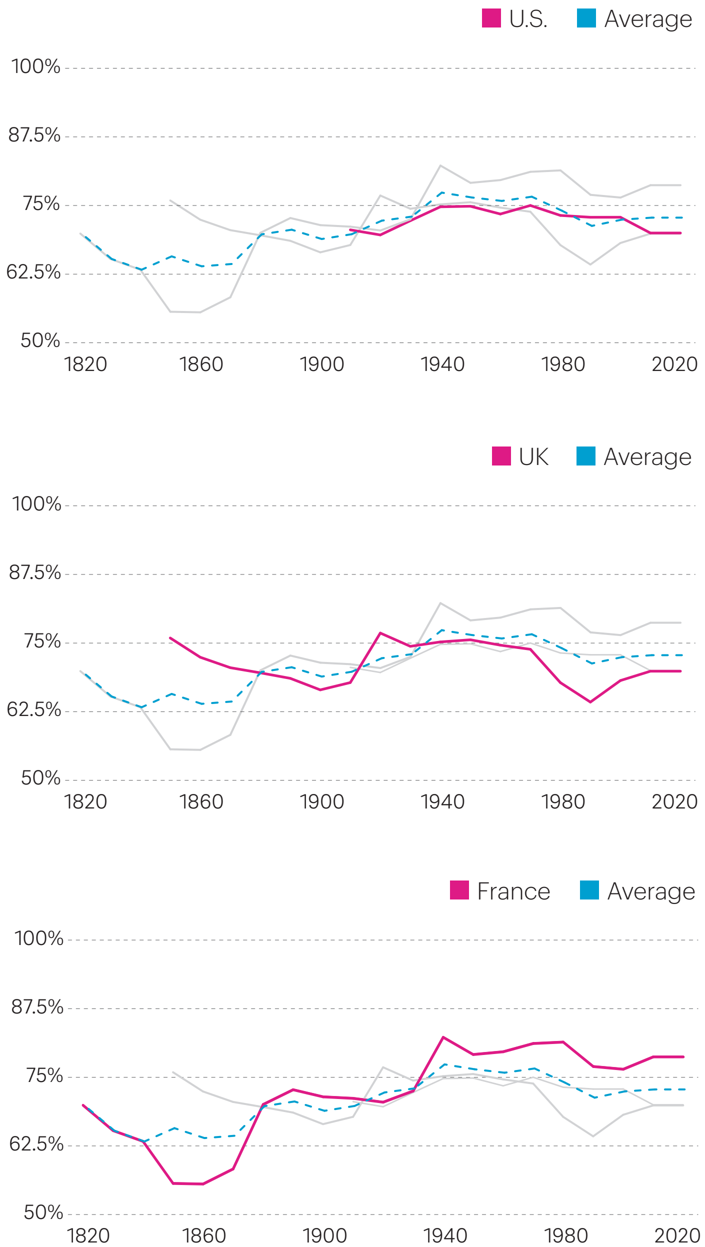

Data Sources: For all countries, 1855–2020: World Inequality Database (https://wid.world/wid-world) and Distributional National Account (Table S.A1), as detailed in Piketty, Saez, and Zucman (2018). For France, 1820–1850: Piketty (2014).

Note: Labor's share of income is defined as the ratio of compensation of employees and 70 percent of mixed income over the factor price national income. The data for this figure come from raw tax data. Labor's share of national income for all countries except the U.S., which is labor's share of personal income.

ing firm much money.

Since profits increase only marginally after automation, the price of output remains $100, and for this reason the final effect could not be more different. Payments to capital increase by almost $2 per unit of output ($1.99 to be precise); labor, by being automated out of the task, loses $2 in displaced income. Consequently, labor's share of income drops by nearly 2 percentage points, and capital's share rises by the same amount.

This *displacement effect of capital-productivity-augmenting technologies* counters the productivity effect, reducing the labor share. More-realistic situations feature both effects, but we can still decompose their net effect into the two elementary effects. Say, for example, the cost of the machine falls to $1 rather than $1.99. We can equivalently think of all such cases in two steps: The cost of using a machine first falls to $1.99, triggering the *displacement effect* associated with automation, and it then drops further to $1, triggering the *productivity effect* within an already automated task. The combined effect reduces labor's share of income by about 1 percentage point: The displacement effect lowers it by about 2 percentage points but the productivity effect then raises it by 1 percentage point.[15]

The remarkable property of such mixed cases is that labor's share of income nonetheless *always* falls—although to a varying degree depending on the strength of the offsetting productivity effect. Consequently, technologies that trigger task automation *always* displace labor's share of income, and hence for the labor share to remain constant it must be offset by innovations that raise the productivity of machines used in already-automated tasks. To see this, note that even in the most extreme case where the cost of using the new machine drops to zero instead of $1.99—that is, capital becomes infinitely more productive in our task—labor's share of income still does not increase. Although it falls by 2 percentage points due to the displacement effect, it then rises by 2 percentage points, for a roughly nil net effect.

What the Past Tells Us About the Future

Given these contrasting effects, it's clear that labor's share of income can go up or down after the arrival of innovations that lead to the emergence of new, capital-productivity-enhancing technologies in the economy. This makes it so much more surprising that the labor share lacked any distinct trend over the last century. Our research attributes this remarkable phenomenon to the random and cumulative nature of innovations across a vast number of tasks in the economy, which, if true, has consequences for how we think about the effects of AI.[16]

Innovation may spring from directed research and development (R&D) in a specific area. But at the macroeconomic level the vast number of tasks in the economy, the incentive to invest in R&D in various areas, the random outcome of research, structural transformation that reshuffles the importance of individual tasks, and the varying impact of new innovations on individual tasks implies that capital productivity within individual tasks may well be a random process that exhibits fairly stable statistical properties across both tasks and time. As a result, sometimes an innovation enhances capital productivity in an already automated task by some random percentage of its value, and at other times it enhances productivity in a task yet to be automated, with the odds of R&D producing either type of innovation, the average size of (relative) productivity increments, and their variability all being stable. If this is so, then, according to our research, labor's share of income will stabilize despite ongoing automation because, as we note above, automation, which reduces labor's share of income, will eventually be balanced out by improvements in already-automated tasks, which raise labor's share of income.

To better understand this phenomenon, imagine it's the dawn of the Industrial Revolution: All tasks are performed by labor, as machines are almost nonexistent; all innovations affect tasks that are not yet automated, simply because very few tasks are done by machines. Initially, the labor share declines, but as more tasks become automated, subsequent innovations more and more frequently improve the productivity of already-automated tasks rather than automating tasks operated by labor. As a result, the odds of either innovation stabilizes, and our research shows that the economy enters a phase of "scale invariant" growth: Even though innovations increase the productivity of the whole economy, the share of tasks that need a lot of capital relative to labor, or need little capital relative to labor, stabilizes. As this happens, the labor share of income also stabilizes.

This mechanism's relevance to the economywide effects of R&D is not that surprising–it already explains many seemingly unrelated phenomena in the natural sciences. For example, an analogous mechanism is believed to govern population growth across cities, employment growth across firms, wealth across people, the number of species across genera, the number of citations of academic papers, and even the propagation of cracks in fracturing materials. These phenomena are unrelated, but their statistical properties are strikingly similar: They all exhibit the so-called "power law property," which in all these cases is similarly attributed to scale invariant growth or propagation.[17]

Why Labor's Share of Income May Decline

But if R&D leads to a new class of innovations that affect nonautomated tasks significantly more than automated tasks relative to the historic pattern, then the natural mechanism that stabilized labor's share will have changed, and labor's share of income may decline. Our research suggests that this is one key concern about AI's long-run effect.

Until recently, one of the most significant barriers to automation has been the need for cognition. This barrier prevented machines from operating autonomously or made it costly to automate many production processes. Cognitive abilities are a significant barrier because we humans are largely incapable of supplying algorithms that would enable machines to mimic our own cognitive abilities.[18] As an example, until the advent of AI, there was no algorithm for a cognitive ability as basic as the general computer-based visual recognition of objects. It's not difficult to see how this could keep machines from replacing humans.

If AI successfully and universally overcomes this hurdle, it will unleash a torrent of long-delayed automation-focused innovations, and for some time new innovations may be heavily biased toward those that enhance the productivity of machines in tasks that have not yet been automated as opposed to those that are already automated. As that happens, labor's share of income may permanently decline despite having been constant in the past. The size of the decline in labor's share will depend on the exact nature as well as the productivity of AI-driven innovations, and it will happen irrespective of whether displaced labor finds employment.

Depending on how productive AI is, however, labor's *income* may still grow even if its *share of income* declines. As an example, consider the case where the cost of employing a machine drops to zero in tasks previously operated by labor. In that scenario, labor's share of income stays constant because the displacement effect is offset by the productivity effect, as we have discussed. But because the economy has become more productive (that is, it can produce more output from the available resources), labor's overall income must have risen. (Here it is important that all resources, including labor, remain fully employed, or else labor income will not rise.)

What bodes well for growth in labor's income is what also makes AI so good at automating jobs. AI software is a nonrival technology, meaning that its use by one entity does not preclude another entity from using it at the same time. Thanks to this nonrival nature, big economies of scale will likely make AI highly productive in many areas. For example, even though it costs billions of dollars to develop self-driving cars, the necessary software can be deployed to every car once it becomes available, and its cost will likely decline as software providers compete for customers.

But even if AI proves productive enough to fuel labor income growth, the declining labor share is a concern. A declining share means that the total purchasing power of those who supply labor grows weaker than the total purchasing power of owners of capital, contributing to inequality between these groups. In many areas where the two groups compete for a fixed supply of resources, the relative purchasing power is all that matters, and a change in relative purchasing power can trigger social unrest due to the lower welfare experienced by one group. For example, land (in good locations) is largely in fixed supply, since it cannot be manufactured. As a result, the price of land depends on the average "bidding power" of those who seek to purchase it, and even if labor income continues to grow, housing may become unaffordable for those who supply labor. It is perhaps no coincidence that labor unrest coincided with a persistent but ultimately transient decline in labor's share of income in the 19th century.[19]

Conclusion

In the early 1980s, Nobel laureate Wassily Leontief famously compared the ongoing microcomputer and software revolution to the advent of steam, electric, and later internal combustion power in production and transportation. He questioned whether humans could avoid being replaced by computers much as horses had been displaced by mechanical power. This forecast didn't come to pass during his lifetime, but the question he posed has gained renewed relevance in the face of AI's rapid development.

AI's potential seems limitless, and perhaps one day AI will take all our jobs, in which case Leontief's forecast will come to pass. But we think that's unlikely in the foreseeable future and perhaps even impossible. We humans possess a wide range of skills and can perform many different tasks. If we maintain a cost advantage over machines in a significant fraction of these tasks, we should avoid the fate experienced by horses, at the very least for quite some time.

In summary, although some jobs may be lost in the short and medium run, we agree with many commentators who argue that AI, in the long run, will transform rather than displace jobs. But this conclusion alone misses an important point. If our theory is correct, the more pressing issue is AI's likely effect on labor's share of income. We need to carefully monitor and address the growing inequality as those jobs transformed by AI receive a smaller share of the economic pie.

What Is Generative AI?

Generative AI and AI in general refer to the set of machine learning methods that use artificial neural networks (ANNs) to obtain computer algorithms for tasks normally requiring human intelligence or computer software written by a human programmer (for example, the task of facial recognition). In other words, these tools write computer software on their own.

An ANN is a mathematical structure inspired by the human brain and emulated by a computer.[20] Like the human brain, an ANN consists of an interconnected network of nodes, called artificial neurons, each of which fires a new signal based on the sum of the signals it receives from the neurons it is connected to. In an ANN, these signals are numerical; in a brain, they're electrical.[21]

Take, for example, an ANN designed to distinguish between images of dogs and cats (Figure 2). This ANN's artificial neurons take grayscale pixel values as inputs. Subsequent neurons process these signals to arrive at two classification signals returned by the last two neurons, one for "dog" and one for "cat." The more confidence this network has in its classification, the greater the signal returned by the output neuron associated with that classification (as shown in the figure). Parameters of this network, which are mostly weights that each neuron assigns to each connection, encode the underlying algorithm for classifying images, and these parameters are obtained algorithmically by computer software in the process of training.

Training an ANN involves feeding it thousands of examples (labeled images in the depicted case) and algorithmically adjusting the network's internal parameters to minimize the average prediction error on that training data set. The key principle behind this algorithm is to iteratively adjust all the parameters of the network by small increments in the direction of the steepest descent in terms of the average prediction error on the training set. Just as water flowing downhill finds the lowest point of a surface, this algorithm can obtain an error-minimizing combination of the network's parameters.

Advanced AI architectures—such as the transformer architecture powering the latest AI models—involve much more complex structures and may involve a more complex training procedure, but the basic principle of operation and training is similar. What makes ANNs so successful is still a bit of a mystery, but it almost certainly has to do with the fact that this structure is conducive to identifying the patterns in the data that do well on both the training data and the new data (for example, learning to recognize a cat by its whiskers, long flexible tail, and pricky ears, so that the ANN can classify new images as well as those in the training set).

Generative AI refers to an approach that trains ANNs to create content. For instance, generative AI could learn from images of cats how to animate an image of a cat, or how to fill in the missing details of an image. A generative AI model of this sort acquires its "knowledge" of what a cat is by being trained to fill in and predict missing pixels based on neighboring pixels, across thousands of blanked-out images that contain images of cats. By doing this, the ANN powering this model, instead of merely learning how to recognize a cat, can learn how to recreate its image when prompted to do so.

But the most revolutionary leap in generative AI methods came from natural language processing, specifically through large language models (LLMs) like OpenAI's ChatGPT and its latest product, GPT-4. Generative LLMs are trained analogously, but by using text instead of pixels, with each word represented by a unique array of numbers. The ANNs powering these models have proven useful for learning about the meaning of words, language structure, and ultimately the human world at large. They do this by being trained to predict blanked-out words based on the surrounding text. LLMs are additionally trained to use their acquired "knowledge" to generate complete responses to human queries. This is done by feeding the LLM ratings for generated responses and teaching the model, which is already pretrained on text, how to sequentially add to the text those words that will maximize the rating for the sequentially generated answer.[22]

What makes the most advanced generation of large-scale LLMs revolutionary is their apparent ability to "reason" based on concepts. By doing so, these models showcase glimmers of an early form of artificial general intelligence, or AGI—the capacity to understand, learn, and execute a broad range of intellectual tasks, much as humans do.[23] This was once considered a distant goal, and the use of AI has been thus far limited to narrowly defined tasks, such as image classification. But more people today believe that generative AI will eventually pave the way for an advanced form of AGI. It's unlikely that AGI will soon match or exceed human intelligence, but even a limited AGI-like capability can transform the economy.

FIGURE 2

An ANN Decides: Dog or Cat?

The connections between neurons show how signals travel through the network, and hence how information is processed by it.

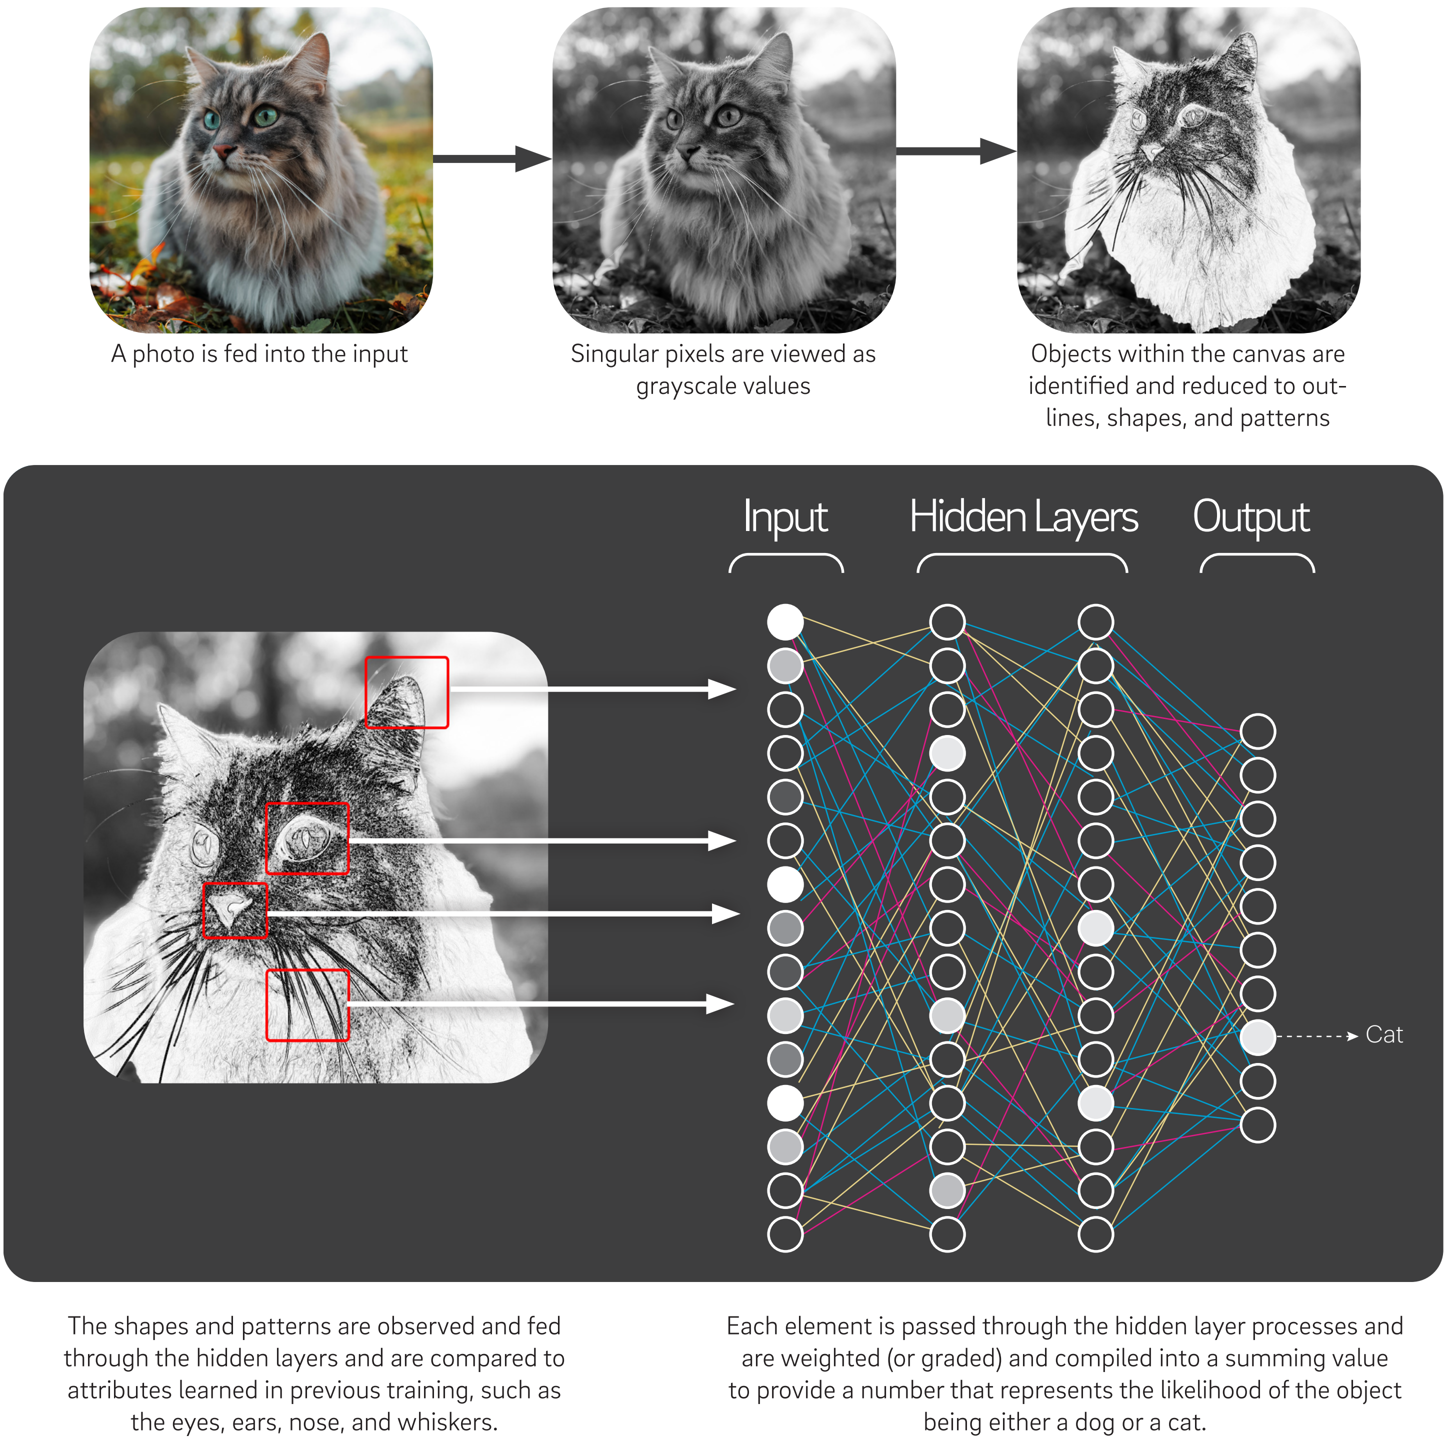

AI's Automation Potential

Several studies confirm AI's immense automation potential. Given the important role of humanlike cognition in production—a role that AI research may offer a substitute for—it's not surprising that AI may have such an effect.

According to one recent study, up to 19 percent of jobs could see at least half of their tasks automated at or near the current state of AI technology over the coming decades.[24] Overall, about a third of all tasks could be automated. According to another study, the explosive growth of generative AI tools may lead to the automation of up to 30 percent of labor hours by 2030. Generative AI increases these estimates by at least 8 percentage points, again at or near the current state of technology.[25] Higher-income jobs are more susceptible to automation—a departure from the past, when lower-income jobs were the primary target of automation.[26] Perhaps this will lessen AI's impact on income inequality among workers, but the scale of these changes may exacerbate inequality between workers and the owners of capital.

Business people generally agree with these assessments. For example, 75 percent of companies recently surveyed by the World Economic Forum plan to use AI to automate jobs within the next five years, and they expect a reduction of approximately 20 percent of existing jobs, many of which are skilled jobs.[27] Casual evidence suggests that AI is also becoming a strategic consideration in hiring decisions across the economy. For example, IBM CEO Arvind Krishna announced that IBM may freeze hiring in all positions that are likely to be affected by generative AI. According to Krishna, if IBM follows through on the freeze, attrition will reduce IBM's workforce by 30 percent in the next five years. This represents an estimated loss of about 7,800 jobs.[28]

These analyses do not consider whether it makes economic sense for firms to use AI, and firms may be underestimating the costs of employing AI. But the massive surge of investment in AI and the fact that AI software is a nonrival good—meaning that nothing prevents two companies from using AI software simultaneously—does not bode well for labor. Nonrival technologies can bring massive economies of scale, so even costly hurdles may eventually prove economical to overcome. And because the cost of the marginal provision of AI software to another customer is close to zero, it will be hard for labor to compete once these hurdles are overcome. Yes, there is the cost of the hardware powering AI—and for the most advanced applications, that cost can be substantial—but the cost of cloud computing has been declining for decades, and that trend is likely to continue.

Workers, in response to a major wave of AI-driven automation, will eventually acquire new skills and find new jobs. Perhaps full employment will even be maintained in the short term and the medium term. But this does not mean that AI won't have a lasting impact on labor. Labor's share of income may still shrink. And that's a cause for concern separate from the employment losses or the labor force's skill obsolescence that AI may bring about.

Notes

1 See, for example, "Pause Giant AI Experiments: An Open Letter" (2023); the press interview with AI expert Geoffrey E. Hinton, widely regarded as the "godfather of AI," published by the *New York Times* on May 1, 2023; and hearings by the Senate Judiciary Subcommittee on Privacy, Technology, and the Law, held to discuss the safety and ramifications of AI.

2 Some income, such as income earned by proprietors, is challenging to classify. For this and other reasons, a precise measurement of the labor share is not possible and all existing measures of this share are approximate. For an overview of the challenges of measuring the labor share of income, see Gomme and Rupert (2004).

3 For a detailed account of the apparently transient labor share decline in the 19th century, see Acemoglu and Johnson (2023).

4 We draw on our recent research with Cornell University Assistant Professor of Economics Mathieu Taschereau-Dumouchel. See Drozd, Taschereau-Dumouchel, and Tavares (2022).

5 Some of the economic literature finds that labor's share has already begun to decline due to the advent of modern computer software and modern industrial robots. For a detailed discussion of the decline in labor's share of income over the last few decades due to software, see Aum and Shin (2020). See also Acemoglu and Restrepo (2020), who analyze the effect of industrial robots on local labor markets. Other hypotheses include the role of market power and concentration, the rise of superstar firms, the effects of trade and outsourcing, and mismeasurement.

6 This definition courtesy of Lipsey, Carlaw, and Bekar (2005).

7 Lipsey, Carlaw, and Bekar (2005).

8 As an example, see the discussion of AI in Brynjolfsson and McAfee (2016) and references in footnote 1.

9 This is supported by the decline in the relative price of capital goods in relation to consumption goods over at least a century of data. See Greenwood, Hercowitz, and Krusell (1997).

10 The stability of labor's share in the late 1930s and 1940s has been recognized as a statistical fact in the writings of John Maynard Keynes, Paul Samuelson, Robert Solow, and Michal Kalecki. It was known as Bowley's Law. See Krämer (2011) for a more detailed discussion.

11 This discussion is grounded in a simplified version of our more general task theory of labor share stability and growth involving automation. See footnote 4.

12 This decomposition was first proposed by Acemoglu and Restrepo (2018). Aghion, Jones, and Jones (2017) were the first to highlight the role of complementarity in the context of automation. They proposed a different naming convention because Acemoglu and Restrepo's productivity effect is essentially the classic Baumol disease effect.

13 Say labor is initially paid $66 and capital is paid $34 on each unit of output sold for $100 in the market. Note that labor can then purchase 66/(66+34) units of output for the income it earns on each unit of output it helps produce ($66). Its share is thus 0.66, or 66 percent. Accordingly, if capital is subsequently paid $33, and labor income is unchanged, labor can purchase 66/(66+33) = 66/99 units, which is its new share. Therefore, labor's share of income increases by approximately 0.66 x 1 percent, since 66/99 = 66/100 x 100/99 ≈ 0.66 x 1.01, and so the share changes by about (0.66 x 1.01) - 0.66 = 0.66 x 1 percent.

14 Another aspect that limits substitutability across tasks is that we consumers want a variety of goods and services.

15 This is an approximation—a number that declines by x percent and subsequently rises by x percent does not return to the previous level, unless x is very small, which is implicitly assumed throughout.

16 For an alternative hypothesis for labor share's stability, see Acemoglu and Restrepo (2018). They stress the emergence of new goods and services and the fact that the new tasks associated with them may be difficult to automate at first. In the process of churning tasks in and out of existence, the labor share can be sustained despite ongoing automation.

17 For example, the resulting frequency of, say, cities with at least N inhabitants, or wealth of at least N dollars, is inversely proportional to N. In other words, cities that are double the size of Philadelphia (or larger) occur in the data half as often as cities that are as large as Philadelphia (or larger). In the context of cities, scale invariant growth means that cities will grow in population randomly and independently of their current population. See Gabaix (2009) for a more detailed discussion of these and other examples.

18 An algorithm is a set of instructions for a machine to operate autonomously. Before AI, these instructions were written exclusively by humans. In essence, AI is a set of tools that writes these instructions autonomously. See the ***What Is Generative AI?*** sidebar in this article for more details.

19 For a detailed account of this turbulent period for labor, see Acemoglu and Johnson (2023).

20 Some of the earliest and most significant research into the intersection of neuroscience and neural networks, particularly regarding the implementation of logical gates, was conducted by Warren McCulloch and Walter Pitts in 1943. They considered a simplified neural network model that used basic logic to describe neural activity. Their seminal research inspired further exploration of neural network structure and ultimately led to the development of the first general neural network model, the Perceptron, by Frank Rosenblatt in 1958. This laid the groundwork for many modern neural network concepts and algorithms. An important milestone for generalizing the Perceptron model to multilayer ANNs with neurons that can modulate signals was the backpropagation algorithm, introduced to the AI community by Rumelhart, Hinton, and Williams in 1986.

21 The known similarities likely end there, although some new research suggests that there might be a tighter connection than previously thought. See Richards (2018) and Richards et al. (2019).

22 For a more technical and detailed description of generative AI technology, see Wolfram (2023).

23 Bubeck et al. (2023) analyze AGI capabilities of generative LLMs and provide many illustrative examples. These examples showcase the ability of AI to "reason" based on the acquired understanding of concepts underlying our human world. They refer to these capabilities as "sparks of intelligence."

24 Eloundou et al. (2023).

25 Ellingrud et al. (2023).

26 See Pizzinelli et al. (2023) for a discussion of how these features may affect automation across countries depending on each country's level of development.

27 Baschuk (2023).

28 Baschuk (2023).

References

"Pause Giant AI Experiments: An Open Letter," Future of Life Institute, March 22, 2023, https://futureoflife.org/open-letter/pause-giant-ai-experiments/.

"Policymaking in the Pause: What Can Policymakers Do Now to Combat Risks from Advanced AI Systems?" Future of Life Institute, April 19, 2023, https://blog.biocomm.ai/2023/04/19/the-future-of-life-institute-policy-making-in-the-pause-what-can-policymakers-do-now-to-combat-risks-from-advanced-ai-systems-19th-april-2023/.

Acemoglu, Daron, and Pascual Restrepo. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108:6 (2018), pp. 1488–1542, https://doi.org/10.1257/aer.20160696.

Acemoglu, Daron, and Pascual Restrepo. "Robots and Jobs: Evidence from U.S. Labor Markets," *Journal of Political Economy*, 128:6 (2020), pp. 2188–2244, https://doi.org/10.1086/705716.

Acemoglu, Daron, and Simon Johnson. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs, 2023.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. "Artificial Intelligence and Economic Growth," National Bureau of Economic Research Working Paper 23928 (2017), http://www.nber.org/papers/w23928.

Aum, Sangmin, and Yongseok Shin. "Why Is the Labor Share Declining?" Federal Reserve Bank of St. Louis *Review*, 102:4 (2020), https://dx.doi.org/10.20955/r.102.413-28.

Baschuk, Bryce. “Tech, AI Driving Job Changes for Nearly a Quarter of All Workers,” Bloomberg, April 30, 2023.

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Company, 2016.

Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, et al. “Sparks of Artificial General Intelligence: Early Experiments with GPT-4,” Microsoft Research 2023, https://doi.org/10.48550/arXiv.2303.12712.

Drozd, Lukasz A., Mathieu Taschereau-Dumouchel, and Marina M. Tavares. “Understanding Growth Through Automation: The Neoclassical Perspective,” Federal Reserve Bank of Philadelphia Working Paper 22-25 (2022), https://doi.org/10.21799/frbp.wp.2022.25.

Ellingrud, Kweilin, Saurabh Sanghvi, Gurneet Singh Dandona, et al. “Generative AI and the Future of Work in America,” McKinsey Global Institute Report, July 26, 2023, https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” manuscript, March 21, 2023.

Gabaix, Xavier. “Power Laws in Economics and Finance,” *Journal of Economic Perspectives*, 30:1 (2009), pp. 255–294, https://doi.org/10.1146/annurev.economics.050708.142940.

Gomme, Paul, and Peter Rupert. “Measuring Labor’s Share of Income,” Federal Reserve Bank of Cleveland Policy Discussion Paper 7 (2004), https://dx.doi.org/10.2139/ssrn.1024847.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. “Long-Run Implications of Investment-Specific Technological Change,” *American Economic Review*, 87:3 (1997), pp. 342–362, https://www.jstor.org/stable/2951349.

Keynes, John M. “Economic Possibilities for Our Grandchildren (1930),” in *Essays in Persuasion*, New York: W.W. Norton & Co., 1963, pp. 358–373.

Krämer, Hagen M. “Bowley’s Law: The Diffusion of an Empirical Supposition into Economic Theory,” *Cahiers d’économie Politique / Papers in Political Economy*, 61:2 (2011), pp. 19–49, https://doi.org/10.3917/cep.061.0019.

Lipsey, Richard G., Kenneth I. Carlaw, and Clifford T. Bekar. *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*. Oxford, UK: Oxford University Press, 2005.

Metz, Cade. “‘The Godfather of A.I.’ Leaves Google and Warns of Danger Ahead,” *New York Times*, May 1, 2023, https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html.

McCulloch, Warren S., and Walter Pitts. “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *Bulletin of Mathematical Biophysics*, 5 (1943), pp. 115–133.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. *Capital in the Twenty-First Century*. Cambridge, MA: Belknap Press of Harvard University Press, 2014.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. “Distributional National Accounts: Methods and Estimates for the United States,” *Quarterly Journal of Economics*, 133:2 (2018), pp. 553–609, https://doi.org/10.1093/qje/qjx043.

Pizzinelli, Carlo, Augustus Panton, Marina M. Tavares, et al. “Labor Market Exposure to AI: Cross-country Differences and Distributional Implications,” IMF Working Paper 2023/216 (2023), https://www.imf.org/en/Publications/WP/Issues/2023/10/04/Labor-Market-Exposure-to-AI-Cross-country-Differences-and-Distributional-Implications-539656.

Richards, Blake A. Speech given at the RAAIS Conference, London, June 29, 2018, https://raais.co/speakers-2018.

Richards, Blake A., Timothy P. Lillicrap, Philippe Beaudoin, et al. “A Deep Learning Framework for Neuroscience,” *Nature Neuroscience*, 22 (2019), pp. 1761–1770, https://doi.org/10.1038/s41593-019-0520-2.

Rosenblatt, Frank. “The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain,” *Psychological Review*, 65:6 (1958), pp. 386–408, https://doi.org/10.1037/h0042519.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. “Learning Representations by Back-Propagating Errors,” *Nature*, 323:6088 (1986), pp. 533–536, https://doi.org/10.1038/323533a0.

Wolfram, Stephen. “What Is ChatGPT Doing and Why Does It Work?” Wolfram Research, LLC, March 9, 2023.

World Economic Forum. *Future of Jobs Report 2023*. Geneva, Switzerland, May 2023, https://www.weforum.org/publications/the-future-of-jobs-report-2023/.

It's Not Just Who You Know, It's How You Know Them

What role do referrals play in matching workers with firms? And what are the consequences?

Benjamin Lester
Senior Economic Advisor and Economist
FEDERAL RESERVE BANK OF PHILADELPHIA

The views expressed in this article are not necessarily those of the Federal Reserve.

Since (at least) the seminal work of Nobel laureates Peter Diamond, Dale Mortensen, and Christoper Pissarides, labor economists have been trying to understand how workers and firms find each other and form a "match."[1] After all, this process determines which worker gets hired by which firm, how productive they are, how long the match lasts, and how much the worker gets paid. According to surveys of workers and firms, a *referral* is used somewhere in the hiring process in about half of all jobs.[2] Of course, a referral can take many forms, including a phone call to an employee at the hiring firm, a casual conversation with someone in the human resources department, or a formal letter of reference. Depending on the context, the person making the referral could be providing a variety of services: They could simply be connecting the worker and the firm; they could be sharing information about the candidate with the firm; they could be sharing information about the job with the worker; or they could be vouching for the worker should that worker get the job.

So, what do referrals do? And how does a worker's labor market outcome change when using a referral? Importantly, the answers to these questions shed light on one important source of economic inequality–namely, inequality that arises from differences in labor income (that is, wages). In particular, some economists argue that referrals *exacerbate* wage inequality, because well-connected workers help each other find and retain high-paying jobs while less-connected workers struggle.[3] However, others argue that referrals are an important channel for low-income, low-skilled workers seeking a job–a "last resort" of sorts–and thus an important force for *ameliorating* economic inequality.[4]

In a recent paper, David Rivers of Western University, Giorgio Topa of the Federal Reserve Bank of New York, and I offer a new perspective on the role of job referrals in the hiring process and their effect on wage inequality. One of our key insights is that the effects of referrals become much clearer if you make two key distinctions in the data. First, it's important to distinguish between different types of referrals: those from someone in the worker's social network (a friend or relative), and those from a contact in the worker's business or professional network. Second, it's important to distinguish between different types of jobs, as measured by the skill requirements to perform the job.

After making these distinctions, we find that business referrals are used most frequently by highly productive workers to find high-paying, high-skill jobs. As a result, these types of referrals tend to increase inequality. Referrals from family and friends, alternatively, are used most frequently by workers in low-skill jobs who struggle to find work through other, more traditional channels. Hence, this type of referral tends to reduce income inequality.

Finding Good Data

An important obstacle to understanding the effects of referrals on labor market outcomes has been the availability of good data. Few data sets contain detailed information about how a worker found their current job or how a firm found its current employees. Moreover, of the few data sets that do contain such information, most are drawn from a narrow range of demographics and occupations, which makes it difficult to draw general conclusions. Given these limitations, the literature has found mixed evidence regarding even basic facts, such as the characteristics of workers who use referrals most frequently or whether using a referral has a positive or negative impact on a worker's starting wages.

For our research, Rivers, Topa, and I used the Job Search Survey, a supplement to the Survey of Consumer Expectations, which is administered by the Federal Reserve Bank of New York. The survey is unique in that the respondents are drawn from a representative sample of workers employed in different industries, and the questions provide explicit information about how each employed worker found their current job. In particular, the survey asks whether the worker used a referral to find their current job and, if so, the worker's relationship to the person who provided it.

Using the information contained in the Job Search Survey, we make two key distinctions that are crucial for understanding the role of referrals in the hiring process: We distinguish between

TABLE 1

Skill Requirements Differ Greatly Across Occupations

The effects of referrals depend partly on these differences.

The Nam–Powers–Boyd index across two-digit occupation codes

Occupation	NPB Index
Farming, Fishing, and Forestry Occupations (FARM)	10
Food Preparation and Serving Related Occupations (FOOD)	17
Building and Grounds Cleaning and Maintenance Occupations (BLDG)	17
Personal Care and Service Occupations (PERS)	27
Transporation and Material Moving Occupations (TRSP)	32
Production Occupations (PROD)	33
Construction and Extraction Occupations (CSTR)	34
Health Care Support Occupations (NURS)	39
Sales and Related Occupations (SLS)	43
Office and Administrative Support Occupations (ADMN)	47
Installation, Maintenance, and Repair Occupations (MNT)	47
Protective Service Occupations (PROT)	55
Arts, Design, Entertainment, Sports, and Media Occupations (ART)	64
Community and Social Service Occuaptions (SOC)	72
Eduation, Training, and Library Occupations (EDU)	75
Health Care Practitioners and Technical Occupations (DOC)	78
Business and Financial Operations Occupations (BUS)	81
Life, Physical, and Social Science Occupations (LIFE)	83
Management Occupations (MGT)	84
Architecture and Engineering Occupations (ENG)	86
Computer and Mathematical Occupations (COMP)	87
Legal Occupations (LEGL)	88

Data Sources: Calculated by Monica Boyd and Charles B. Nam from the 2010–2012 American Community Survey (ACS), Ruggles et. al. 2010

Note: The occupational classification is from the ACS Occupation Code (OCC) variable in the ACS. Scores are aggregated to 2-digit occupation level. See Monica and Nam (2015) for more details.

different types of jobs and *different types of referrals*.

To classify different types of jobs, we measured the skill content of each employed worker's stated occupation using the Nam-Powers-Boyd (NPB) occupation index, which ranks occupations from 0 to 100 based on the earnings and educational attainment of workers in that occupation (Table 1).

To classify different types of referrals, we distinguished between workers who indicated they were referred by a friend or relative and workers who indicated they were referred by what we call a "business contact," which could include a former coworker, supervisor, or business associate. We classified the remaining workers, who say they used another job-finding channel (such as the employer's website, an online search engine, or a headhunter), as nonreferred.

frequently for high-skill jobs.[5] This suggests that referrals from friends and relatives and those from business contacts are qualitatively different.

We confirmed the importance of this difference by studying the relationship between each type of referral and subsequent labor market outcomes. We first examined the relationship between the use of a referral and workers' wages. We found that employed workers who used a referral from a business contact to find their current job had, on average, about 15 percent higher starting wages than nonreferred workers, controlling for observable characteristics of the worker, the occupation, and the time period. In contrast, workers who used a referral from a friend or relative had, on average, about 5 percent lower starting wages than nonreferred workers. Then, we examined the relationship

FIGURE 1

Referrals from Family and Friends Are More Associated with Low-Skill Jobs

Whereas referrals from business associates are more associated with high-skill jobs.

The relationship between the fraction of currently employed workers who used a referral and the position's skill requirements (as measured by the Nam–Powers–Boyd index); the size of each dot is proportional to the number of individuals within each occupation

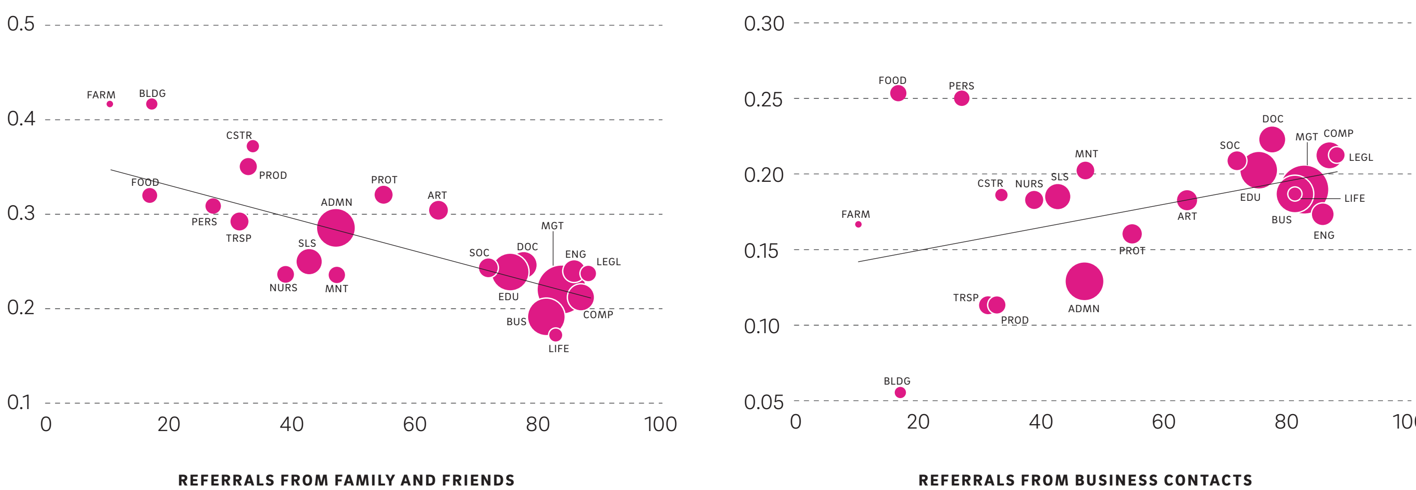

Data Sources: Nam–Powers–Boyd index: calculated by Monica Boyd and Charles B. Nam from the 2010–2012 American Community Survey (ACS), Ruggles et. al. 2010; currently employed workers who used a referral: author's calculations based on data from the Survey of Consumer Expectations, Center for Microeconomic Data, Federal Reserve Bank of New York

It may seem obvious that these distinctions would be helpful for understanding the data. After all, the manager of a fast-food restaurant and the hiring manager at a tech firm presumably find referrals useful for different reasons. And, of course, a referral from your grandmother is very different from a referral from your former boss. However, previous researchers didn't make these distinctions, in large part because of the data limitations described above.

New Facts About Referrals

To start, we analyzed which type of referral is used most frequently in each type of occupation (Figure 1). Referrals from family and friends are used most frequently for low-skill jobs, while referrals from business contacts are used relatively more

between whether a worker used a referral—and, if so, what kind of referral they used—and their tenure at the job. We found that workers who found their current job through a business contact tended to leave the job more quickly, whereas workers who found their current job through a friend or relative stayed at the job for longer.

If using a referral from a business contact is associated with finding a good, high-paying job, why do these workers tend to leave that job quickly? Similarly, if using a referral from a friend or relative is associated with a low-paying job—even after controlling for the worker's occupation—why do these workers stay longer? To answer these questions, we exploited a unique feature of the survey that allowed us to study the experiences of these workers after they found a job. In particular, the survey contains information about the arrival rate of new job offers received by currently employed workers. We find that workers who found their current job through a referral from someone in their business network are subsequently offered jobs by other employers more often. However, we found that workers who found their current job through a referral from a friend or relative are subsequently offered jobs by other employers less often.

To summarize, our analysis reveals that referrals from business contacts are used more frequently at high-skill jobs; they are associated with higher starting wages; and they produce jobs with shorter tenures, on average, because workers who use a referral from a business contact often continue to receive outside offers after they start a job. In contrast, referrals from family and friends are used more frequently at low-skill jobs; they are associated with lower starting wages; and they produce jobs that last longer, on average, because workers who use a referral from a family member or friend tend to receive fewer outside offers from other employers.

What Referrals Do

In the previous section, we described statistical relationships between the channel that workers use to find a job and their subsequent labor market outcomes. However, these relationships alone are not sufficient to identify the effect of referrals on workers' wages and tenure, or on economic inequality in the aggregate. Consider, for example, the positive relationship between the use of a business referral and a worker's starting wages. This relationship is consistent with at least *three* theories of referrals.

See ***Theories of Referrals and Their Implications*** ↓

First, some economists argue that a referral simply helps create a good match by communicating information about the worker to the firm and about the job to the worker. A second theory is that referrals are a way for firms to use their network of employees to find the best, most productive workers–that is, workers who are not just a good match for their vacancy but good at most jobs. Finally, some economists posit that being hired through a referral can help make any worker more productive–perhaps because the referrer serves as a mentor, or even because the new employee works extra hard to avoid

Theories of Referrals and Their Implications

The literature provides different theories about the role of referrals in the hiring process and the subsequent effects of referrals on labor market outcomes.

One theory is that referrals simply help workers and firms find each other. More specifically, "matching frictions" make it difficult for workers to know enough about all the available jobs, and for firms to know enough about all the workers interested in their position. According to this theory, the primary role of referrals is to overcome these matching frictions by connecting workers with firms—that is, a referral could be nothing more than a current employee telling their firm, "I have a friend who is looking for work."

A second theory is based on the idea that it's often difficult to predict whether a worker and a firm will be a good match. Economists call this phenomenon "symmetric uncertainty," since both the worker and the firm are unsure about the prospects for an employment relationship. However, someone who knows both the firm and the worker might be able to determine whether it would be a good match. In this case, the referrer is not only making the connection but also sharing information with the worker and the firm. Thus, a referral can help overcome symmetric uncertainty.

Another theory of referrals is based on the concept of "asymmetric information." Economists typically believe that it's hard for firms to identify which workers will be highly productive. In this case, a referral from a trustworthy source—such as an employee at the hiring firm—is informative about a worker's productivity, and this helps the firm identify and hire more-productive workers.

According to a fourth theory, a referral can change a worker's behavior after they are hired. For example, in some jobs, workers are tempted to slack off after they are hired—a phenomenon economists call "moral hazard." However, when a worker is referred to the firm by, for example, a relative who would suffer embarrassment if the worker performed poorly, the worker may have extra incentive to work hard. Similarly, if a new hire is referred by a current employee who might serve as a particularly good colleague or mentor for the new hire, the referral could generate a highly productive relationship.

An additional theory posits that using informal networks to help someone get a job is a form of nepotism. That is, a referral is a request for a favor and, as a result, the worker gets a job they otherwise would not get.

These theories lead to very different predictions about the labor market outcomes of workers who use a referral. For example, if the primary role of referrals is to overcome matching frictions, workers hired through referrals should be no more productive than those hired through other channels. Hence, the wages and tenure of referred workers should be like those of the nonreferred.

However, if referrals are used to overcome symmetric uncertainty, asymmetric information, or moral hazard, then referred workers should be more productive—or a "better match"—and hence earn higher wages and stay in the position longer, on average. And if referrals are a form of nepotism, the predictions are exactly the opposite: Less qualified workers should be less productive, earn lower wages, and be quicker to leave the firm, either because they are fired or because they find a better match elsewhere.

embarrassing the person who provided the referral.

In our paper, we used an economic model to help us identify which of these theories are consistent with the relationships we found in the data. We find that referrals from business contacts and those from friends and family play very different roles in the match formation process; that is, to be consistent with the trends we see in the data, we have to apply *different theories* to these two types of referrals.

According to our model, whether a worker uses a business referral is highly sensitive to that worker's underlying productivity. This could be the case if, for example, a business contact is willing to recommend an applicant only if they know the applicant is productive and hard working. Hence, our analysis suggests that referrals from business contacts support the theory that a referral's primary role is to help firms screen applicants and find the best workers. In the parlance of economic theory, business referrals help resolve *asymmetric information*.

Alternatively, referrals from friends and relatives seem to create good matches for all types of workers. This could be the case if, for example, a worker's friend knows which skills they have and helps them find a job that requires those skills. Or it could be the case that a worker's friend or family member agrees to help them find a job, but only with the expectation (explicit or implicit) that the worker shows up on time and works hard. Thus, in the language of economic models, our findings suggest that referrals from family and friends more closely support the theory that referrals overcome *symmetric uncertainty* or *moral hazard*.

Implications for Inequality

Without a consensus on how referrals are used in the hiring process, economists also disagree about the effect of referrals on inequality. Naturally, this makes it hard to formulate policy advice. For example, nepotism laws that prohibit workers from referring a family member could either exacerbate or ameliorate income inequality.[6] Our results suggest that the relationship between referrals and inequality depends on the type of referral and the type of job.

On the one hand, referrals from business contacts are used most frequently at high-paying, high-skill jobs and by highly productive workers. According to our model-generated counterfactuals, these workers may use a business referral to generate offers, but they also frequently find opportunities through other, formal channels. Hence, though business referrals help form good matches, they contribute to earnings inequality by helping well-paid workers increase their wages even further.

Referrals from family and friends, on the other hand, are quite different. They are used more frequently at low-paying, low-skill jobs and by workers who struggle to generate offers through other channels. In other words, these referrals are often a worker's last resort. Therefore, referrals from friends and relatives, like referrals from business contacts, help form good matches—albeit through a different mechanism. However, unlike referrals from business contacts, they tend to ameliorate earnings inequality by helping workers at the bottom of the wage distribution find a job with decent wages.

Conclusion

Why do some workers find jobs quickly while others struggle? Why do similar workers get paid different wages? What determines how long a worker stays in their job? To answer these fundamental questions, labor economists seek to better understand the process that connects a worker with a firm.

In surveys of workers and firms, a referral is often cited as an ingredient in this process. However, the precise role of referrals and the implications for labor market outcomes have been unclear, in part because of data limitations. Our recent research uses a new survey to show that the role of a referral—and its effect on workers' wages and tenure—becomes clear once we distinguish between different types of referrals, and how each type of referral is used to find different types of jobs.

Our new insights could help rationalize a variety of puzzling facts about the labor market. For example, economists often struggle to understand why certain workers don't leave their hometowns in search of better work prospects. Our findings suggest a reason: Workers who depend on family and friends to find jobs are understandably reluctant to leave this network behind. EI

Notes

1 Diamond, Mortensen, and Pissarides were awarded the 2010 Nobel Memorial Prize in Economic Sciences for their work on markets with search frictions, of which the labor market is a prime example. See, for example, Diamond (1971, 1982), Pissarides (1985), and Mortensen and Pissarides (1994).

2 Topa (2011) provides an extensive review of the prevalence of referrals in surveys of both workers and firms. Most surveys of job seekers find between 50 and 60 percent of workers report having used a referral to find their current job. Surveys of firms also indicate widespread use of referrals or word of mouth: Those results vary from just under 40 percent to significantly more than 50 percent.

3 See, for example, Calvo-Armengol and Jackson (2004), who developed a theoretical model to explain how referrals through networks can exacerbate inequality. Of course, there are other important sources of economic inequality, including those that arise from differences in *capital* income, but these are beyond the scope of our focus on labor markets.

4 See Loury (2006) for a more detailed discussion of the role of referrals as a last resort for certain workers.

5 In the paper, we use regression analysis to confirm that the trends in Figure 1 are statistically significant, even after controlling for observable characteristics of the worker, along with time and region fixed effects.

6 A deeper question is whether there is "too much" inequality and, if so, whether it's wise for policymakers to address this issue with policies that affect the matching process or with policies that redistribute income after matches have been formed. We focus here on the relationship between referrals and inequality, without taking a stand on this (admittedly important) question.

References

Boyd, Monica, and Charles B. Nam. "The Newest Nam–Powers–Boyd Occupational Scale: Development and Insights," paper presented at the Southern Demographic Association annual meeting, San Antonio, October 14–16, 2015, www.npb-ses.info.

Calvo-Armengol, Antoni, and Matthew O. Jackson. "The Effects of Social Networks on Employment and Inequality," *American Economic Review*, 94:3 (2004), pp. 426–454, https://doi.org/10.1257/0002828041464542.

Diamond, Peter A. "A Model of Price Adjustment," *Journal of Economic Theory*, 3:2 (1971), pp. 156–168, https://doi.org/10.1016/0022-0531(71)90013-5.

Diamond, Peter A. "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90:5 (1982), pp. 881–894, https://doi.org/10.1086/261099.

Lester, Benjamin R., David Rivers, and Giorgio Topa. "The Heterogeneous Impact of Referrals on Labor Market Outcomes," Federal Reserve Bank of New York *Staff Report*, 987 (2021), https://doi.org/10.2139/ssrn.3937491.

Loury, Linda Datcher. "Some Contacts Are More Equal than Others: Informal Networks, Job Tenure, and Wages," *Journal of Labor Economics*, 24:2 (2006), pp. 299–318, https://doi.org/10.1086/499974.

Mortensen, D. T., and C. A. Pissarides. "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61:3 (1994), pp. 397–415, https://doi.org/10.2307/2297896.

Pissarides, Christopher A. "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, 75:4 (1985), pp. 676–690.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]. Minneapolis: University of Minnesota, 2010.

Topa, Giorgio. "Labor Markets and Referrals," *Handbook of Social Economics*, vol. 1. North-Holland: Elsevier, 2011, pp. 1193–1221, https://doi.org/10.1016/B978-0-444-53707-2.00005-0.

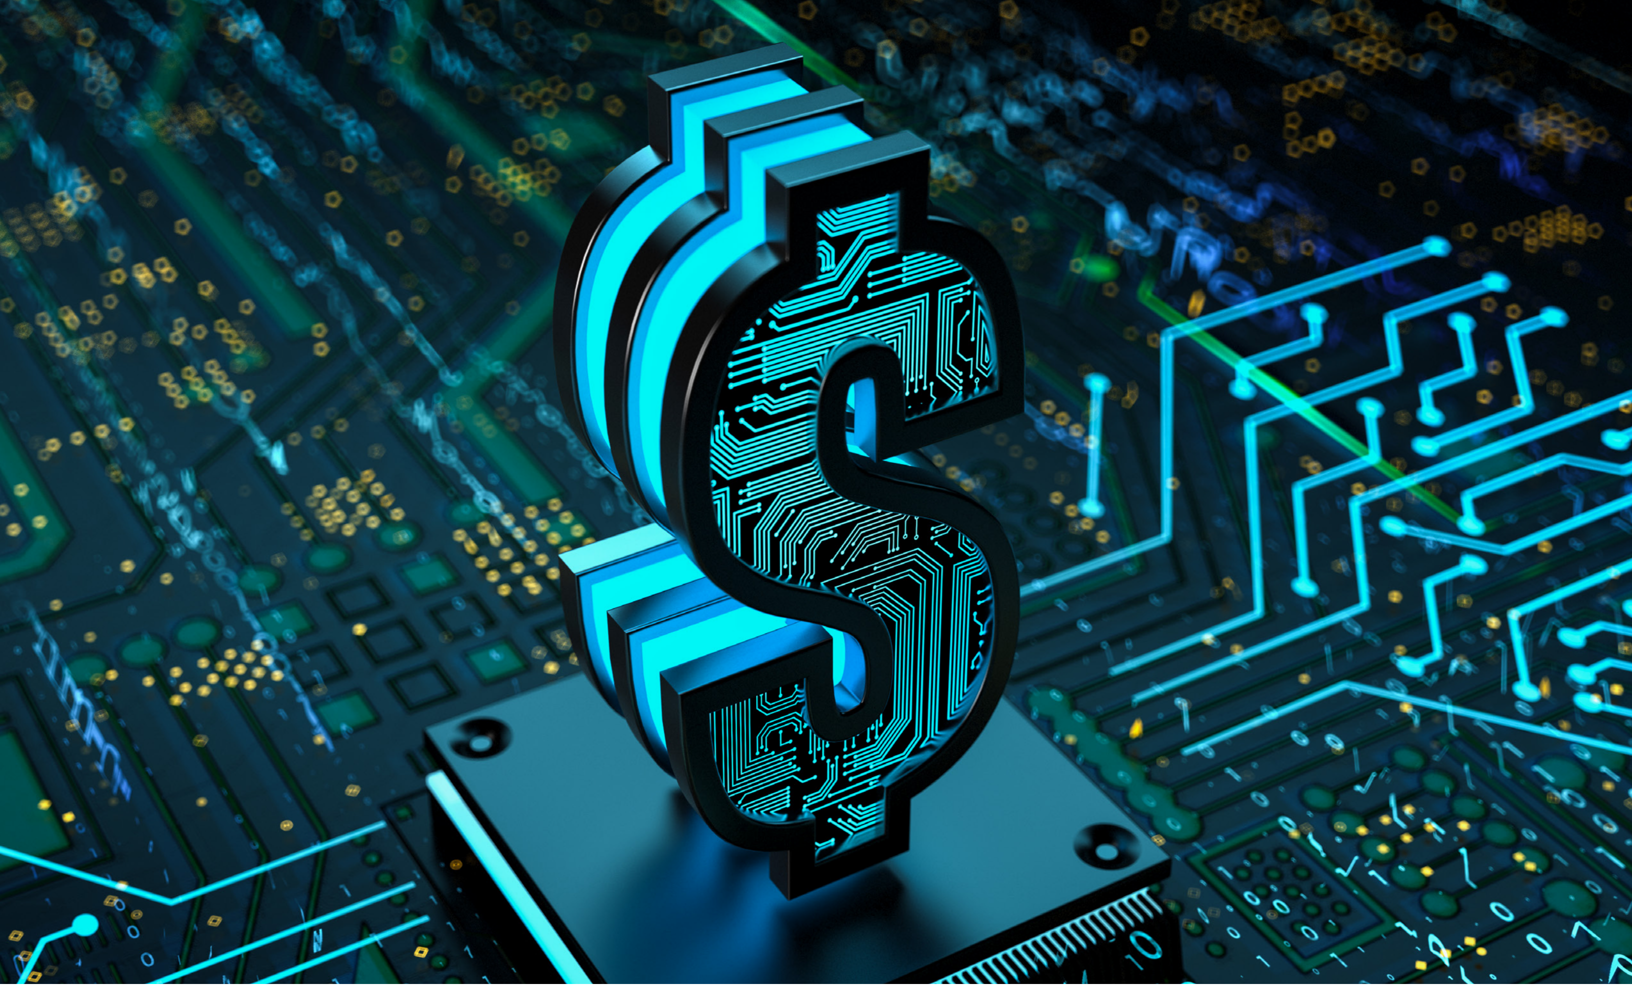

Making Sense of Decentralized Finance

What is DeFi? And how does it differ from traditional finance?

Joseph Abadi
Economist
FEDERAL RESERVE BANK OF PHILADELPHIA

The views expressed in this article are not necessarily those of the Federal Reserve.

Decentralized finance (DeFi) is a system of financial platforms built on *public blockchains*—immutable, open-access ledgers that record the ownership of cryptocurrencies and other digital assets. Just as cryptocurrencies aim to provide an alternative to traditional currencies such as the dollar, DeFi platforms are an alternative to traditional finance (TradFi) platforms such as stock exchanges or credit card payment systems.[1]

DeFi's proponents argue that TradFi is rife with inefficiencies and rent-seeking intermediaries. DeFi's value proposition, therefore, is to create a new financial system that will better serve users. To fulfill this mission, DeFi platforms differ from TradFi platforms in two key respects: their transaction technology and their ownership structure. DeFi platforms automate transactions, thereby eliminating the need for any centralized intermediary that executes transactions, such as an exchange. This also circumvents TradFi's legacy transaction-processing systems, which are often older and less technologically advanced. Unlike a shareholder-owned TradFi platform, a DeFi platform is collectively owned and governed by its users as well as key insiders such as the founding team and software developers. Users thus have some authority to run a DeFi platform according to their own interests.

DeFi platforms have been deployed across a wide range of traditional financial applications as well as some new ones. Some platforms facilitate collateralized, peer-to-peer lending, providing a substitute for bank-intermediated credit. Other platforms ("decentralized exchanges") allow users to trade cryptocurrencies with one another directly on the blockchain, just as investors would trade stocks on a traditional stock exchange. Still other platforms issue "stablecoins," digital assets whose value is pegged to that of a fiat currency, such as the dollar or the euro.

summarized the optimists' case for DeFi: "There [are] a lot of intermediaries that end up charging 20-30%, and if the concept of decentralization takes off, then those [fees] are also going to decline to near zero."

Others, however, believe that DeFi's potential to cut costs and give authority to users is greatly overstated. A recent report by the Bank for International Settlements (BIS) summarizes the pessimists' stance: "There is a 'decentralization illusion.' First and foremost, centralized governance is needed to take strategic and operational decisions. In addition, some features in DeFi... favor a concentration of power."[2] Indeed, despite its recent and rapid rise, DeFi still does not play a central role in the broader financial system, and it is unclear whether it will eventually provide a widely used alternative to TradFi.

FIGURE 1

DeFi Platforms Increased by a Factor of 10 Amid a Surge in Cryptocurrency Prices

Total dollar value of assets deposited on DeFi platforms, in millions, 2020–2023

Data Source: DeFi Llama (2023)

Although DeFi has existed since the popular Ethereum blockchain launched in 2015, DeFi activity took off in earnest in 2020, when transaction volumes on DeFi platforms increased by a factor of 25 amid a general surge in cryptocurrency prices (Figure 1). This drove interest in DeFi products among both retail and institutional investors, but the DeFi sector has faced some headwinds since then. The collapse of the Terra stablecoin in May 2022 and of the FTX cryptocurrency exchange in November of the same year shook investors' confidence in the safety of cryptocurrency and DeFi products. Nevertheless, DeFi has defied predictions that it would quickly die in the wake of these crises, and transaction volumes remain well above their pre-2021 levels.

Economists and practitioners disagree about the future of DeFi. Optimists believe that in the long run, DeFi's transaction technology and governance structure will prove vastly superior to TradFi's. Vitalik Buterin, founder of the Ethereum blockchain,

Will DeFi's automated transaction technology result in lower fees charged to users? Will its decentralized ownership structure succeed in redistributing decision-making power to users? To answer these questions, and to judge the merits of the pro and con arguments, we need to compare DeFi with TradFi in terms of costs and of ownership structure.

Lowering Costs by Automating Transactions

To manage transactions, TradFi uses intermediaries such as banks or centralized exchanges–but DeFi's proponents argue that these intermediaries often use their market power to take advantage of users. (For example, they can *extract rents* by charging users high fees.) Instead of using intermediaries, DeFi users transact with one another via *smart contracts*–software protocols that automatically execute trades once a sequence of if-then conditions is met. Hence, the main benefit of automated transactions is to cut out rent-seeking intermediaries. Furthermore, because trades are executed by code, there is no need for costly court proceedings or arbitration when a contract is breached. Rather, the smart contract automatically imposes a penalty for misbehavior (for example, by seizing collateral from the defaulting party).

However, DeFi transactions are not costless. Users must pay transaction fees to *validators*, a set of users who run computer programs that certify blockchain transactions. These fees are often substantial and in principle can exceed those paid to TradFi intermediaries. There is also an implicit cost of automatically executed DeFi contracts because there is no way to renegotiate a contract if unforeseen circumstances arise. Moreover, there is no recourse in cases of fraud or theft, as when a malicious counterparty exploits a vulnerability in a smart contract's code.

Governance by Users

DeFi's smart-contract-based settlement system is not the only feature that sets it apart from TradFi. Just as important is DeFi platforms' ownership structure: One of DeFi's ambitions is for platforms to be owned and governed by a "decentralized" community of users.

A TradFi platform is typically controlled by a manager who acts primarily in the shareholders' interests. The interests of other stakeholders, such as workers, suppliers, and creditors, are protected by contractual claims on specific payments. Workers have employment contracts with agreed-upon salaries, suppliers sell services at contractually specified prices, and creditors are owed payments of a fixed maturity.

Shareholders, on the other hand, are *residual cash flow claimants*: They receive whatever is left over after contracts with all other constituencies are paid out. Shareholders therefore need some degree of influence over the firm's management to ensure that they receive a return on their investment. Otherwise, shareholders might not receive anything from management. The protection of shareholders' interests has traditionally been viewed as the central problem in corporate governance and a primary focus of legislation. Although shareholders do not make day-to-day strategic decisions, managers legally have a fiduciary duty to act on shareholders' behalf.

A DeFi platform has no residual claimants whose interests need to be protected. The platform's code governs how all cash flows are distributed to stakeholders. For example, the platform's smart contracts specify the transaction fees charged by the platform, compensation for the platform's software developers, and the profits that will be distributed back to users. In contrast to TradFi, then, there is no problem of protecting shareholders. The main question in DeFi governance, rather, is who gets to write the platform's code. DeFi's governance model specifies that users themselves should decide how the code is written.

DeFi platforms delegate decision-making power to users and key insiders by issuing digital assets called *tokens*. On most DeFi platforms, each token is worth one vote, so users' voting power is directly proportional to their token holdings. To ensure that some tokens end up in users' hands, many platforms reward users with tokens when they provide liquidity to the platform. A user can provide liquidity by making an asset available for others to purchase or borrow. The user deposits the asset in an escrow account owned by a smart contract, at which point other users can purchase or borrow it at a contractually specified price. Liquidity providers are rewarded with newly minted tokens. The smart contract specifies how many new tokens a liquidity provider will receive for depositing an asset for a fixed period.

This incentive scheme, known as "yield farming" (or sometimes "liquidity mining"), has been a significant institutional feature in the emergence of DeFi. The total quantity of assets deposited on a DeFi platform (called the total value locked, or TVL) is widely considered to be the most reliable metric of that platform's popularity. The "DeFi summer" of 2020, indeed, followed shortly after the introduction of yield farming on decentralized lending platforms and exchanges: From June to October 2020, the aggregate TVL on DeFi platforms jumped from $1 billion to $10 billion. Today, DeFi platforms continue to promote their yield farming policies to attract new users.

However, not all tokens are awarded to users. A DeFi platform will typically also issue new tokens as compensation for founders, software developers, and the venture capitalists that initially funded the platform. As such, users do not necessarily have full ownership of a platform; rather, they usually share ownership with these insiders.

Token holders vote on key policy decisions, such as liquidity providers' rewards, transaction fees, and the transaction protocol for the platform's smart contracts. A user community known as a *decentralized autonomous organization* (DAO) runs the voting process. Token holders can propose changes to the platform's policies in a DeFi platform's DAO, at which point the policy change is put to a vote. If the measure passes, then the DAO immediately updates the platform's smart contracts to reflect the new policy.

An Example of DeFi

To better understand how DeFi works, let's examine a DeFi lending platform that enables collateralized lending across several digital assets. Most DeFi lending platforms, such as Compound, Aave, or Cream Finance, share the same basic design. Unlike in traditional bank-intermediated credit markets, borrowing and lending rates on a DeFi lending platform are not set by financial intermediaries. Instead, a type of smart contract called a "lending pool" determines interest rates algorithmically.

Consider the following scenario with a hypothetical DeFi platform (henceforth "the platform"). Lenders have 200 units of popular stablecoin USD Coin (USDC) they would like to lend. Borrowers would like to borrow 100 units of USDC, and they have some Ether (the cryptocurrency issued by the Ethereum blockchain) they can post as collateral. Lenders deposit their $200 of stablecoins in the USDC pool on the lending platform. Borrowers then borrow $100 from the pool and deposit enough Ether to cover the required margin—say, $150 (Figure 2). The interest rate paid by borrowers is algorithmically determined by the pool's *utilization rate*, which is the ratio of borrowed USDC to deposited USDC (in this case, 0.5). The higher the utilization rate, the higher the interest rate paid by borrowers. Lenders receive the interest rate paid by borrowers minus a spread collected by the platform. The pool automates default penalties: If a borrower fails to repay a loan, or if the value of their collateral declines too much relative to the value of the loan, the smart contract will close out the borrower's position and deliver the collateral to lenders.

The platform issues new tokens as a reward to lenders. In our earlier example, USDC lenders receive newly minted tokens as additional yield on their loans. For instance, if lenders receive an interest rate of 2 percent from borrowers, and the platform provides a 50-basis-point subsidy in tokens, then the total return to lenders is 2.5 percent. Tokens also grant users the right to propose policy changes, receive a share of the platform's transaction fees, and vote on governance decisions in the platform's DAO.

the DAO determine the design of the platform's smart contracts. For instance, the platform's policies determine the interest rate schedule faced by borrowers, the spread charged by the platform, collateral requirements, and the token subsidy received by lenders. The platform's users therefore have some power to set interest rates themselves, whereas a bank's borrowers and depositors typically do not.

Living Up to Its Promise: Costs

DeFi's proponents argue that by cutting out the middleman, DeFi will be much cheaper for users than TradFi. However, the reality for now looks quite different. DeFi transactions usually incur a substantial fixed cost: Users must pay validators a fee to include their transactions in the blockchain. The cost of a small transaction is currently much higher in DeFi than in TradFi. For example, on the Ethereum blockchain, the average transaction

FIGURE 2

How a DeFi Platform Can Enable Collateralized Lending

Lenders receive the interest rate paid by borrowers minus a spread collected by the platform.

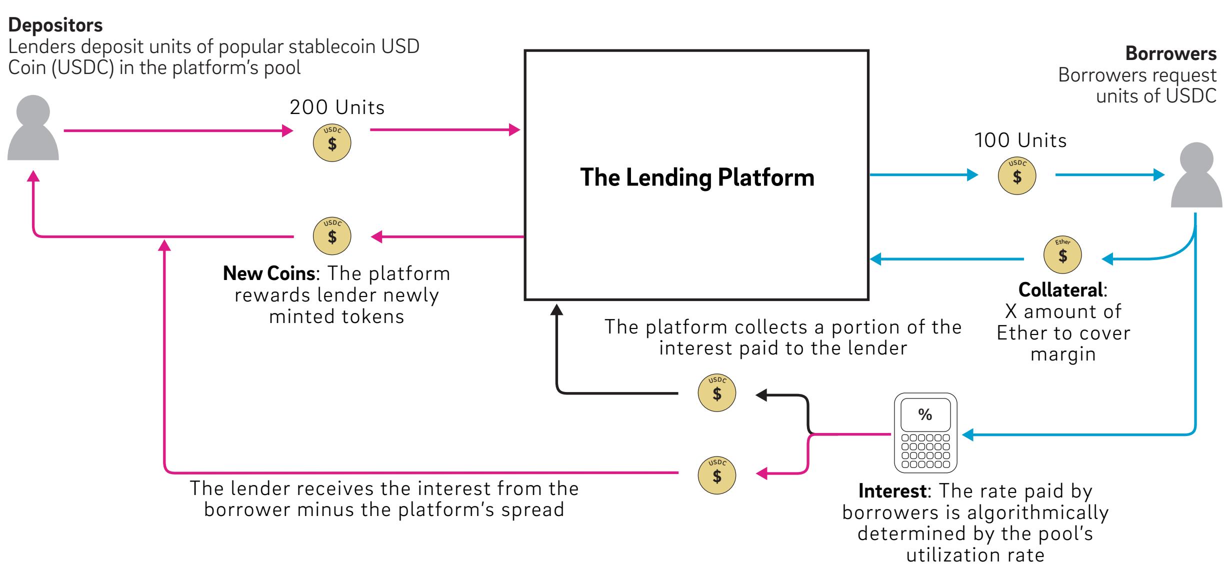

The lending platform sets interest rates on its loans and earns a spread, just like a traditional bank. But it's different from a bank in two important ways.

First, the platform sets interest rates algorithmically, whereas a bank has the discretion to set interest rates however it likes. A bank can therefore *price-discriminate*: It can charge higher interest rates to borrowers it perceives to have a greater need for credit. The platform, on the other hand, cannot price-discriminate. The interest rate charged to a borrower is defined by a set of observable loan-specific characteristics, such as the size of the loan or the quality of the collateral.

Second, the platform is governed in part by its users (who hold tokens), whereas a bank is not necessarily governed by its borrowers or depositors.[3] The key governance decisions made in

fee is $32.[4] For comparison, a typical fee charged to a vendor in a $100 credit card transaction would be about $2. DeFi's high fees are especially detrimental to retail consumers who would like to use DeFi for everyday financial transactions.

DeFi transaction fees are high primarily because DeFi platforms, unlike TradFi platforms, have not scaled. The problem is neither a lack of demand for DeFi transactions nor a lack of validators who want to process them. Rather, popular DeFi platforms have run up against technological limits on their transaction-processing capacity. The Ethereum network, for instance, processes 15 to 20 transactions per second,[5] whereas Visa's network can process 65,000.[6] Due to these scale constraints, DeFi transaction fees have grown as these platforms have become more popular: As transaction demand increases, users must pay higher transaction fees to ensure their transactions will be included in the blockchain. Several proposals have been put forth to enhance DeFi's scalability, such as breaking up large blockchains into smaller pieces that only occasionally communicate with one another (a process called "sharding"). These proposals, however, are untested, so it is not clear whether DeFi will be able to scale and sustain broad-based use by retail consumers.

In other areas, though, DeFi shows promise. Tobias Adrian, an economist with the International Monetary Fund, argues that DeFi platforms incur much smaller *marginal* lending costs than banks.[7] Unlike TradFi platforms, DeFi platforms do not have to cover significant labor, operational, or regulatory compliance costs. As a result, DeFi platforms can charge smaller lending spreads. DeFi borrowing may therefore be attractive to large borrowers, such as firms: On a large enough loan, the additional cost of a $32 transaction fee is easily compensated for by a lower spread.

DeFi platforms have a long way to go before they are substantially more effective than TradFi, but there is no reason to believe that DeFi's transaction technology can't surpass TradFi's. Moreover, even if DeFi never overtakes TradFi in popularity among retail customers, it may nevertheless prove superior in some specific applications, such as lending to large borrowers.

FIGURE 3

Most Tokens Go to Insiders

As a result, decision-making on DeFi platforms remains highly concentrated.

The share of tokens distributed to insiders in 15 large initial coin offerings (ICOs) as of 2022

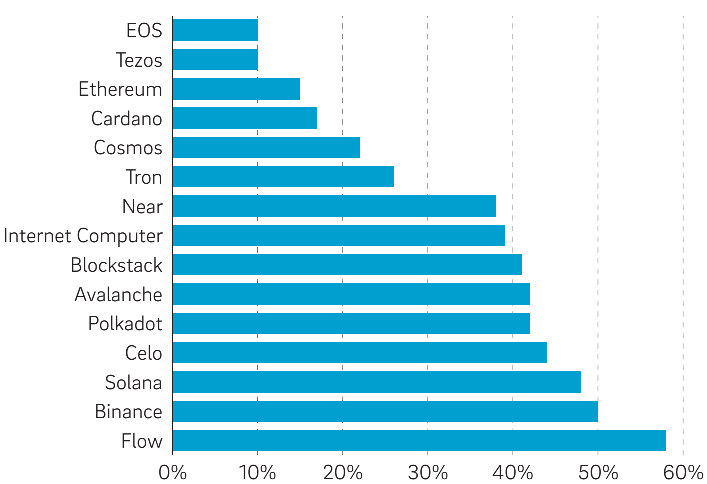

Data Source: Watkins (2021)

Living Up to Its Promise: Governance

DeFi's proponents argue that a platform is more likely to be run in users' interests if it is governed by users themselves rather than by profit-seeking intermediaries. Of course, for this argument to have merit, DeFi must succeed in distributing decision-making power to users. So far, however, it has largely failed to do so.

Decision-making power in DeFi remains highly concentrated for two reasons. First, token holdings tend to be concentrated in a few hands; insiders such as the founding team and venture capital funders often retain a substantial fraction of a platform's tokens as compensation. In several large *initial coin offerings* (which are the equivalent of initial public offerings for traditional corporations), the share of tokens distributed to insiders exceeded 30 percent (Figure 3).

Even large liquidity-mining rewards do not necessarily resolve this issue. If liquidity mining activity is highly concentrated, then a handful of users will earn a large share of newly issued tokens. Also, most small-scale users are reluctant to vote on platform policies. The average user has little incentive to participate in a vote, since that vote is unlikely to influence the outcome. There may also be formal or informal barriers to small users' participation: Some platforms require users to hold a minimum quantity of tokens before they can vote, and users sometimes lack the technical expertise necessary to understand policy proposals.

Due to concentrated token holdings and low user participation, a DAO's decisions are often made by only a small set of individuals. For the average DAO proposal, the majority of voting power is controlled by three or fewer individuals.[8] Decision-making power in a DAO may thus be even more concentrated than on the board of a traditional corporation.

Without broad user participation in DeFi governance, how can DeFi accomplish its goals? The entire point of DeFi is to avoid the concentration of decision-making power and "decentralize" authority. But there may still be hope for DeFi's governance model. When a platform passes a DAO proposal that decreases barriers to voting, the platform's token price tends to increase.[9] Because insiders hold large quantities of these tokens, this incentivizes them to encourage broad-based user participation. Indeed, voting power in DAOs is becoming less concentrated: The Herfindahl-Hirschman index (HHI), a measure of the concentration of voting power in DAOs, declined from 0.45 in 2020 to 0.30 in 2022.[10]

Conclusion

DeFi platforms aim to improve upon the traditional financial system by combining two innovations: a "smart contract" transaction technology and a "token holder" governance model. Although DeFi's transaction technology promises to reduce costs in some applications, it cannot, on its own, redistribute these economic gains to users. To achieve their goal, DeFi platforms must be governed in a way that is consistent with users' interests. So far, however, DeFi has not decentralized decision-making power to users in the way its proponents had hoped. Future progress in DeFi will require not only technical advances in smart contract and distributed ledger design but also economic solutions to the governance problems faced by these platforms.

Efficient Governance

Economists say that a platform's governance structure is *efficient* if it maximizes the total economic value the platform generates for its stakeholders. DeFi's proponents argue that user-owned platforms are likely to be governed more efficiently than traditional shareholder-owned platforms.

The main argument in favor of user ownership is that it protects users from rent extraction, making for a more efficient platform. Both TradFi and DeFi platforms exhibit *network effects*: The platform's service gains value as its user base grows. For instance, a credit card is useful to consumers only if enough merchants are willing to accept it. Similarly, a peer-to-peer lending app is useful to lenders only if borrowers use it. Thus, financial transaction platforms require a minimum number of users to function, which limits the extent of cross-platform competition. Credit card payment processing, for example, is a highly concentrated market, with credit card payments in the U.S. processed by just three companies: Visa, Mastercard, and American Express.

Financial transaction platforms can leverage their market power to boost profits at users' expense. In practice, platforms often do so by charging users various fees, such as the credit card payment processing fee charged to merchants (which is 2.2 percent on average).[11] Platforms may also extract rents by selling users' transaction data to third parties. For instance, stock trading apps sometimes sell user order flow data to high-frequency trading firms.[12] This type of rent extraction can be inefficient: High fees or other costs can dissuade users from transacting on a platform.

DeFi's token holder governance model views rent extraction as the main source of inefficiency in the governance of financial platforms. *Transaction cost* theories of organizational structure emphasize a similar argument: If a group of stakeholders stands to be exploited by a firm's market power, then it is sometimes efficient for those stakeholders to own the firm.[13]

However, other theories of corporate governance argue precisely the opposite: Platforms should be governed by an accountable group of knowledgeable insiders rather than a dispersed community of users. Governing a platform is difficult. It requires technical expertise and a capacity to coordinate, both of which users often lack. DeFi users may not be familiar enough with a platform's code to understand proposed policy changes. Moreover, deliberations among members of a DAO, which typically take place informally on message boards run by a platform, can lead to deadlocks. Since *every* proposed policy change must pass through a DAO, DeFi platforms may struggle to adapt to changing conditions if they rely on broad-based user participation in governance.

DAOs therefore tend to rely on founders and developers to guide upgrades to the platform. These insiders' technical expertise grants them informal authority that exceeds their voting power, as predicted by the theory of organizations.[14] Users defer to insiders' judgment on questions of protocol design—but this concentrated decision-making power may not serve users' best interests. Unlike the managers of a corporation, these insiders do not have a fiduciary duty to anyone: They may pursue their own interests, even at users' expense. For example, a platform's developers may be reluctant to upgrade the platform's code to make transactions faster. Doing so would benefit users, but it could also be costly for developers.

It is thus unclear whether the token holder governance model can decentralize authority and advance users' interests. Token holder governance could mitigate rent extraction and benefit users, but a DAO's reliance on a decentralized community of users raises inefficiencies of its own. DeFi platforms must deal with the same governance problems as other types of organizations, and those problems cannot be eliminated by smart contracts.

Notes

1 For additional discussion of the competition between cryptocurrency and traditional currencies, see Sanches (2018).

2 Aramonte, Huang, and Schrimpf (2021).

3 Mutual savings banks are an important exception to this rule. Instead of distributing profits to shareholders, these banks distribute profits back to depositors.

4 Makarov and Schoar (2022).

5 Data on Ethereum's transaction throughput are taken from Blockchair's Ethereum Transactions per Second Chart, https://blockchair.com/ethereum/charts/transactions-per-second.

6 See Visa UK's Fact Sheet, https://www.visa.co.uk/dam/VCOM/download/corporate/media/visanet-technology/aboutvisafactsheet.pdf.

7 The marginal cost of lending is defined as the additional cost incurred for each additional $1 of lending. See Adrian (2022).

8 Appel and Grennan (2023b).

9 Appel and Grennan (2023a).

10 Appel and Grennan (2023b).

11 Daly (2023).

12 Williams (2021).

13 Hansmann (1988).

14 Aghion and Tirole (1997).

References

Adrian, Tobias. "Discussion of *Cryptocurrencies and Decentralized Finance*," Bank for International Settlements 21st Annual Conference (June 24, 2022).

Aghion, Philippe, and Jean Tirole. "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105:1 (1997), pp. 1–29, https://doi.org/10.1086/262063.

Appel, Ian, and Jillian Grennan. "Decentralized Governance and Digital Asset Prices," Darden Business School Working Paper 4367209 (2023), https://dx.doi.org/10.2139/ssrn.4367209.

Appel, Ian, and Jillian Grennan. "Control of Decentralized Autonomous Organizations," *AEA Papers and Proceedings*, 113 (2023), pp. 182–185, https://doi.org/10.1257/pandp.20231119.

Aramonte, Sirio, Wenqian Huang, and Andreas Schrimpf. "DeFi Risks and the Decentralisation Illusion," *BIS Quarterly Review* (2021), p. 21–36 https://www.bis.org/publ/qtrpdf/r_qt2112b.pdf.

Daly, Lyle. "Average Credit Card Processing Fees and Costs in 2023," *The Ascent: A Motley Fool Service* (2023).

DeFi Llama, retrieved from defillama.com on December 18, 2023.

Hansmann, Henry. "Ownership of the Firm," *Journal of Law, Economics, & Organization*, 4:2 (1988), pp. 267–304, https://doi.org/10.1093/oxfordjournals.jleo.a036953.

Makarov, Igor, and Antoinette Schoar. "Cryptocurrencies and Decentralized Finance (DeFi)," National Bureau of Economic Research Working Paper 30006 (2022), https://doi.org/10.3386/w30006.

Sanches, Daniel R. "Bitcoin vs. the Buck: Is Currency Competition a Good Thing?" Federal Reserve Bank of Philadelphia *Economic Insights* (Second Quarter 2018), pp. 9–14, https://www.philadelphiafed.org/the-economy/bitcoin-vs-the-buck-is-currency-competition-a-good-thing.

Watkins, Ryan. "Power and Wealth in Cryptocurrencies," Messari (2021), https://messari.io/report/power-and-wealth-in-cryptoeconomies.

Williams, Michael. "Robinhood and the CCPA: Selling Your Personal Data to the Rich?" *Clym* (2021).

Research Update

These papers by Philadelphia Fed economists, analysts, and visiting scholars represent preliminary research that is being circulated for discussion purposes.

The views expressed in these papers are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Philadelphia or Federal Reserve System.

The Changing Polarization of Party Ideologies: The Role of Sorting

Ideology scores derived from U.S. congressional roll-call voting patterns show that the ideological distance between the two parties along the primary dimension changes inversely with the ideological distance along the secondary dimension. To explain this inverse association, a model of party competition with endogenous party membership and a two-dimensional ideology space is developed. If the distribution of voter preferences is uniform on a disk, equilibrium ideological distances along the two dimensions are inversely related. The model can quantitatively account for the historical movements in ideological distances as a function of changes in the ideological orientation of the two parties.

WP 24-4. Satyajit Chatterjee, Federal Reserve Bank of Philadelphia; Burcu Eyigungor, Federal Reserve Bank of Philadelphia.

CECL Implementation and Model Risk in Uncertain Times: An Application to Consumer Finance

I examine the challenges of economic forecasting and model misspecification errors confronted by financial institutions implementing the novel current expected credit loss (CECL) allowance methodology and its impact on model risk and bias in CECL projections. We document the increased sensitivity to model and macroeconomic forecasting error of the CECL framework with respect to the incurred loss framework that it replaces. An empirical application illustrates how to leverage simple machine learning (ML) strategies and statistical principles in the design of a nimble and flexible CECL modeling framework. We show that, even in consumer loan portfolios with tens of millions of loans, like mortgage, auto, or credit card portfolios, one can develop, estimate, and deploy an array of models quickly and efficiently, and without a forecasting performance penalty. Drawing on more than 20 years of auto loans data and the experience from the Great Recession and the COVID-19 pandemic, we leverage basic econometric principles to identify strategies to deal with biased model projections in times of high economic uncertainty. We advocate for a focus on resiliency and adaptability of models and model infrastructures to novel shocks and uncertain economic conditions.

WP 24-3. José J. Canals-Cerdá, Federal Reserve Bank of Philadelphia.

Inference Based on Time-Varying SVARs Identified with Sign Restrictions

We propose an approach for Bayesian inference in time-varying SVARs identified with sign restrictions. The linchpin of our approach is a class of rotation-invariant time-varying SVARs in which the prior and posterior densities of any sequence of structural parameters belonging to the class are invariant to orthogonal transformations of the sequence. Our methodology is new to the literature. In contrast to existing algorithms for inference based on sign restrictions, our algorithm is the first to draw from a uniform distribution over the sequences of orthogonal matrices given the reduced-form parameters. We illustrate our procedure for inference by analyzing the role played by monetary policy during the latest inflation surge.

WP 24-5. Jonas E. Arias, Federal Reserve Bank of Philadelphia; Juan F. Rubio-Ramírez, Emory University, Federal Reserve Bank of Atlanta; Minchul Shin, Federal Reserve Bank of Philadelphia; Daniel F. Waggoner, Emory University.

Price-Level Determination Under the Gold Standard

We present a microfounded monetary model of a small open economy to examine the behavior of money, prices, and output under the gold standard. In particular, we formally analyze Hume's celebrated price-specie flow mechanism. Our framework incorporates the influence of international trade on the money supply in the Home country through gold flows. In the short run, a positive correlation exists between the quantity of money and the price level. Additionally, we demonstrate that money is nonneutral during the transition to the steady state, which has implications for welfare. While the gold standard exposes the Home country to short-term fluctuations in money, prices, and output caused by external shocks, it ensures long-term price stability as the quantity of money and prices only temporarily deviate from their steady-state levels. We discuss the importance of policy coordination for achieving efficiency under the gold standard and consider the role of fiat money in this environment. We also develop a version of the model with two large economies.

WP 24-6. Jesús Fernández-Villaverde, University of Pennsylvania, NBER, and CEPR; Daniel Sanches, Federal Reserve Bank of Philadelphia.

Research in Focus

Summaries of Working Papers

A Quantitative Model for Mapping the Consequences of Public Housing Demolitions

Using Chicago's housing stock as the basis for a quantitative model, the authors reveal the household-level economic consequences of demolishing public housing.

Exploring an Unexpected Driver of Racial Segregation and Its Long-Term Effects on Economic Mobility

A study of racial segregation finds negative consequences for long-term economic mobility, yielding evidence of a little-studied enabler of segregation itself: railroad location.

Measuring National Economic Activity by Race

A Philadelphia Fed research associate has developed an index of U.S. Economic Activity by Race (EAR), identified differential findings by race, and explored the uses of EAR.

Is Loan Volume Becoming Concentrated Among Too Few Online Mortgage Providers?

By studying a period of substantial regulatory change and technological innovation, three researchers identify shifts in market concentration within the U.S. mortgage market.

How Effective Are Place-Based Industrial Policies?

A Philadelphia Fed economist and his coauthors study place-based subsidy programs, including their impact on reducing wage inequality between Turkey's various provinces.

Data in Focus

Tri-State Tracking

The Philadelphia Fed collects, analyzes, and shares useful data about the Third District and beyond. Here's one example.

Community outreach is a key function of a Federal Reserve district bank. Staff from several departments regularly traverse the district, visiting community banks, chambers of commerce, educational institutions, and trade organizations to learn about "on the ground" economic conditions and to share our research with an interested public.

Several years ago, we developed Tri-State Tracking as a handout for audiences attending these outreach events. By summarizing key data in this handout, we educate Third District stakeholders about current economic conditions in each of the district's three states.

It was only natural to add the product to our Bank website so that anyone, anywhere can see these data. All the data in Tri-State Tracking come from other sources, including monthly employment data from the U.S. Bureau of Labor Statistics (BLS) and quarterly personal income data from the U.S. Bureau of Economic Analysis. By collecting different state-level data in one place, we save you the trouble of having to visit multiple websites and then isolate the data for whichever state you're interested in.

In this Data in Focus, we feature one of the most popular and relevant variables included in Tri-State Tracking: Year-over-Year Payroll Employment Growth. Payroll job gains or losses over the year can vary considerably between states and among economic sectors, as well as from one month to the next, so this snapshot is useful for anyone interested in understanding how their state's industrial structure is shifting. In the December 2023 BLS state employment release, for example, payroll employment in the Professional and Business Services sector grew in Pennsylvania and Delaware but fell in New Jersey. This is precisely the sort of detail that policymakers and public advocates in New Jersey might miss if they were looking only at national (or total state) payroll employment growth.

FIGURE 1

Employment Estimates

The 12-month difference in payroll employment estimates, number of jobs in thousands, by sector for each Third District state from the December 2023 BLS state employment release (released January 23, 2024)

Source: U.S. Bureau of Labor Statistics via Haver Analytics.

Learn More

Online: https://www.philadelphiafed.org/surveys-and-data/regional-economic-analysis/tristate-tracking
E-mail: Elif.Sen@phil.frb.org

This page was intentionally left blank.

FEDERAL RESERVE BANK
PHILADELPHIA

Ten Independence Mall
Philadelphia, Pennsylvania 19106-1574

PRESORTED STANDARD
U.S. POSTAGE

PAID

PHILADELPHIA, PA
PERMIT #583

ADDRESS SERVICE REQUESTED

From Philadelphia, Pennsylvania, the birthplace of American finance, comes *Economic Insights*. Continuing our 100-year tradition of sharing cutting-edge research with the general public.

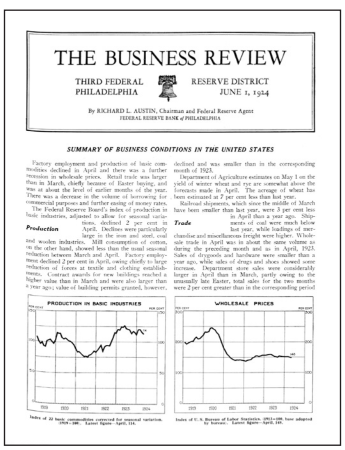

1920s

1950s

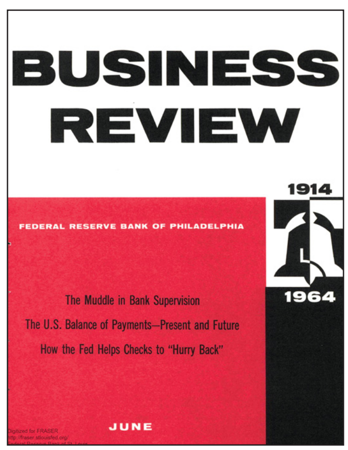

1960s

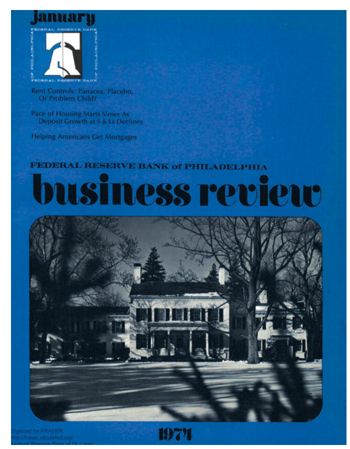

1970s

1980s

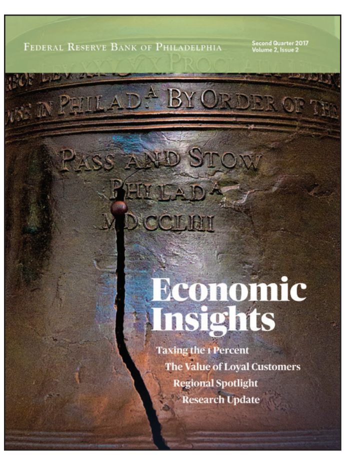

2010s

2020

2023