RESEARCH BRIEF

Taxes, Automation, and the Future of Labor

Daron Acemoglu (MIT), Andrea Manera (MIT), and Pascual Restrepo (BU)

Taxes, Automation, and the Future of Labor*

Daron Acemoglu (MIT), Andrea Manera (MIT), and Pascual Restrepo (BU)

* This research brief was written for the MIT Task Force on the Work of the Future. We thank David Autor and David Goldston for comments and Timothy Aeppel for editorial help.

Introduction

Employers often ask whether a machine or a person is better suited to do a job, based on factors such as speed, precision, and what's technically possible. But one of machines' biggest advantages over a worker may be how they are taxed.

While labor has been taxed at an average rate of about 25% for the past four decades, things like equipment, software, and buildings that are classified as capital are taxed at a lower rate — and this rate has fallen steadily in recent years. For every dollar a worker receives, the business that employs them has to pay an additional 25 cents in taxes, effectively making the cost of labor to employers 25% higher. In contrast, the average tax rate on software and equipment stood at about 15% in the 1990s and fell to about 5% after a series of tax reforms in the 2000s and 2010s.

In essence, the U.S. tax system encourages companies to buy machines while discouraging them from adding workers.

While lower taxes on capital is a long-standing feature of the U.S. tax code, the issue now is more acute as new automation technologies move into a larger variety of worker tasks performed throughout the economy, from warehouses and factories to insurance firms and stores. Automation, which involves the substitution of machines and algorithms for tasks previously performed by workers, could be an engine of growth. But when it is excessive — for example, driven by tax incentives rather than for efficiency gains — it harms labor and fails to improve productivity. In this context, it is notable that the biggest drop in tax rates in recent years has been for software, followed by equipment, exactly the types of capital goods used to build automated systems, including robots, numerically controlled machinery, specialized software, and, more recently, artificial intelligence.

Acemoglu and Restrepo (2019) have argued that there has been faster progress in automation technologies over the last 20 years and much slower progress in technologies complementary to humans. If this view is correct, then U.S. technology looks much more substitutable for human labor than it did in previous decades. In the past, it was more common for every new machine to require at least one or more human operators, and subsidies to capital indirectly helped labor as well. This may no longer be so with rapid advances in automation technologies, making the asymmetric taxation of capital and labor much worse news for labor.

In this brief, we first examine the different forms of automation and then turn to how the U.S. tax system has led to excessive reliance on machines in ways that can hurt workers and subdue wages. We subsequently look at how the tax system evolved to favor capital and what U.S. policymakers could do to reverse this negative impact and level the playing field between capital and labor.

Automation and Work

Technological advances have been the main driver of economic growth and improvements in living standards and wages over the last two centuries. Yet, not all technologies are created equal. Automation can have very different effects on employment and wages than technologies that increase human productivity.

It is thus critical to distinguish automation from other types of technologies and capital deepening (meaning an increase in the amount of capital and machinery used in production). First, businesses may invest in machinery that is complementary to workers — think of upgrading the machine tools workers are using. Second, firms may buy more or better machines for performing tasks that are already automated — think of replacing an automated welding machine with an advanced robot. Such (non-automated) capital deepening is generally helpful for labor, because it increases the firm's productivity and encourages it to demand more labor.

Automation works very differently from capital deepening, however, because it is substituting machines for workers in the tasks that they were previously performing. One implication of this substitution of machines for humans is that automation has an unambiguous negative impact on the labor share — how much of a firm's, an industry's, or the economy's revenue goes to labor. By reducing the set of tasks performed by workers and increasing those performed by capital, automation always increases the capital share and reduces the labor share. This change has obvious distributional implications beyond the automation's impact on employment and wages.

What about automation and employment? The same substitution implies that automation displaces workers from these tasks, creating a force toward lower labor demand. Nevertheless, the productivity effect is still operational. When capital can perform some tasks more cheaply than labor, automation reduces the firm's costs and in effect makes it more productive. As the firm expands to take advantage of these productivity gains, it also demands more labor. We thus see the two countervailing effects of automation on labor: task displacement, which reduces labor demand, and productivity improvements, which increase labor demand. If the former force dominates, automation will reduce not only the labor share but also employment and earnings. However, if the productivity gains that automation brings are more powerful than the displacement that it causes, then jobs and earnings will tend to rise, even as workers are displaced from some tasks.

So-So Technologies

The problem is that not all automation creates this coveted productivity effect. In some cases, which Acemoglu and Restrepo (2019) call "so-so technologies," the new automation isn't much better in terms of productivity than the human labor it is replacing.[1] Because so-so technologies do not improve productivity much, their end result is an overall decline in demand for labor.

An example of so-so technology is self-checkout machines increasingly found at retail stores. Those machines shift the task of checking out groceries to consumers, eliminating the need for cashiers. And, since an average shopper is probably less adept at tallying groceries than a skilled cashier who doesn't have to struggle to figure out if the tomatoes are organic, or if a can of soup is really on sale, it likely slows the checkout line rather than accelerates it. And yet, the development and installation of those self-service machines benefit from the same tax advantages for capital as productivity-enhancing automation, and they likely result in fewer jobs.

Another ubiquitous example of so-so technology is automated customer service, which displaces workers but is not always satisfactory for users. Even very advanced automation technologies could be so-so, as witnessed by the very high level of automation adopted in Tesla factories that later turned out not to have improved productivity.

So-so technologies used for automation are a double whammy for labor: They displace workers and do not add much to productivity. If they have been prevalent, this would explain, at least in part, why labor demand and wages have been near stagnant. The question, which we turn to next, is why businesses may have been so keen on adopting so many so-so technologies.

Taxes and Automation

This framework is helpful in understanding how the tax rate on capital works its way through the economy to impact whether the number of jobs grows, stagnates, or declines. Let us start with automation and consider a simple thought experiment: Suppose we lived in a world without any taxes and let firms choose which tasks to automate and which to assign to humans. In this world, if the technology allows, firms will automate tasks only where capital is more productive than labor (meaning that the cost of performing a task with the machine is less than the cost of performing it with labor). Alternatively, if capital doesn't hold that edge, the job will continue to be done by workers. There is no incentive to push one over the other without the corresponding cost reduction or productivity gain.

So, what happens when we subsidize capital — or equivalently, put a relatively heavier tax on labor? It will naturally encourage firms to prefer machines instead of labor, even in "marginal" tasks — those that formerly would have been considered not worth automating, because it would not increase productivity.

However, this also implies that the asymmetric tax treatment of capital and labor can lead to "excessive automation" — automation beyond what a pure cost-minimization logic would dictate.

In this context, our earlier contrast between capital deepening and automation leads to an important observation: The fact that there is too much automation in specific cases like a grocery self-checkout does not imply that the ratio of capital to labor is too high in the economy as a whole or even in all automated tasks. Recall that workers benefit when there is more machinery that is complementary to them, or when better machines are introduced in already automated tasks. So, an overall increase in the amount of capital in the economy will benefit labor by raising productivity and labor demand. But this is not necessarily so when new machines are used for automating tasks previously performed by labor: In that case, as we previously noted, labor may benefit or lose out, depending on the extent of the displacement and productivity effects.

The important lesson for our discussion of the tax structure is that a blanket reduction in the capital-labor ratio would not be beneficial for labor. That would merely hamper productivity growth, harming workers in the process. The goal should not be to reduce capital as a whole, but rather to identify those cases where automation is being used for marginal tasks simply because the tax system encourages it.

Excessive Automation

To be sure, there may be reasons other than tax policy for excessive automation. For instance, corporations whose business model or vision is centered on automation may overlook promising non-automation technologies as they race to introduce new automation machines or algorithms. This problem would become

even more severe if these companies were to dominate many sectors or even the economy itself. There could also be a bias in the research and development process that favors too much automation. For example, machinery, software, or algorithms used for automation may come to be viewed as more interesting and important to the research community; or worse, researchers may start to believe that an automation-focused path of technological progress is inevitable.

Fads and visions in the research community are not the only reason why there may be excessive automation. When firms decide whether to use machines or workers, they may ignore the possibility that workers themselves are much better off when employed, rather than when they are "on the dole" or out of the labor force. Firms also may not consider the benefits that hiring workers creates for their families, the broader community, and perhaps even our democracy.

All of these considerations exacerbate the concerns we may have about a tax policy that encourages greater automation: When we are already automating excessively, taxes leading to even more automation would be doubly costly.

In fact, a tax policy that leads to excessive automation not only costs jobs, but also is likely to reduce paychecks for the workers who remain, since, as we have seen, automation depresses the labor share. In 1980, labor's share of national income stood at 63.6%, meaning that nearly two-thirds of every dollar earned in the economy went into the pocket of workers. It fell to 56.6% by 2017.[2] The drop in labor's share of income isn't the result of low productivity. From 1980 to 2017, total output per capita doubled. Instead, a lot of that productivity growth has gone to capital rather than to labor. And, the workers who *have* benefited are those at the top of the skill and wage distribution. While the paycheck of the highly educated workers with specialized skills has continued to grow, the median wage has been by and large stagnant, and the real wages of low-education workers, especially low-education men, have fallen sharply.[3]

In the meantime, the U.S. economy experienced the rapid spread of automation technologies. Automated machinery and numerical control became commonplace in U.S. manufacturing, and industrial robots started being installed in large numbers in the 1990s. In 1993, the United States had a mere 2.5 industrial robots for every 1,000 factory workers. Over the next 16 years, this number increased eightfold, to 20 per 1,000 workers. Industrial robots were just the tip of the automation iceberg, however. It wasn't just blue-collar jobs that witnessed rapid automation. Various specialized software and subsequently algorithms started automating white-collar jobs, a process likely to continue with the advent of artificial intelligence and machine learning. At the same time, the demand for workers with advanced skills continued to grow.

Of course, many forces have contributed to stagnant wages in recent decades, especially for the lowest-paid workers, including the offshoring of jobs and competition from cheap imports.[4] Nevertheless, many experts see a close connection between the growth of automation and some of the aforementioned

adverse labor market developments.[5] Recent work by Acemoglu and Restrepo (2020) suggests that automation may have been the most important factor boosting inequality in the U.S. labor market.

And yet, the most common perspective among economists is that even if automation is contributing to a decline in labor's share of national income, the adoption of these new technologies is beneficial. According to this view, the best way to deal with such adverse consequences is through a combination of ex-post redistribution to help those left behind by automation and investments in education and training so that fewer people are left behind in the first place. While both redistribution toward the less fortunate and broad-based education and training investments are much needed in the U.S. economy today, they are not the solution to excessive automation. If there is indeed excessive automation, then we as a society are obtaining a lower rate of return on our capital investments and failing to maximize productivity relative to what we would be able to achieve with "appropriate automation." The right policy response to these trends should then be to make automation more appropriate, if necessary by reducing incentives for automating marginal jobs with social technologies.

A first step would be to start redressing the highly asymmetric nature of the U.S. tax system. However, it is also important to understand how we came to grant capital such an edge over labor.

Capital and Labor Taxes in Recent Decades

Figure 1 shows the evolution of the effective tax rates on labor and different types of capital since 1980. The asymmetric treatment of capital and labor is clear: The advantage of capital relative to labor has not been constant over the decades. Rather, capital, especially software and equipment capital, is increasingly taxed far less heavily than labor.

Figure 1: Effective Taxes (In Percent) on Labor and Software, Equipment and Structures Capital, 1981–2018.

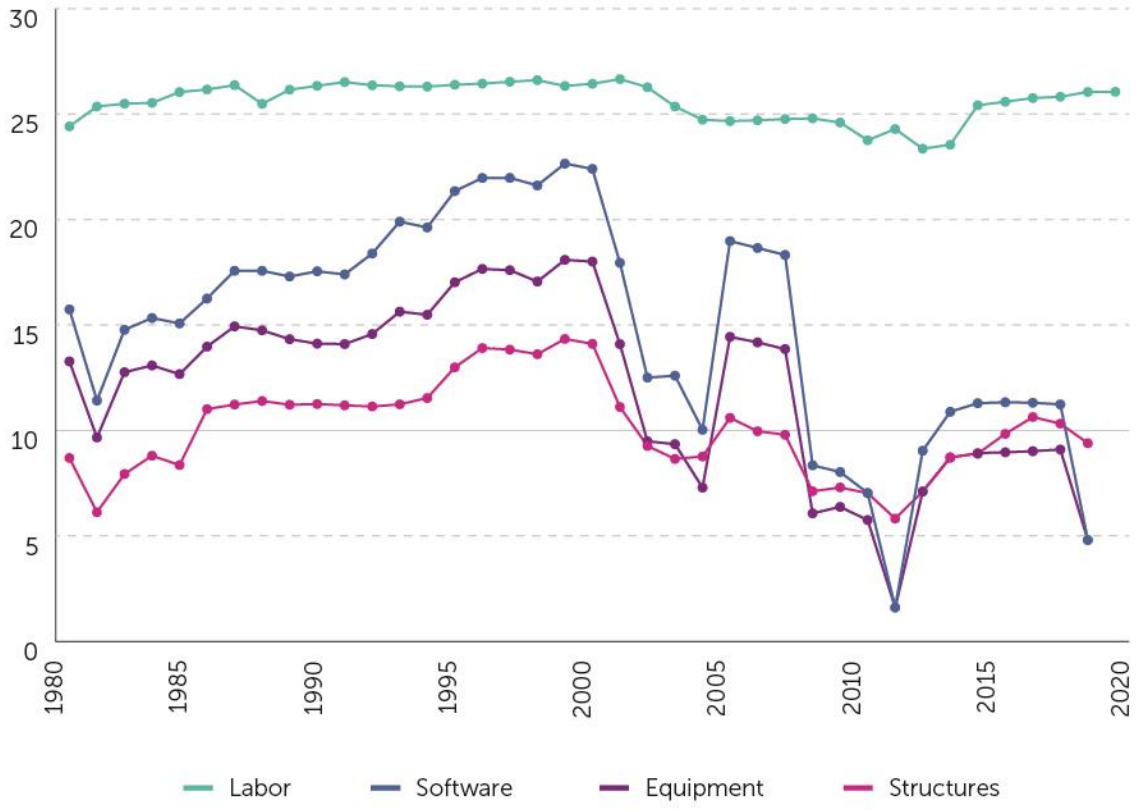

Source: Acemoglu, Manera and Restrepo, 2020

The "effective" qualifier here is important, and explaining how we calculated these figures helps to show the complexity faced by anyone attempting to revise U.S. tax policies. There isn't a single tax on labor or capital, and there's a great deal of heterogeneity in how corporations and different types of capital income are treated by the U.S. tax code. There are similarly diverse taxes on labor. Any attempt to construct a single effective tax rate for one or the other requires making a number of assumptions.

Let's start with how we estimated labor taxes, which are a bit more straightforward than taxes on capital. These consist of payroll taxes and federal and local income taxes, both of which create a gap between what employers pay for labor and what workers receive. Federal and state income taxes are progressive — meaning that those with lower incomes are taxed less heavily. Payroll taxes are regressive, however,

since employees pay 6.2% of their earnings for Social Security until a maximum of 132,000 (they pay the 1.45% Medicare tax on all wages).

In other words, the top 5% of earners don't pay this tax at all. For the purposes of this study, we took the average income taxes faced by those below the 95th percentile and combined that with payroll taxes. Since automation generally doesn't replace those at the very top of the income pyramid, focusing on the bottom 95% makes sense.

Figuring a single effective tax for capital is more challenging. In Acemoglu, Manera, and Restrepo (2020), we do this by distinguishing between different types of corporations — which fall into three categories: C-corporations, S-corporations, and other privately held corporations — and equity-financed and debt-financed investments. C-corporations pay corporate taxes, while the two other types of corporations do not. In recent years, more companies have sought to transition their classification from C-corporation to S-corporation in order to take advantage of beneficial tax treatment. Whether a business is financed with debt or equity also matters, because interest paid on debt is tax deductible, further reducing the tax burden of debt-financed capital investments. Capital owners pay federal and income and state taxes as well; because they tend to be rich, they face the higher marginal tax rates that apply at the top of the distribution. Recent declines in these high-income marginal tax rates have also reduced the effective tax rate on capital.

Even more important than the type of corporation in recent decades has been depreciation allowances, which enable companies to reduce their tax bill by deducting what they spend on capital. Depreciation allowances are complicated in practice, because they depend on the type of capital involved, its estimated life, and the tax rate faced by the corporation claiming the deduction. All the same, they play a crucial role in understanding the level and evolution of the effective tax rate in the United States, and their broad outlines are uncontroversial. Depreciation allowances enable corporations to deduct capital expenditures from their tax obligations and have been a major factor in reducing taxes on capital. But they have become much more generous over the last two decades, allowing firms to write off their expenditures immediately and fully.[6]

The reason is simple: Depreciation allowances are particularly popular with policymakers, who view them as a way to quickly and decisively spur capital spending (not to mention the fact that they are loved by corporations as well). This is especially important at times of economic recession, when the government looks for ways to revive growth. The periodic sharp drop in taxes on capital, followed by upturns (see Figure 1) in large part reflect these special provisions. However, these ebbs and flows do not take away from the long-term downward trend in taxes on capital.

Is Taxing Capital a Bad Idea?

If we are correct that the tax code can spur excessive automation, why has the tax code evolved to favor capital over labor — and increasingly so? We don't know for sure, but two factors seem to be important. Capital owners often have greater sway over policymakers than workers, so it is reasonable for them to receive more favorable treatment. Most commentators believe that the power of the rich, and thus of capital, has increased over the last several decades, due to both the decline of organized labor and the increased role of money in politics.[7] These developments may have contributed to making the tax code more "capital-friendly." But equally important has been the policy advice from many economists that taxing capital is always a bad idea.[8] From this perspective, in the 1980s the United States was taxing capital not too little but too much, and the cuts in taxes on capital in the 2000s and beyond were a key policy objective.

The argument undergirding this view is that capital taxes distort savings and discourage capital accumulation, harming everybody in the economy. But economic taxes are ultimately borne by the factors of production — labor, capital, and land — so there is simply no way of avoiding distorting something through taxation.

Because taxation is (almost) inevitably distortionary, how does one balance the distortion among the various sources of taxation? A longstanding economic insight is that factors such as labor, capital, and land should be taxed in inverse proportion to their supply elasticity — that is, how responsive their availability is to shifts in taxation. Put differently, factors that don't respond much to taxation, such as the supply of land, are good candidates for taxation because one distorts their supply very little by taxing them. One can't create more land to take advantage of lower taxes (though some land can come in and out of the market in response to taxes). Conversely, higher taxes won't reduce the supply of land. By this same logic, factors that respond significantly to taxation may be less ideal candidates.

This brings us back to the classic argument for *not* taxing capital. It's often assumed that the supply of capital is infinitely elastic, meaning that it is very responsive to taxes, and in fact even a small tax will heavily distort spending on it. If that were true, then taxing capital would indeed be a bad idea. But the evidence does not support this assumption.

In Acemoglu, Manera, and Restrepo (2020), we looked at the most reliable studies on the elasticities of capital and labor and concluded that both capital accumulation and labor supply are modestly responsive to taxation. In other words, taxation discourages both and in fairly comparable ways. This suggests that under more realistic conditions, they should both be taxed relatively similarly, and that's one of the key arguments we make in this brief.

In fact, there may be reasons, if anything, to err on the side of taxing labor more lightly than capital, because of the potential reasons for excessive automation we spelled out earlier, and the fact that labor market imperfections seem to already distort labor supply significantly.

Even if one is dubious that capital should be taxed no more favorably than labor, there are good prudential reasons for tending toward equal taxation in practice. For example, it's hard for tax rules to actually distinguish capital from labor.[9] If capital is taxed less heavily than labor, experience shows that business owners will soon find ways to relabel the returns to their labor as returns to physical capital in order to gain favorable tax treatment.

Take a successful tech entrepreneur, such as Mark Zuckerberg. The success of a company such as Facebook critically depends on the thousands of hours of labor its founders invest early on as well as their ideas and skills. Yet, these founders would receive much of the fruits of that success as capital income; and because capital is taxed more lightly than labor, this would significantly reduce their tax obligations. So, the asymmetric treatment of capital most likely ends up being not just distortionary, but also another handout to corporations, providing us with another reason for treating capital and labor comparably in the tax code.

The reality is that governments have to levy taxes to raise revenue, and there is a strong conceptual and empirical case for taxing labor and capital about equally. And yet, the United States has been moving precisely in the opposite direction in recent decades.

The ultimate solution is to alter U.S. tax policy to level the playing field, or at least bring the trade-off between capital and labor closer into balance. The most direct approach would involve rolling back depreciation allowances and clamping down on evasion of corporate income taxes by curbing the ability of companies to classify themselves as S-corporations. These policy actions would broaden the capital income tax base, and that might be enough to bring capital and labor income taxes in line with each other. If this is insufficient to create balance, somewhat higher corporate income taxes and higher taxes on the highest income earners could be necessary.

In practice, there are likely to be limits on how far policymakers can go in correcting the bias of the tax code in favor of capital. High taxes on capital may encourage capital flight and may run into opposition from powerful corporations. Tax reforms are difficult in the best of times, and may have to be gradual.

Still, there are obvious tax reforms that can be implemented. First, the depreciation allowances, which are responsible for as much as half of the decline in the effective tax rate of software and equipment,[10] were meant to be temporary, and thus the opposition to lifting them may be less vociferous. Second, much of the effective tax on labor comes from payroll taxes, which could be reduced (though this would then call for

other taxes to fund Social Security and Medicare at a time when government budget deficits have already reached new highs).

Automation Taxes

If there are limits on how far policymakers can go in such reforms, there may be other remedies. These would have to get to the root cause of the problem: excessive automation.

Alas, discouraging excessive automation is not easy either. It is challenging to stimulate productivity-enhancing automation without at the same subsidizing so-so technologies. The underlying capabilities needed for both forms of automation are often the same. Research into advanced sensor technology, for instance, is crucial for developing many new AI-powered applications that promise to create big productivity gains in coming years, and sensors can be a crucial complement to human operators and drivers. But the same technology would likely be just as useful in developing even better self-checkout machines that just displace more workers.

Acemoglu, Manera, and Restrepo (2020) argue that one possibility would be to impose "automation taxes," which would simultaneously correct some of the distortions created by the asymmetric tax code and leave businesses to decide which automation technologies are the ones they should be investing in. The idea behind the automation tax is not to target all forms of capital automation, but rather to focus taxes on automation technologies that are being used for marginal tasks where it does not bring much productivity gain. More explicitly, automation taxes should not be applied to cases where automation boosts productivity and adds jobs. Of course, it is a challenge for the tax code to identify these "marginal" cases.

All of this implies that automation taxes need to be implemented with care. One has to start with a road map of which types of tasks and activities are the marginal ones where there is excessive automation.

Though this question requires additional research, it is clear that structures capital (buildings) as well as capital types that work closely with labor, such as machine tools and many computer applications, should not be the target of automation taxes. Neither should automation taxes target tasks that have been successfully automated in the past, such as those now performed by industrial robots or numerically controlled machinery. This would amount to taxing capital deepening — which, as we argued earlier, would hinder rather than help labor, because more capital in already-automated tasks increases labor's productivity.

Our discussion makes it clear that automation taxes are different from taxes on capital. Capital taxes apply to all types of capital, including those such as structures or manual machine tools that are

complementary to labor. The point, however, is not to reduce the capital intensity of the economy, but to discourage the automation of marginal tasks. On the contrary, as we have seen, capital deepening helps labor, and discouraging it would be yet another factor holding back labor demand.

The same argument applies to wealth taxes. Imposing taxes on the very wealthy may have political justification — such as to limit excessive political power of the rich. But it is not the type of tax that we are calling for here either. This is because, like capital taxes, wealth taxes would apply to all types of capital; and, further complicating matters in this case, they would only do so when these capital goods are bought or owned by the wealthiest investors.

Automation taxes are also different from "robot taxes," which have entered the public discourse in recent years. For one, automation taxes do not apply to all robots — they target marginal tasks, not all automation. In contrast, robot taxes would go after some of the more well-established types of automation in manufacturing (which are unlikely to be marginal anymore or hardly likely to be reversed). Equally important, automation technologies go far beyond robots — and in the decades to come, algorithms will be much more central to the process of automation. Hence, the case for concentrating taxes on robots but not numerically controlled machines, specialized software, and artificial intelligence is not compelling.

Changes to the tax code such as these not only would help bolster income for low-wage workers and create more jobs, but also could help lessen inequality.

How would all this work in practice? That's the trillion-dollar tax question. It's not feasible to write tax policies that say "you get a tax break on capital investment — but not if it's funding machinery that makes a process only marginally more productive." There is a similar problem with trying to steer investments through the use of R&D tax credits. It becomes unmanageable to control whether R&D is only being used to build artificial intelligence that is used for marginal applications.

Meanwhile, implementing an automation tax or cutting depreciation allowances might prove politically impossible in an environment where any reform that goes against corporations or capital is likely to encounter powerful opposition. Perhaps the best answer is that the tax code cannot tell these things apart, which is another good reason to simply treat capital and labor symmetrically for purposes of taxation. Even if it does not restore full efficiency and eliminate all excessive automation, a neutral tax system is at least agnostic in its treatment of these factors and goes some way toward redressing the deep-rooted asymmetries in the U.S. tax system.

Summing up

The U.S. tax code treats capital more favorably than labor and has become even more asymmetric over the last 20 years. This is likely to be highly distortionary, especially when it encourages more and more automation.

There is a belief in some policy and academic circles that eliminating taxes on capital is optimal from the viewpoint of maximizing output and average welfare. This isn't so. Taxes on both labor and capital should be guided by their relative supply elasticities. Our assessment of the literature is that the elasticity of the two are comparable, and hence a broadly equal treatment of capital and labor in tax policy is likely to be optimal. The goal should be neither subsidizing labor nor discouraging capital formation. Rather, the ideal is to ensure a level playing field.

Tilting the playing field in favor of capital is costly in the best of times. It is doubly so at a time of rapid advances in and enormous enthusiasm about automation technologies, because the favorable treatment of capital and tax policy encourages firms to substitute machines for workers even when this is unwarranted. We end up with a double whammy: distortions in the production process and lack of jobs and wage growth.

Redressing the asymmetries of the U.S. tax system can therefore have major benefits. Our estimates suggest that tax reforms that level the playing field between capital and labor can discourage so-so automation, stimulate employment and significantly improve welfare. According to our baseline calculations, moving toward an optimal balance in taxes between capital and labor would raise employment by 5.85% and increase labor's share of income by 0.53 percentage points, without any reduction in government revenues. More modest reforms would increase employment and labor's share, although not to the same degree. Even if these numbers are speculative, the fact that there is much to gain by removing the imbalances between capital and labor remains.

References

Acemoglu, Daron, and David H. Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4, edited by Orley Ashenfelter and David Card. Amsterdam: Elsevier.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. 2020. "Does the U.S. Tax Code Favor Automation?" Brookings Papers on Economic Activity, Spring 2020 Edition.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33, no. 2: 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2020. "The Labor Share, Displacement and Wage Inequality." Mimeo. Massachusetts Institute of Technology.

Atkeson, Andrew, V. V. Chari, and Patrick J. Kehoe. 1999. "Taxing Capital Income: A Bad Idea." *Federal Reserve Bank of Minneapolis Quarterly Review* 23: 3–18.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103, no. 6: 2121–68.

Autor, David H., David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics* 135, no. 2: 645–709.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton & Company.

Ford, Martin. 2019. *Rise of the Robots*, Basic Books, New York.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2016. *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. "Capitalists in the Twenty-First Century." *Quarterly Journal of Economics* 134, no. 4: 1675–1745.

End Notes

1 Here, by "productivity" we are referring to "total factor productivity," which takes into account that capital is also being used in production and thus evaluates output relative to the whole bundle of inputs. Automation mechanically increases labor productivity (output divided by labor), because it reduces the number of workers as it substitutes machines for humans. Total factor productivity is the relevant notion for evaluating both how much these technologies are contributing to our overall national income and generating additional demand for labor.

2 See, for example, Autor, Dorn, Katz, Patterson, and Van Reenen (2020).

3 Acemoglu and Autor (2011).

4 Autor, Dorn, and Hanson (2013).

5 See, for example, Brynjolfsson and McAfee (2014) and Ford (2019).

6 Acemoglu, Manera, and Restrepo (2020).

7 McCarty, Poole, and Rosenthal (2016).

8 Atkeson, Chari, and Kehoe (1999).

9 Smith, Yagan, Zidar, and Zwick (2019).

10 Acemoglu, Manera, and Restrepo (2020).