

Taxation of Automation and Artificial Intelligence as a Tool of Labour Policy

By

Vincent OOI*

Glendon GOH**

Abstract

Rapid developments in automation technology pose a risk of mass displacement of human labour, resulting in the need to support and retrain displaced workers (a negative externality). We propose an “automation tax” that would slow the adoption of automation technology in appropriate circumstances, giving workers and social support systems time to adapt. This could be easily implemented through changes to the existing schedular system of depreciation/ capital allowances, reducing the uncertainty of its application and implementation costs. Such a system would be flexible enough to keep up with rapid technological developments. Two main dimensions may be adjusted to produce intended distortionary effects: 1) accelerated depreciation, and 2) bonus depreciation. While the benefits of efficiency gains mean that the automation tax is unlikely to have widespread application, it does provide a useful tool for specific situations where the rate of automation needs to be slowed due to its resultant social costs.

A. Introduction

The positive impact of developments in technology on the economy has historically outweighed the disruptive impact on employment. However, in the present case, the pace of development of the “Fourth Industrial Revolution” presents a risk of mass displacement of human labour, particularly in tasks that are repetitive and menial. Section B of this article lays out the background to these developments and considers their consequences. It goes on to submit that while most workers will be able to continue in their roles after job alteration, some workers may be unable to retain their jobs post-automation because they lack all the skills required to perform the higher-value redesigned job (structural unemployment).

Section C of this article moves on to consider the case for an automation tax. It argues that automation, by inducing worker displacement, results in social costs arising from the need to support and retrain displaced workers, which constitute a negative externality. Thus, an “automation tax” is required to correct this market failure. The aims of an automation tax are two-pronged: first, to slow the introduction of automation technology in industries which

* BA (Oxon), Lecturer, School of Law, Singapore Management University. This research is supported by the National Research Foundation, Prime Minister’s Office, Singapore under its Emerging Area Research Project Funding Initiative. The authors gratefully acknowledge the support of the Centre for Artificial Intelligence and Data Governance, School of Law, Singapore Management University.

** BA, MFE (Oxford). The authors would like to thank Denise Lim and Esther Wong for their insightful comments and assistance. The authors also gratefully acknowledge the capable research assistance of Chloe Chen and Shari Huang.

would otherwise suffer rapid and massive unemployment otherwise, so as to provide as much time as possible for governments, welfare systems, and workers to prepare for the impending effects of structural unemployment; second, to impose a tax on companies that automate so as to generate revenue for the support and re-skilling of displaced workers.

Section D of this article considers the viability of basing this automation tax on existing emissions pricing models, eventually rejecting such models after considering the considerable differences between them and the requirements of an automation tax (in section E)

Section F considers how an automation tax could be implemented. It proposes that the two main dimensions that may be adjusted to produce intended distortionary effects are: 1) accelerated depreciation, and 2) bonus depreciation. Finally, it notes that while the benefits of efficiency gains mean that the automation tax is unlikely to have widespread application, it does provide a useful tool for specific situations where the rate of automation needs to be slowed due to its resultant social costs.

B. Background

Technological advancement has always had a significant effect on employment and the economy.¹ The near-simultaneous introduction of clusters of related technologies with broad applications – in other words the appearance of a technological wave or revolution – has historically been particularly disruptive.² The recent explosion in progress in a closely-linked cluster of areas such as robot dexterity, machine learning, processing power, and sensor capabilities appears to herald a new technological wave with profound economic implications. Commentators have dubbed this the “Fourth Industrial Revolution”.³

The central notion of autonomy unites the cluster of technologies that comprise this technological revolution. The cost of automation has fallen dramatically, making automation more cost-effective for many more industries and companies, while the capabilities of automation technology have expanded on multiple fronts, including information collection, data processing, and physical action.⁴ To illustrate the magnitude and implications of these improvements, we will briefly explore three examples of the dramatic progress and impact of automation technology on the economy.

¹ Theodore Lewis, “Studying the impact of technology on work and jobs”, *Journal of Industrial Teacher Education*, (1996), 33(3), 44-65, 44-46.

² Joseph Bower and Clayton Christensen, "Disruptive Technologies: Catching the Wave", *Harvard Business Review* 73, (1995), 43–53.

³ Ben Vermeulen, Jan Kesselhut, Andreas Pyka and Pier Paolo Saviotti, “The Impact of Automation on Employment: Just the Usual Structural Change?”, *Sustainability*, (2018), 10(5), 1661.

⁴ Carl Frey and Michael Osborne, “The Future Of Employment: How Susceptible are Jobs to Computerisation?” (2013), 3; McKinsey Global Institute, “A Future That Works: Automation, Employment, And Productivity” (“**McKinsey**”) (January 2017), 6 (Exhibit E3); and Jonathan Tilley, “Automation, robotics, and the factory of the future”, *The great re-make: Manufacturing for modern times*, (2017), 67-71.

Manufacturing Robots

Significant progress has been made in the physical capabilities and dexterity of robots, such that they are presently able to perform tasks, such as fabric sewing, that could previously only be performed by humans. This progress has been enabled by improvements in sensor capabilities, which enable robots to have a greater awareness of their environment, improvements in processing power, which enable robots to perform the calculations required to execute complex tasks, and improvements in machine learning, which enable robots to identify the appropriate techniques for executing complex tasks.⁵

Artificial Intelligence

Algorithms have become dramatically better at identifying patterns and making judgments due to the greater availability of the data used as raw material for these algorithms, as well as an increase in processing power that has made it possible to process and interpret the vast quantity of available data.⁶ Algorithms are used in artificial intelligence programmes, which, due to their inherent speed, reliability, and scalability, now possess an advantage over humans in areas such as securities trading.⁷

The Internet of Things (“IoT”)

Sensors, processors, and networking capabilities are now sufficiently cheap and sufficiently miniaturised to be embedded into everyday objects, creating wide networks of interconnected objects that are able to independently collect, process and transmit information.⁸

C. The Case for an Automation Tax

Automation and Employment

Technological waves have historically created significant employment and economic opportunities, the positive impact of which has generally outweighed the disruptive impact of major technological change.⁹ Some positive effects of technological revolutions include greater aggregate economic output, a reduction in the need for menial labour, higher labour productivity and higher wages, as well as the creation of new job opportunities.¹⁰ However,

⁵ *ibid.*

⁶ Dave Cliff, “The impact of technology developments”, in *The Future of Computer Trading in Financial Markets*, Government Office for Science Working Paper (2011), 43-46; and McKinsey Global Institute, “Jobs Lost, Jobs Gained: Workforce Transitions In A Time Of Automation” (“**McKinsey 2**”) (December 2017), 24.

⁷ Cliff, (n 6), 49; and McKinsey 2 (n 6), 24.

⁸ McKinsey 2 (n 6), 24; and Felix Wortmann and Kristina Flüchter, “Internet of Things” *Business & Information Systems Engineering* (2015), 57(3), 221-224.

⁹ Martin Ford, *The Rise of the Robots: Technology and the Threat of a Jobless Future*, (2015) 27; and Daniel Smihula, “The waves of the technological innovations of the modern age and the present crisis as the end of the wave of the informational technological revolution”, *Studia politica Slovaca*, (2009), 32-47.

¹⁰ Sarah Kessler, “The Optimist’s Guide To The Robot Apocalypse” (9 March 2017) (available at <https://qz.com/904285/the-optimists-guide-to-the-robot-apocalypse>) (accessed on 28 September 2018); and Smihula (n 9), 31-42.

technological revolutions may also result in the mass displacement of human labour rendered obsolete by technological advancements.¹¹

In the short-term, automation is likely to alter existing jobs by catalysing changes in the scope and nature of a large proportion of existing occupations. Automation is able to promise better results for lower costs in the performance of specific tasks, giving firms a strong incentive to automate these tasks and subsequently redesign jobs such that workers complement technology by performing other tasks that cannot yet be practicably automated. These redesigned jobs will be constituted by a different mix of tasks: for instance, it is conceivable that tasks that are repetitive and menial will be replaced by tasks that require critical thinking or the management of interpersonal relationships.

Job Displacement as a Negative Externality of Automation

The negative effects of job alteration are likely to be significant. Some workers may find that their skills are no longer relevant due to the full automation of tasks associated with those skills; they may thus face structural unemployment.¹² The impact of automation is not evenly distributed: business owners may succeed in capturing the gains from automation instead of raising employment or wages, and highly-skilled workers that better complement automation may benefit more than the low-skilled workers whose jobs are easily automated.

The Severity of the Problem

While previous technological revolutions have also resulted in structural and frictional unemployment, it is likely that the present wave of automation will be more disruptive than before, for several reasons. Firstly, while previous technological innovations did not eliminate the need for human labour to operate and control technology, the autonomous nature of the present wave of technologies threatens to substitute human labour to a greater extent by fully eliminating the need for human intervention in the autonomous execution of a given task.¹³ Secondly, unlike previous technological innovations that have been limited in the scope of their applicability, autonomous technology, is a general-purpose technology with a much wider set of capabilities – ranging from physical action to information processing – and hence the potential for disruptive impact across a wider range of sectors.¹⁴ Finally, there is less time for governments to react to automation due to the rapid rate of progress in automation technology; Moore’s Law, for instance, predicts that the computing power of a microchip doubles every two years.¹⁵ The substitutability, scale, and speed of automation distinguish it from previous waves of technological innovation and strengthen the impetus for government intervention.

¹¹ Vivian Dong, “Policy Responses to Technological Unemployment” (30 September 2017), 1; and Frey and Osborne (n 4), 3.

¹² Frey and Osborne (n 4), 3; and Vermeulen et. al. (n 3), 6-7.

¹³ Ford (n 9), xvi-xviii, and 17.

¹⁴ *ibid*, p xvi-xviii, and 26-27.

¹⁵ Gordon Moore, “Cramming more components onto integrated circuits”, *Electronics*, (1965), 38(8), 114-117.

The Pattern of Job Displacement

Automation technology is currently still maturing. Recent breakthroughs or advancements have not all already translated into commercially-viable systems or equipment, and the adoption rate for automation technology will differ across firms and industries.¹⁶ Hence, the impact of automation at any given point in time is unlikely to be significant across the entire economy, but is instead likely to be disproportionately large for some sectors or job classes.¹⁷ To illustrate, sectors that are likely to be early adopters of automation include insurance and manufacturing, while job classes that are vulnerable to automation include clerical or administrative jobs.¹⁸

The concentrated nature of automation's effects means that worker displacement is likely to be limited to a few sectors or job classes at any point in time.¹⁹ However, there is a strong need for intervention in such sectors as displaced workers are likely to find it more difficult to transit to new jobs if most jobs in the same sector or job class are contemporaneously vulnerable to automation.²⁰ In other words, there is no need as yet for a radical overhaul of corporate or labour policy, but there is a strong need for policy measures targeted at mitigating the impact of automation on specific sectors or job classes.

In some sectors, technological advancements may make it feasible and economically compelling for companies to fully automate an entire class of jobs.²¹ A major problem ensues if the workers that perform these jobs are not typically required to possess skills that would allow them to perform alternative tasks or jobs. An archetypal example of this is the potential effect that self-driving trucks pose to truckers. Unlike other cases of automation like the introduction of automated teller machines, which were unable to perform all of the functions that human tellers performed, self-driving trucks could fully automate the roles of today's truck drivers. Furthermore, as the main skill of a truck driver is driving, these truckers are unlikely to possess skills that would allow them to perform alternative jobs or tasks, unlike bank tellers that would have been able to find employment in other customer service roles. Hence, automation in cases similar to that of self-driving trucks is likely to result in the long-term structural unemployment of a large class of workers. The need for intervention in such sectors is therefore particularly acute.

¹⁶ Ryan Kiggins, "The Political Economy of Robots Prospects for Prosperity and Peace in the Automated 21st Century" (2018), 62-63; and David Gruen, "The Future of Work", *Policy*, (2017), 33(3), 4-5.

¹⁷ Gruen (n 16), 6-7.

¹⁸ Frey and Osborne (n 4), 57-72.

¹⁹ Vermeulen et. al. (n 3), 13-15.

²⁰ Dong (n 11), 3; and Vermeulen et. al. (n 3), 13-15.

²¹ Brian Heater, "Technology is killing jobs, and only technology can save them" (27 March 2017) (available at <https://techcrunch.com/2017/03/26/technology-is-killing-jobs-and-only-technology-can-save-them/>) (accessed on 28 September 2018); and Vermeulen et. al. (n 3), 13-15.

An Automation Tax

The upshot is that automation, by inducing worker displacement, results in social costs arising from the need to support and retrain displaced workers. As these costs are borne by the worker or by society, instead of being borne by the company that makes the decision to automate, these social costs constitute a negative externality.²² The causal relationship between the acts of automation performed by firms and the resulting negative externality of job displacement creates a *prima facie* case for the state to intervene by discouraging or penalising these externality-generating acts of automation. Such an intervention could take the form of a Pigouvian tax,²³ which is imposed on an agent responsible for an externality for the purpose of mitigating that externality. In this case, such a tax would be an automation tax that would apply to all the technologies that make up the present wave of technological innovation – imposed on companies that automate through the deployment of automated systems or equipment in their production process.

However, the appropriate policy response is not to impose a blanket tax on automation, but instead to recognise the distinction between automation's employment-substituting²⁴ and employment-complementing²⁵ effects, so as to reward instances of the latter while disincentivising instances of the former. Such a policy response best exploits the potential of automation to raise productivity and generate employment opportunities, while limiting its potential for job displacement.

The aims of an automation tax are two-pronged: first, to slow the introduction of automation technology in these industries, so as to provide as much time as possible for governments, welfare systems, and workers to prepare for the impending effects of structural unemployment; second, to impose a tax on companies that automate so as to generate revenue for the support and re-skilling of displaced workers. Regulators must take care to only adopt these policy responses in a small number of cases where structural unemployment is widespread, irreversible, and clearly attributable to automation.

Reforming the Existing Tax System

The existing tax system is likely to be ill-suited for dealing with the challenges posed by “Fourth Industrial Revolution”. The existing tax systems of most developed nations give employers an incentive to make capital investments, but do not give similar incentives for employment.²⁶ This imbalance in incentives arises from the fact that tax deductions on

²² Nicolaus Tideman and Florenz Plassmann, “Pricing Externalities”, *Journal of Economic Literature*, (2010), 26(2), 176-184.

²³ Jonathan Masur and Eric Posner, “Toward a Pigovian State”, *Coase-Sandor Working Paper Series in Law and Economics*, (2015), 1-38.

²⁴ Vermeulen et. al. (n 3), 3-4.

²⁵ *ibid.*

²⁶ “Tax Incentives and Foreign Direct Investment: A Global Survey”, *United Nations Conference on Trade and Development ASIT Advisory Studies*, (“UNCTD”) (2000), (No. 16), 14-19.

investment in physical assets, such as capital allowances, are granted in many jurisdictions,²⁷ while payroll taxes are commonly imposed for every human employee. Historically, capital allowances and other incentives for investment have encouraged firms to make investments in capital that in turn raise the competitiveness and productivity of human labour, resulting in a boost to the economy's productive capacity while providing opportunities for workers.²⁸ Moving ahead, however, it is not clear if capital investment will continue to complement human labour in this way, as automation displays an increasing potential to substitute and displace workers.

Furthermore, while structurally unemployed workers could adapt by reskilling and finding opportunities in new industries, the rapid rate at which automation technology is improving²⁹ could leave workers with insufficient time to adjust. Given the increasing potential and rapid progress of modern technological innovation, tax systems should provide employers with an incentive to moderate the pace of displacement while continuing to encourage firms to make capital investments in ways that provide opportunities for workers. There is an urgent need to reform the tax system to shift from the present blanket incentive for investment, and instead move towards encourage labour-complementing investments while discourage labour-substituting investments.

D. Emissions Pricing as a Model for Automation Taxation

Theoretical Optimality of Pigouvian Taxation

We will explore the concept of emissions pricing as a model for automation taxation. Emissions pricing, or the imposition of a price on greenhouse gas emissions that is payable by emitters, is an example of a Pigouvian tax and has been adopted by several jurisdictions as a response to the negative externality of climate change arising from greenhouse gas emissions.³⁰ Analogously, a price could be imposed on the deployment of automated systems or equipment by firms, as a response to the negative externality of worker displacement caused by such automation.

Such a policy of pricing the externality by imposing a tax constitutes a market-based solution to the externality. The intended policy outcome of minimising the size of the externality is achieved not by prescribing or proscribing the actions of agents, but by adjusting the market prices associated with certain actions in order to fully reflect the externalities of these actions. While the intent is to influence agent behaviour by adjusting prices, these agents retain the ability to freely choose their actions and to generate the externality under market-based solutions. These market-based solutions have the advantage of achieving socially-optimal outcomes through the decentralised choices of independent and incentive-driven

²⁷ *ibid*, 19-22.

²⁸ Stanley Surrey, "Tax Incentives as a Device for Implementing Government Policy: A Comparison with Direct Government Expenditures" *Harvard Law Review*, 83(4) (Feb 1970), 705-738; UNCTD (n 26) 14-19.

²⁹ Vermeulen et. al. (n 3), 1-3.

³⁰ Masur and Posner (n 23), 1-6.

agents, eliminating the need for governments to undertake the challenging task of determining and dictating what the socially-optimal behaviour would be for each individual agent.³¹

Implementation Approaches

Regulators must determine the size of the externality generated by each agent, so as to tax it and cancel out the externality. There are two possible ways of doing so. The first approach is for regulators to directly estimate the social costs arising from the actions of each individual agent.³² It is often difficult to identify or attribute the individual causal contribution of each agent to the social costs that are collectively generated.³³

The alternative to directly observing the size of the externality generated is to infer its size from observations of the intensity or extent of the agent's externality-generating actions. This is usually achieved in the following way. First, the intensity or extent of the agent's externality-generating actions is quantified and measured in terms of a chosen unit of taxation, and then used as a tax base. Next, a specified tax rate is imposed on each unit of the tax base, such that the overall tax imposed on each agent is equivalent to the tax rate per unit multiplied by the tax base as measured in the same units. This overall tax payable will be a close approximation of the size of the externality generated by the agent, and hence will be economically optimal, if the tax meets the following conditions.

First, the agent's externality-generating actions must be quantified in terms of a single, standardised unit that applies to all externality-generating agents. This common unit of quantification allows the tax base to be computed for each agent. Second, it must be feasible to accurately measure the intensity or extent of the agent's externality-generating actions in terms of the specified unit of taxation. This allows the tax base to be determined for each agent, hence enabling the practical enforcement and administration of the tax. Finally, the size of the tax base, or the intensity or extent of the agent's externality-generating actions as measured in terms of the specified unit of taxation, must be proportional to the size of the externality. This will allow regulators to impose a constant amount of tax for each unit of the externality-generating action.

These conditions are met in the case of the EU's Emissions Trading System ("**ETS**"),³⁴ which imposes a tax on each company based on its contribution to climate change, which is quantified in terms of the total potential warming effects of the greenhouse gases it emits.

³¹ *ibid*; and Joseph Petrucelli and Jonathan Peters, *Preventing Fraud and Mismanagement in Government: Systems and Structures* (2016), 87-88.

³² Joel Slemrod and Christian Gillitzer, "Optimal Observability and Complexity" *Tax Systems* (2014), 145-156.

³³ Warren Samuels, *Economics, Governance and Law: Essays on Theory and Policy* (2002), 160.

³⁴ Denny Ellerman, Claudio Marcantonini and Aleksandar Zaklan, "The European Union Emissions Trading System: Ten Years and Counting" *Review of Environmental Economics and Policy*, 10(1), (2016), 89-107; and Bram Borkent, Alyssa Gilbert, Erik Klaassen, Maarten Neelis and Kornelis Blok, "Dynamic allocation for the EU Emissions Trading System Final Report", (2014).

The unit of quantification used under the ETS is typically MMTCDE, or million metric tonnes of carbon dioxide equivalents.³⁵ The tax base can be measured in terms of these units using a three-step procedure: first, the amount in million metric tonnes (“MMT”) of each greenhouse gas emitted by an agent is measured; second, to account for differences in the environmental effects of different greenhouse gases, the amount in MMT emitted for each greenhouse gas is converted into MMTCDE using the associated global-warming potential (“GWP”) for that gas; third, the total size of the tax base is obtained by summing the amount emitted in MMTCDE across all greenhouse gases.³⁶ MMTCDEs are a suitable unit of quantification, as they provide a single measure of the total effect of the many different types of greenhouse gases that could be emitted, and as such can serve as a common unit of quantification across companies that emit different types of gases.³⁷ Hence, the first condition of a quantifiable tax base is met in the case of the EU ETS.

The second condition, which is the practicability of measuring the tax base, is also fulfilled in the case of the EU ETS. The amount of greenhouse gases emitted by each company can be easily and accurately measured by detectors installed on its properties, facilitating the computation of the MMTCDE emitted by each company.³⁸

Finally, the EU ETS meets the condition of having a tax base that is proportional to the size of the externality. Since the warming effect for each MMTCDE emitted is a scientifically-knowable physical fact, the total contribution of a company’s emissions to climate change can be obtained simply by multiplying this warming effect by the amount emitted in terms of MMTCDEs. Hence, the final condition is fulfilled, as the tax base as measured in terms of MMTCDEs is proportional to the externality of climate change generated by each company.

The ability of the EU ETS to meet the conditions for an efficient Pigouvian tax demonstrates its potential to effectively address negative externalities by improving both allocative and distributive efficiency. This makes emissions pricing a promising model, at least in theory, for tackling the negative externality of job displacement that is associated with automation. Through the allocative mechanism, the automation tax could reduce or slow the displacement of workers by automation. Through the distributive mechanism, the automation tax could generate revenue to support and re-skill workers that are displaced by automation.

While emissions pricing may be a promising model for an automation tax in theory, an automation tax adopted according to the model of emissions pricing would suffer from significant practical and theoretical issues. In the following sections, we will explore first the

³⁵ Mark Miner, “Neutral Energy Consulting” (5 September 2011), (available at <http://www.neutralenergy.info/2011/06/pricing-carbon.html#EU>) (accessed on 28 September 2018); and Directive 2003/87/EC of the European Parliament and of the Council of 13 October 2003 establishing a scheme for greenhouse gas emission allowance trading within the Community and amending Council Directive 96/61/EC, Article 3(j).

³⁶ *ibid.*

³⁷ *ibid.*, Article 6.

³⁸ Mirabelle Muûls, Jonathan Colmer, Ralf Martin and Ulrich Wagner, “Evaluating the EU Emissions Trading System: Take it or leave it? An assessment of the data after ten years” (October 2016) Grantham Institute Briefing Paper No. 21, 1-12; and “The Monitoring and Reporting Regulation – Guidance on Sampling and Analysis MRR Guidance document No. 5”, *European Commission*, (2017), 1-35.

practical design and implementation challenges faced by an automation tax, before exploring the theoretical flaws that undermine its economic optimality.

E. Design and Implementation Challenges for an Automation Tax

Automation taxes could be designed and implemented based on the emissions pricing model. Under such a model, the actual or potential employment-displacing effects of automation would serve as the base for an automation tax, analogous to emissions pricing systems which use the actual or potential warming effects of greenhouse gases as the tax base. Subsequently, companies would pay a constant amount of tax for each unit of the tax base.

The key design and implementation challenge for an automation tax that is based on the emissions pricing model is the difficulty in identifying an appropriate tax base. An appropriate tax base should meet the three above-mentioned conditions of quantifiability, measurability and proportionality: it should be quantifiable in terms of a common unit, it must be feasible in practice to measure the size of the tax base in terms of this common unit, and the tax base as quantified in terms of this unit should be proportional to the size of the externality. If these conditions are not fulfilled, the automation tax may not be practically enforceable. It may also fail to be theoretically efficient, as the size of the tax payable by each agent would not be proportional to the size of the externality generated by that agent.

We will explore if the possible tax bases for an automation tax meet all of the above three conditions. We will do so by grouping the wide range of possible tax bases into two categories. The first category is outcome-related tax bases, under which companies are taxed according to the actual outcome, in terms of employment or job displacement, of the automation that they have implemented. Tax bases in this first category seek to quantify the actual employment-displacing effects of automation.³⁹ The second category is that of action-related tax bases, under which companies are taxed according to the extent or type of the automation that they have implemented – in other words, the intensity of automation. Tax bases in this second category seek to quantify the potential employment-displacing effects of automation based on the intensity of automation implemented.⁴⁰

Overall, we will argue that no tax base in either category meets all three of the conditions of quantifiability, measurability, and proportionality. The lack of an appropriate tax base from either category poses a significant design and implementation challenge for an automation tax based on the emissions pricing model.

³⁹ Xavier Oberson, “Taxing Robots? From the Emergence of an Electronic Ability to Pay to a Tax on Robots or the Use of Robots”, *World Tax Journal*, (2017), 254-255.

⁴⁰ *ibid.*, 256-257.

Outcome-Related Tax Bases

The most direct approach to determining the size of the externality is to determine the number of workers retrenched due to the implementation of automation within a firm. Abbott and Bogenschneider argue that this approach is conceptually similar to existing systems where employers are taxed, in the form of payments into unemployment insurance schemes, based on the number of workers that have been retrenched from their firms.⁴¹ Some modifications could suffice to transform this existing system into an automation tax. For instance, tax authorities could determine the extent to which each firm's layoffs can be attributed to automation, and then accordingly adjust the amount each firm contributes in unemployment insurance payments.

Tax bases in this outcome-related category could, in theory, meet the three conditions of quantifiability, measurability and proportionality. The tax base could be quantified in terms of either the number of layoffs, or the total monetary extent of any reduction in worker wages. It would also be easily measurable based on the financial data, payrolls, or other internal records of companies. If layoffs could be accurately attributed to automation instead of other causes, the condition of proportionality would also be met; the size of the tax base, as measured in terms of layoffs attributable to automation, would be proportional to the extent of externality of automation-induced job displacement.

The main challenge for outcome-related tax bases lies in the difficulty of determining if layoffs should be attributed to automation instead of other possible causes such as poor business conditions or productivity improvements unrelated to automation. One potential approach for establishing a prima facie causal relationship between automation and layoffs is to use a multi-pronged test to determine if the displaced employee was substituted with automation. This test could include elements such as whether the period of time between automation and the layoff was sufficiently short and whether the tasks automated were sufficiently similar to the tasks performed by the displaced employee. The presence of all elements in this test could form the basis for a presumption that the layoffs in a given case were attributable to automation.

The interpretation and application of such a test, however, would pose significant practical challenges due to the level of technical expertise and industry knowledge required to determine, for instance, whether the tasks automated were sufficiently similar to that performed by the displaced employee. Given the widespread extent of automation, it is unlikely that tax authorities would possess the requisite expert capabilities to an extent sufficient to cope with the likely volume of cases. More importantly, the presumptive nature of this approach runs the risk of over-estimating the extent of layoffs attributable to automation. Should other layoff-inducing factors such as a downturn or productivity improvements coincide with the implementation of automation, the layoffs caused by these other factors may mistakenly be attributed to the impact of automation. This problem may be addressed by removing the presumption when other layoff-inducing factors are present, but this opens the possibility for firms to disguise the employment-displacing effects of automation by implementing such

⁴¹ Ryan Abbott and Bret Bogenschneider, "Should Robots Pay Taxes? Tax Policy in the Age of Automation", *Harvard Law & Policy Review*, (2018), 170.

automation in a manner that coincides with a downturn or with the implementation of other productivity improvements. Overall, the approach of using such a test may work well in a limited number of situations, such as the replacement of truck drivers with self-driving trucks, where the relationship between automation and employee displacement is evident. In the overwhelming majority of cases, however, where automation bears a more complex relationship with employment, this test runs the risk of being either too strict or too lenient.

An alternative to determining the number of workers retrenched due to automation is to use capital, labour or profit ratios as the tax base for the automation tax. This alternative falls within the outcome-related category of tax bases because these capital, labour and profit ratios are observable outcomes of a firm's employment and automation decisions, and because these ratios could be indicative of the employment-displacing effect of automation implemented within a firm. To illustrate, high capital-profit or capital-revenue ratios indicate that a firm uses significant amounts of capital, which could come at the expense of employing labour. Conversely, low labour-profit or labour-revenue ratios indicate that a firm uses relatively little labour to generate sales or profits, which might imply the substitution of capital for labour. To determine if a firm has capital ratios that are too high or labour ratios that are too low, tax authorities could compare the firm's present capital or labour ratios either to its own capital or labour ratios in previous periods, or to the capital or labour ratios of other firms in the same sector.

Both alternatives have their drawbacks. The former method of comparing a firm's present capital and labour ratios to its own past ratios does not account for inherent differences in the capital-intensiveness or profitability of different industries. For instance, automobile manufacturing is an inherently more capital-intensive industry than hospitality and tourism, while firms in the information technology sector tend to enjoy significantly larger profit margins than firms in retail. Consequently, it gives insufficient credit to firms that have consistently created more employment opportunities than other companies in the same industry. In contrast, the latter method of comparing the firm's capital and labour ratios to that of its peers in the same industry does not account for the way firms have changed over time.

Overall, it is not clear that outcome-related tax bases, while promising in their attempt to directly determine the size of the externality of automation-induced job displacement, can overcome the crucial challenge of determining if observed employment outcomes should be attributed to automation or to other layoff-inducing factors. Crucially, the relationship between observed employment outcomes and automation is too complex and contestable, making it difficult for tax authorities to make judgments on attributability that are both sufficiently efficient to cope with the high volume of cases and sufficiently accurate to achieve the desired economic effects of the automation tax.

In other words, the condition of proportionality can only be fulfilled, if at all, at the expense of the condition of measurability. The upshot is that none of the possible outcome-related tax bases can meet all three of the conditions for the optimal design of an automation tax.

Action-Related Tax Bases

Instead of relying on observable employment outcomes, tax authorities could use the extent or type of automation implemented, in other words the overall intensity of automation, as a proxy for the employment-displacing impact of automation. For this approach to succeed, the same three conditions must be met: what is meant by the “intensity of automation” would have to be specified in terms of quantifiable units, it must be feasible to measure the intensity of automation in terms of these quantifiable units, and the intensity of automation as measured using this quantifiable unit should be proportional to the employment-displacing impact of automation, and by extension the size of the externality.

The first condition is best fulfilled by using the value added by automation as the quantifier for the intensity of automation. Since the intensity of automation is intended to reflect the extent to which the firm relies on automation, as opposed to labour, to perform value-generating tasks, it makes sense for the intensity of automation to be quantified in a way that captures the total sum of such tasks that are performed by the firm, or the total amount of work performed using automation. When summing these tasks, each task should not carry equal weight, as tasks vary in terms of their importance to the firm. Instead, the sum should be weighted based on the value of each task. This sum of the value of automated tasks is conceptually equivalent to the total value added by automation. Furthermore, if it holds true that automated tasks could have been performed equally well by human labour, then the value added by automation in performing these tasks is indicative of the value of the human labour that could have been employed in place of automation, and therefore further indicative of the employment-substituting effect of automation.

Indeed, using the value added by automation as a quantifier is superior to alternative means of quantifying the intensity of automation. The cost of automation equipment, for instance, is a poor proxy for the intensity of automation. One reason is that the cost of the non-autonomous components of automation equipment may contribute disproportionately to the overall cost of the equipment. For instance, a self-driving vehicle may only be slightly more expensive than a comparable conventional vehicle, because the cost of its autonomous components and software makes up only a fraction of the car’s total cost. In this case, it seems inappropriate that the entire cost of the car should be taxable. Even if the tax base included only the cost of the automation-related components of equipment, this cost is a poor reflection of the extent to which tasks have been automated by this piece of equipment. Expensive equipment may not necessarily have greater capabilities or contribute to the automation of more tasks; indeed, while physical capital tends to have a constant and significant cost per marginal unit of output, the scalability and low marginal cost of automation technologies such as algorithms means that the impact of these technologies is decoupled from their cost.⁴² The upshot is that automation taxes ought not to take a form similar to that of a sales tax, since that would entail using the cost of automation as a tax base.

⁴² Paul Romer, *Capital, Labour and Productivity* (1990), 337-367, 339; and Tilley (n 4), 67-72.

Another alternative way to quantify the intensity of automation is to determine the potential capabilities of automation systems and equipment purchased by a firm. Should a firm invest in automation systems and equipment that have significant capabilities or that are able to perform a significant proportion of the tasks involved in the production process, it is reasonable to assume that they are automating to an intensive degree. To facilitate comparisons between systems and equipment with different capabilities, technical experts could estimate the amount of human labour that could be replaced by automation systems or equipment with a particular feature or a particular category of automation system or equipment. The tax base could then be computed by adding up the estimated quantity of human labour that could be replaced by each individual system or piece of equipment, so as to obtain the total amount of human labour that could be replaced by the potential capabilities of the entirety of a firm's automation systems or equipment. However, this capability-based approach of quantifying the intensity of automation suffers from a problem similar to that of the cost-based approach discussed above. Automation systems and equipment with greater capabilities have greater potential for performing tasks, but the actual scale of the tasks performed by automation may not be as significant as the full potential of such automation. Again, the impact of automation depends not only on the intrinsic capabilities of the equipment or systems used, but also on the size of demand for the goods and services that are produced through automation. In the context of automation, the capabilities of equipment or systems give a sense of what types of tasks they can perform, but may not necessarily shed light on the scale or significance of these tasks. The value added by automation remains superior to alternatives, such as the cost or capabilities of automation systems and equipment, for quantifying the intensity of automation.

Having shown that the value added by automation is the best available means of quantifying the intensity of automation, we will now examine if the value added by automation is easily measurable and observable. Here, we run into a problem similar to that faced by outcome-related tax bases: the problem of attributing outcomes, such as employment outcomes in the previous case or value added in this case, to automation as opposed to other factors that might also be causally related to these outcomes. While it is possible to determine the change in the profitability or output of a firm following an instance of automation, it is incorrect to assume that this change in value added is entirely attributable to automation. Any observed change might have been caused, for instance, by process changes that were unrelated to and introduced at the same time as the implementation of automation. Furthermore, the value added by automation depends significantly on the synergies between automation and other inputs or changes such as capital, skilled labour or improvements to business processes. The presence of these synergies precludes a neat decomposition or partition of the total output of a firm into the separate contributions or value added by individual inputs such as automation and labour. The upshot is that the value added by automation is in practice difficult to measure or determine.

The value added by automation fails not only the second condition of being measurable but also the third condition of being strongly correlated with the employment-displacing effects of such automation. The key assumption that is required for this final condition to be fulfilled is that the value added by automation could have been value added by labour. This assumption does not hold in all instances of automation. For example, the value created by Internet search

engine algorithms is value added by automation that could not have been achieved with the use of human labour alone, due to the unfeasibility of making humans search manually through millions of webpages. Similarly, the value added by automated precision manufacturing tools lies in the ability to achieve a level of accuracy and reliability that would be impossible with human labour. Including these instances of value added by automation in the tax base for an automation tax has two implications: first, the desired outcome of ameliorating worker displacement is not achieved, as human labour cannot be used as a substitute for automation in these cases; second, as these forms of automation provide benefits to society that could not otherwise be achieved, there is an economic loss to society as the use or development of such automation is penalised.

To summarise the discussion above, the intensity of automation meets only one of the three conditions that are necessary for it to serve as an appropriate tax base for an automation tax. While it can be quantified in terms of the value added by automation, this value added cannot easily be measured and does not correlate well with the size of the externality, making it a poor tax base both in terms of practical enforcement as well as in terms of theoretical efficiency.

Theoretical Challenges

An automation tax, unlike an emissions pricing regime, may not be theoretically optimal due to the adverse social and economic impacts of such a tax. These adverse impacts are especially significant in open economies that are highly exposed to global competition in technology, trade and investment.

The Cost to Technological Competitiveness

The first of these adverse impacts is that an automation tax threatens to undermine the economy's technological competitiveness. The rapid growth and potential size of the technology sector has resulted in intense global competition for leadership in the development and production of various automation-related technologies, such as artificial intelligence and robotics.⁴³ The importance of rapidly assuming leadership in automation development and production is magnified by two facts: first, technological dominance can translate into broader economic dominance, as automation-related goods and solutions are likely to be widely embedded in the production and consumption of goods and services across all sectors;⁴⁴ second, players that assume early leadership in the automation sector have the opportunity to entrench their dominant positions with their outsize influence over the development of global technology standards.⁴⁵

⁴³ Emanuel Gasteiger and Klaus Prettnner, "A Note on Automation, Stagnation, and the Implications of a Robot Tax" (2017) School of Business & Economics Discussion Paper, 2; and Jan Fagerberg, "Technology and Competitiveness" (1996) *Oxford Review of Economic Policy*, 12(3), 48-51.

⁴⁴ Richard Heeks and Carolyne Stanforth "Technological change in developing countries: opening the black box of using actor-network theory", *Development Studies Research* (2015), 2; and Fagerberg (n 43), 54-56.

⁴⁵ Fagerberg (n 43), 58.

An automation tax could undermine an economy's efforts at assuming this leadership position and strengthening its technology sector in two ways. Firstly, the technology sector is likely to be disproportionately affected by an automation tax, due to the high degree of automation technology used in the development and production of automation technology itself. Secondly, even if the technology sector is exempted from an automation tax, such an automation tax would result in the loss of industrial users of automation technology and the loss of labour with automation-related skills in other sectors of the economy. As a critical mass of both skilled talent and industrial partners is crucial for nurturing a budding automation production and development industry, an automation tax that obstructs the formation of this critical mass also poses an obstacle to the development of the economy's technology sector.

Beyond the longer-term effect on the growth and development of technology sector, however, there are more immediate costs to the wider economy. Firstly, the allocative effect of an automation tax will undermine the economy's trade competitiveness in non-technology sectors. By raising the productivity of the production process and reducing labour costs, automation contributes to a reduced cost of production for goods of similar type and quality, making these goods more competitive in global markets. If companies were to reduce their use of automation due to the allocative effects an automation tax, the economy would forgo the export-boosting benefits of automation. These benefits are particularly large for the developed economies that are the focus of this article: their greater ability to adopt automation technology allows them to compensate for their higher labour costs by raising labour productivity and reducing production costs, allowing them to retain trade competitiveness and by extension employment opportunities in automation-intensive sectors.⁴⁶

Secondly, the distributive effect of an automation tax will undermine the economy's competitiveness for investment and production in non-technology sectors. As globalisation has made it possible for large multi-national corporations ("MNCs") to shift production and investment to economies with the lowest costs and greatest returns, competition among economies for the limited global pie of investment and production is intensifying.⁴⁷ If the potential returns of MNCs investing in a given economy were to be reduced due to the distributive effects of an automation tax, its competitors would become relatively more attractive as destinations for investment and production. All things equal, this would shrink the economic output of as well as the medium-term supply of employment opportunities available within that economy.

These adverse impacts on the output of and supply of job opportunities in non-technology sectors are likely to be particularly acute in automation-intensive sectors – sectors in which the use of automation contributes significantly to output and productivity, and in which companies are likely to make investment and production decisions based on an economy's openness to the use of automation. The adverse economic impact of an automation tax is likely to be magnified by the importance of these automation-intensive sectors to developed economies as these are likely to be sectors that are high-value and highly productive,

⁴⁶ Abbott and Bogenschneider (n 41), 21.

⁴⁷ *ibid*, 32; and "Globalization, Competition, Competitiveness and Development", *United Nations Conference on Trade and Development*, (1997), 2.

in addition to being the sectors in which developed economies are likely to have a comparative advantage.

More significantly, this fact that the negative effects of an automation tax are likely to be concentrated in automation-intensive sectors implies that an automation tax could be highly counterproductive. The intent of an automation tax, after all, is to mitigate the effect of automation in displacing workers, especially in automation-intensive industries. While such an automation may serve as a direct remedy by slowing the pace of automation or by raising funds to support affected workers, it may also indirectly exacerbate the underlying problem by undermining the price competitiveness of their output and the creation of job opportunities in their sectors. In other words, automation is a double-edged sword for workers in automation-intensive industries: it threatens to displace their labour but compensates by safeguarding them from external competition for their output and their jobs. Automation can be both employment-substituting and employment-complementing; in attempting to address the costs of the former, an automation tax may force society to forego the sizeable benefits of the latter.

Differences Between the Cases of Automation and Emissions

There are two crucial differences between the economic benefits of automation and those of greenhouse gas emissions. The first difference lies in the parties to which these benefits accrue. Since the economic benefits of greenhouse gas emissions are fully captured by private agents, the imposition of a Pigouvian tax achieves a socially optimal outcome by compelling these agents to consider the social costs of emissions in addition to their private benefits. In contrast, the economic benefits of automation extend beyond the agent to society in general, and hence are not considered by agents in making decisions on whether to automate. This results in a suboptimal degree of automation in society – a problem exacerbated by an automation tax. The second difference is that an automation tax, unlike an emissions pricing scheme, has the potential to be counterproductive because the employment-substituting effects of automation are directly opposed to its employment-complementing effects. In other words, while it is certain that an emissions pricing scheme will at least to some extent accomplish its objective of mitigating climate change, it is unclear if the net effect of an automation tax will be to decrease or increase employment in automation-intensive industries.

Hence, unlike emissions pricing schemes that are theoretically effective as a means of tackling climate change, an automation tax could suffer from serious theoretical flaws. An automation tax adopted according to the emissions pricing model assumes that automation will inevitably result in the displacement of employment, and therefore generate an externality that must be remedied through a Pigouvian tax. As our discussion above shows, however, automation is not purely a substitute for labour. Like other forms of capital or other productivity improvements, it may have employment-complementing or employment-substituting effects. This is unlike the case of greenhouse gases, which have an unambiguously negative effect on the environment.

F. Proposal: Reverse Depreciation

We propose that any reforms to the corporate tax system be made to the existing system of depreciation/ capital allowances. Instead of a blanket deduction for capital investment, we suggest that the deductibility of capital investments should vary depending on the effect of the capital investment on employment. More specifically, companies that invest heavily in employment-complementing capital will be allowed to deduct a greater proportion of their capital expenditure from their taxable income, while companies that invest more heavily in employment-substituting capital will only be allowed to deduct a smaller proportion of their capital expenditure. We have chosen this approach for several reasons.

This reform directly tackles the need to rebalance existing capital investment incentives for companies. The existing capital allowances system gives companies a blanket incentive to invest in capital, regardless of whether this capital is employment-substituting or employment-complementing. While such a blanket deduction may have been appropriate in a context where new and better employment opportunities were continuously being generated, it is no longer appropriate in the present context where the rate of automation-induced job displacement is likely to exceed the rate at which new job opportunities are created. By targeting capital allowances, this reform ensures that

This reform further tackles the need to expand the revenue base so that the government has sufficient funds to support the growing number of displaced workers. By reducing the deductibility of capital expenditure for a subset of companies, the government can raise more corporate tax revenue without an increase in the headline rates of existing taxes.

Overall, reforming the existing system of capital allowances corrects the imbalances of the existing system and also expands the revenue base to address the increased need for government support of displaced workers. We will now suggest several reasons for why targeting capital allowances is superior to alternative means of achieving these two goals of rebalancing incentives and raising revenue.

First, as the additional revenue raised by our proposed reform is generated by reducing the tax deductibility of some capital investments, the additional tax burden is imposed only on profit-making companies. In contrast, alternatives such as an automation tax or a tax on the capital-profit ratio impose an additional tax burden on both profit-making and loss-making companies. The effect of the latter is to increase the pressure on struggling companies and make it more difficult for them to invest in automation as a means of staying afloat. While this may achieve the intended effect of reducing overall levels of investment into employment-substituting capital, it has the unintended and more significant effect of reducing the competitiveness of local firms vis-à-vis their international competitors. In turn, this may result in greater job losses as uncompetitive domestic firms are forced to shut down or relocate in the face of external competition. This problem can be avoided by ensuring that only profit-making firms – firms that are able to generate profits even in the face of competition – are subject to the increased tax burden. Revising the existing system of capital allowances accomplishes this goal.

Second, the enforceability of this proposal is superior to that of imposing a completely new tax. While companies can avoid a direct tax on automation by concealing their investments into automation, there is no similar means of avoidance for a policy that removes the tax deductibility of some capital investments. Furthermore, there is already a well-established enforcement and administrative mechanism for the existing system of capital allowances, as well as a significant body of case law. Piggybacking on this solid foundation avoids the inefficiency and disruption that would result from the establishment of an entirely new enforcement and administrative mechanism for a new tax.

Finally, this proposal contributes to increasing the overall productivity of the economy, which is necessary to allow domestic industries to face the threat of external and low-cost competition. It continues to preserve the deductibility of capital expenditure if such expenditure complements labour, so companies continue to have an incentive to raise worker productivity by investing in capital. Furthermore, unlike a blanket excise tax on technology or capital or a tax on companies with high capital-profit ratios, it does not distort the market by placing a greater tax burden on companies in capital-intensive industries. Instead, companies in these industries can continue to enjoy the existing tax benefits of capital investment, so long as they ensure that their capital investments are enhance but do not displace existing job opportunities.

Our proposal of making adjustments to the deductibility of capital expenditure in order to achieve a policy objective – in this case the attenuation of the effects of automation on employment – is not without precedent. Depreciation rates have been accelerated to stimulate investment during recessions,⁴⁸ and bonus capital allowances were granted to companies in the Job Creation and Worker Assistance Act of 2002 enacted by the US Congress in the wake of the September 11 attacks.⁴⁹ These precedents suggest that it is feasible and legally justifiable to use capital allowances as a policy instrument to achieve economic goals.

A Complex Problem

The function of the automation tax is as a policy tool for the Government to control the rate at which automation displaces human workers. Given that automation is doing so in a myriad of different ways and in unpredictable and perhaps unimaginable ways, the ideal robot tax would have to be complex enough to apply in all these different ways, and flexible enough to keep up with the rapid developments in automation technology. The demand for a system that works in precisely such a way is not new. Tax authorities around the world have long realized the inadequacies of general tax rules;⁵⁰ draft them too broadly and there is no taxpayer certainty;

⁴⁸ Christopher House and Matthew Shapiro, “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation” *American Economic Review*, (2008), 98:3, 737-738; and John Kitchen and Matthew Knittel, “Business Use of Section 179 Expensing and Bonus Depreciation, 2002-2014” (October 2016), Office of Tax Analysis Working Paper 110, 5.

⁴⁹ Job Creation and Worker Assistance Act of 2002 (Pub. L. 107-147, 116 Stat. 21); and Kitchen and Knittel (n 48), 5.

⁵⁰ Victor Thuronyi, “Tax Law Design and Drafting”, *International Monetary Fund*, (1996), vol 1, chapter 3; and Surrey (n 28).

draft them too narrowly and risk the loss of tax revenue.⁵¹ The solution lay in the use of comprehensive schedules supported by general tax principles; each item on the schedule could be accorded different tax treatment depending on the governmental policy at the time.⁵² Such a system was applied to the tax treatment of assets acquired for use in businesses. The use of such systems diverged over time, with the United States labelling the concept “depreciation of assets” and the United Kingdom calling it “capital allowances”. This section explores the mechanism of “depreciation”, which has been used to provide tax incentives and disincentives arguably since 1878.⁵³ It will argue that the schedular system used by depreciation is extraordinarily well suited for controlling the rate at which automation displaces human workers.

The Concept of Depreciation

The fundamental concept of depreciation is simple enough to grasp. A firm which invests in new capital for its business incurs an expense. However, to allow the firm to deduct that expense from its income for the year in full would be too generous. The firm still possesses the asset, which will continue to have value until the end of its working-life. Thus, most tax authorities will not allow the firm to deduct the full cost of the asset as an expense immediately. Instead, the firm must deduct the expense on a periodic basis, generally mirroring the gradual decrease in value of the asset until it becomes worthless.⁵⁴

The United States (Depreciation)

In the United States, depreciation follows the specifications of the Modified Accelerated Cost Recovery System (“**MACRS**”), which contains comprehensive schedules of the rates at which specified assets or classes of assets may be depreciated at.⁵⁵ While there are multiple methods for calculating the rates of depreciation under MACRS, the most common method is the “declining balance method”,⁵⁶ which allows for faster depreciation when the asset is initially purchased, with the rate of depreciation relative to the cost price of the asset declining over time.⁵⁷

The United Kingdom (Capital Allowances)

The United Kingdom concept of capital allowances has a major conceptual difference from that of depreciation, the adoption of which it has continuously rejected.⁵⁸ Strictly speaking, the English tax system draws a distinction between revenue expenditure and capital expenditure. While revenue expenditure may be deductible in computing the profits of a trade, capital

⁵¹ *ibid.*

⁵² *ibid.*; and Slemrod and Gillitzer (n 32), 7.

⁵³ Dominic de Cogan, “Purposive Interpretation in the Age of Horse Trams”, (2015) 1 BTR 80, 81.

⁵⁴ House and Shapiro (n 48), 744.

⁵⁵ *ibid.*, 744-745.

⁵⁶ Douglas Kahn, “A Proposed Replacement of the Tax Expenditure Concept and a Different Perspective on Accelerated Depreciation”, 41 Fla. St. U. L. Rev. 143 (2013-2014), 153.

⁵⁷ House and Shapiro (n 48), 744-745.

⁵⁸ Glen Loutzenhiser, *Tiley’s Revenue Law* (8th Ed, 2016), para 6.1.3.

expenditure is not. It has been suggested that the reason why the English tax system did not initially have a system for recognizing capital expenditure is because income tax was thought to be only temporary.⁵⁹ An outright refusal to recognize the gradual depreciation in the value of capital assets renders the calculation of the true amount of profits earned from trade inaccurate. Expenditures of capital assets are a cost of doing business in much the same way as revenue expenditures are.⁶⁰

Another effect of disallowing the deduction of capital expenditure is that it has a distortionary effect on the market, disadvantaging capital-intensive businesses. Yet there is no good reason why capital-intensive businesses should suffer as such. The great importance of such capital-intensive businesses eventually led to the implementation of various schemes to recognize the true costs of capital expenditure.⁶¹ The most extensive and established scheme is the capital allowance system,⁶² which broadly functions in a similar manner to the MACRS in the United States.

A Scheduling System

Both MACRS and the capital allowance system are highly prescriptive, relying heavily on comprehensive schedules of assets which provide for differing tax treatment of the various assets or classes of assets. This is done to keep the system reasonably simple,⁶³ for it would be a nearly impossible task to attempt to estimate the exact depreciation in value of each and every asset being claimed by taxpayers. The scheduling system has provided particular advantages in terms of flexibility of economic policy. Where the Government wished to introduce a particular adjustment to the rate of depreciation, instead of changing the general taxing provision or specially enacting a separate section, it was able to simply change the rate for a particular asset or class of assets in the schedule. In this way, a very flexible system could be created where the Government could provide incentives for very specific classes of capital assets depending on what economic policy required.

Two Dimensions of Adjustments

While the primary function of MACRS is to reflect the true costs of doing business, it has been used on multiple occasions to provide firms with an incentive to engage in certain activities. The prescriptive nature of the system, with its extensive schedules make it possible to single out particular activities for special tax treatment. There are two main dimensions by which the

⁵⁹ *ibid*, para 6.1.1.

⁶⁰ *ibid*, para 6.12.1.

⁶¹ Michael Sherry, *Whiteman & Sherry on Income Tax*, (4th Ed) (2016), para 9.001.

⁶² Loutzenhiser (n 58), para 6.1.1.

⁶³ Royal Commission on Taxation, *Report of the Royal Commission on Taxation, Volume 4: Taxation of Income (Continued)*, (1966), 237-238.

standard declining balance method may be adjusted to produce intended distortionary effects: 1) accelerated depreciation, and 2) bonus depreciation.

Accelerated Depreciation

“Accelerated depreciation is the allowance of deductions for declines in the value of an asset at higher rates than are expected to occur in practice.” Conceptually, the total amount of tax deductions attributable to the capital expenditure does not change. Instead, the deductions are brought forward so that they can be made earlier.⁶⁴ In effect, the taxpayer receives an interest-free loan from the Government, equivalent to the amount of tax deferred as a result of the early deduction.⁶⁵ Apart from the interest-free loan from the Government, accelerated depreciation also offers firms several other benefits. First, as the value of money decreases over time due to inflation, the ability to defer one’s taxes raises the net present value of the capital asset,⁶⁶ since the deduction is claimed in present dollars rather than in future dollars (which are worth less).⁶⁷

Second, while future tax deductions are uncertain since they may be affected by a variety of unexpected factors, claiming the deductions immediately locks in the effect of the tax deductions, reducing the risk for the business.⁶⁸ Third, the early deduction of capital expenditure provides cash flow benefits, giving the firm more liquid cash and allowing the asset to breakeven at a faster rate.⁶⁹ Finally, accelerated depreciation provide an important source of funds to firms and reduces the need to obtain external financing.⁷⁰ Presently both the US MACRS and the UK capital allowances system have default depreciation rates that are accelerated.⁷¹ Nevertheless, the depreciation rates are frequently further accelerated to achieve economic objectives, particularly to stimulate capital investment during times of recession.⁷²

Bonus Depreciation

For prescribed categories of capital expenditure, the taxpayer is allowed to deduct more than 100% of the cost of the capital asset. Bonus depreciation is used considerably less frequently than accelerated depreciation as a policy, although Congress did provide for it in the Job Creation and Worker Assistance Act of 2002, in the wake of the September 11 attacks.⁷³ Bonus depreciation is often combined with accelerated depreciation, though it can be effective as a standalone policy.

⁶⁴ Review of Business Taxation, *A Platform for Consultation*, Vol 1, (1999), 117.

⁶⁵ Andrew Harper, “Finance Act Notes: Annual Investment Allowance, etc- Sections 74-76 and Schedule 24”, (2008) BTR 480, 483; and Andrew Harper, “Finance Act Notes: Section 24- First-Year Capital Allowances for Expenditure in 2009-2010”, (2009) BTR 505, 507.

⁶⁶ Rebecca Morrow, “Accelerating Depreciation in Recession” (2016) 19 Fla. Tax Rev. 465, 472, citing Christopher Hanna, “Tax Theories and Tax Reform”, 59 SMU L. REV. (2006), 435, 441.

⁶⁷ *ibid.*, 469.

⁶⁸ *ibid.*, 472.

⁶⁹ Review of Business Taxation Vol 1 (n 64), 117.

⁷⁰ Royal Commission on Taxation, *Report of the Royal Commission on Taxation, Volume 6: Implication of the Proposed Tax Reforms*, (1966), 92.

⁷¹ *ibid.*

⁷² Morrow (n 66), 473.

⁷³ *ibid.*, 481.

Assessing the Two Dimensions

Assessing Accelerated Depreciation

If accelerated depreciation can give taxpayers considerable benefits, it follows that there must be a cost to the Government. These costs largely mirror the benefits that the taxpayer receives. The most obvious cost is the “‘interest’ on this ‘loan’ less any revenue from taxation of the extra income earned by the taxpayer as a result of the ‘loan’”,⁷⁴ after factoring in the decrease in the value of money due to inflation. The Government also bears a “certainty risk”, since there is no guarantee that the firm will still be in business with it is time to collect the deferred taxes. However, the sheer number of taxpayers in a jurisdiction creates very efficient spreading of the risk, putting the Government in the strongest position to bear such a risk. Unless a Government has a persistent and major fiscal deficit, the benefits to the taxpayer in terms of cash flow and availability of funds are not a significant cost to the Government.

It is noted that accelerated depreciation means that in the early years where capital assets are purchased, the Government will suffer revenue losses, which will be recovered in later years. This effect can be very large indeed if the expenditure on capital assets as a society constantly increases, as would be expected in a growing economy.⁷⁵ This may result in a budget deficit if the Government does not estimate the effect of accelerated depreciation accurately.

Assessing Bonus Depreciation

By allowing a deduction greater than the cost of a capital asset, the Government is basically collecting less tax revenue than it otherwise would have. An interesting observation from empirical research is that bonus depreciation can have a very powerful effect on investment in long-lived capital assets. House and Shapiro’s findings indicate that the investment supply elasticities of long-lived capital assets are very high, making bonus depreciation policies very effective. Their study also found no evidence of an increase in market prices as a response to bonus depreciation.⁷⁶ This suggests that bonus depreciation is a very powerful policy tool, provided that the Government is willing to bear the potentially significant costs of allowing for bonus deductions.

General Comments

Adjustments to MACRS basically directly affect only those businesses that are capital - intensive.⁷⁷ As such, this distortionary effect must be understood and factored in when formulating tax policy. Studies suggest that capital investment is more volatile than spending by consumers or governments,⁷⁸ making adjustments to MACRS potentially very effective.

⁷⁴ Review of Business Taxation Vol 1 (n 64), 117.

⁷⁵ *ibid*, 118.

⁷⁶ House and Shapiro (n 48), 762.

⁷⁷ Review of Business Taxation Vol 1 (n 64), 118.

⁷⁸ Loutzenhiser (n 58), para 6.12.2.

Automation Regulation and Depreciation

Automation technology is progressing at a rapid rate and manifesting in so many different forms that it would be a considerable challenge for regulators to attempt to use a general tax provision to encompass all the relevant cases. Further, the management of the rate at which automation displaces human workers is a very delicate matter. The impact of automation in different industries varies considerably and there are good reasons for tailoring the tax treatment of the various assets according to their impact on human workers. As such, the flexibility offered by the MACRS schedular system is particularly suitable in the automation regulation context. Further, the MACRS requires that an asset must be used in a trade in order for its cost to be deducted against trade income. This ensures that it is only the displacement of human workers that is addressed; non-commercial uses of the assets, for example by consumers, would not be caught by the system. This avoids the creation of distortions beyond what is specifically targeted.

Two Dimensions of Adjustment (Automation Regulation)

Our proposed system of automation regulation would involve adjustments to the MACRS along the two dimensions considered above. As the function of automation regulation is to disincentivize specific uses of capital assets, the policy tools available to the Government would be 1) decelerated depreciation, and 2) reverse depreciation (appreciation).

Decelerated Depreciation

This model simply involves a reversal of the declining balance method. Rather than allowing faster depreciation at the time of acquisition of the capital asset, which slows towards the end of the working-life of the asset, a suitable disincentive would be provided by reversing the process and making the depreciation rate start out slowly. A stronger version of this proposal would involve lengthening the overall period over which depreciation takes place. For example, if a capital asset has a working life of 10 years, decelerated depreciation might require a firm to depreciate the asset over 20 years, even after the asset has ceased to exist, or is sold.

The incentives considered above under accelerated depreciation would simply be reversed under our proposal and become disincentives. The taxpayer is forced to make an interest-free loan to the Government, putting some strain on cash flows and potentially requiring the firm to have recourse to external capital. The certainty risk of being able to set off the pre-paid taxes against future taxes would rest on the taxpayer, who may no longer be in business by the time they are allowed to do so. With all of these disincentives (which can be calibrated through the rate of depreciation), the Government has a powerful economic tool which it can use to slow down the displacement of human workers by automation.

Reverse Depreciation (Appreciation)

Conceptually, there are five different positions which the Government can take when it comes to the deductible value of an asset. It can allow deductions of 1) more than 100% of the value (bonus depreciation); 2) exactly 100% of the value (neutral depreciation); 3) less than 100% of the value (reduced depreciation); 4) 0% of the value (no depreciation); or 5) less than 0% of the value (reverse (or negative) depreciation). “No depreciation” is an interesting case, because it treats the acquisition of a capital asset as a non-event for the purposes of income tax. The taxpayer’s income is unaffected by the capital expense. Reverse depreciation is simply the inverse of bonus depreciation. Instead of being allowed to deduct more than 100% of the cost of a qualifying capital asset, a firm acquiring a specified capital asset under the reverse depreciation regime will have a certain percentage of its cost treated as income. This adjustment is similar in effect to the tax recognition of a notional appreciation of the value of a capital asset. Reverse depreciation can be used as a standalone policy or together with decelerated depreciation.

Assessing the Two Dimensions (Automation Regulation)

The costs to the Government under accelerated depreciation (discussed above) become potential sources of revenue under decelerated depreciation. In the case of reverse depreciation, by disallowing certain portions of the cost of capital assets, the Government basically deems the profit of a firm to be higher than its true profit. The additional tax revenue collected from this exercise can then similarly be hypothecated and used to correct the social externalities created by automation.

Distinguishing Employment-Complementing from Employment-Substituting Capital

Having established that the two goals of our proposed tax reform are best accomplished by making adjustments to the system of capital allowances, we will now suggest two approaches for distinguishing between employment-complementing capital, expenditure on which will enjoy higher rates of tax deductibility, and employment-substituting capital, expenditure on which will be relatively less tax-deductible.

Both approaches are similar in that they assess the causal effect of capital on employment in order to determine the rate at which expenditure on this capital should be tax-deductible. The difference between both approaches lies in how the causal effect of the capital investment on employment is determined. The first approach assesses the collective effect of the firm’s capital investments as a whole, while the second approach targets individual instances of capital investment. Another difference is that the first approach takes a backward-looking and empirical approach to measuring the effect of the firm’s capital investment decisions on employment, while the latter approach forecasts the likely effects of individual instances of capital investment based on the features of the asset purchased.

The first proposed tax reform ties the overall deductibility of capital investment to the company's overall record of employing workers. Under this proposal, the percentage of the annual depreciation cost of the capital investment that is deductible as part of the company's capital allowance will vary each year depending on the total net percentage change in employment in that year. Companies that reduce employment levels in a given year will only be allowed to deduct part of the full annual depreciation cost of their investment from their payable income in that year, while companies that increase overall employment levels will instead be granted a bonus deduction of more than 100% of the annual depreciation cost of that investment. This reform has several features and implications.

It considers the company's use of capital and assets holistically. Instead of assessing each piece of capital in isolation to determine its effect on employment, it treats all the capital of the company as an entire system along with the company's processes and operations. This accounts for the possibility that the effect of a piece of capital on employment is determined not only by its nature and type but also by the way it is deployed within the company as well as its interactions with other capital investments.

By adopting a backward-looking and empirical approach of observing employment data, this approach offers greater accuracy in determining the magnitude of the change in employment that follows from a company's investment policies, at the cost of reducing certainty about the cause of this change in employment. The advantage of this approach is that it creates a strong incentive for companies to raise or maintain employment by directly targeting the intended outcome – ensuring that employment levels are maintained in the face of increasing automation. The disadvantage is that it is difficult to determine whether observed changes in employment levels should be attributed to the capital investments made by the company, blunting the incentive for the company to invest in employment-complementing capital.

It provides the greatest incentives to raise employment for capital-intensive companies. As these companies have higher levels of capital expenditure, they stand to gain more from capital allowances should they invest in employment-complementing capital. This contributes to improving the overall productivity of the economy by providing greater incentives for job creation in high-productivity companies.

It provides companies with distinct direct and indirect incentives to maintain or increase net levels of employment. The direct incentive arises because companies that reduce their net employment levels suffer an immediate reduction in the tax-deductibility of their capital expenditure. The indirect incentive arises because companies that face a choice between an employment-complementing capital asset and an employment-substituting capital asset will, all things equal, expect a future tax benefit from choosing the former over the latter.

Companies that have made capital investments in the past will face the same incentives to maintain or raise overall employment levels as companies that are planning to make capital investments, because the deductibility of both past and future capital expenditure depends solely on changes in present employment levels.

By tying the deductibility of capital expenditure to net employment instead of gross jobs displaced, this proposal creates an incentive to hire workers that is symmetrical to the disincentive to retrench workers. The upshot is that labour market rigidities are minimized because firms are not penalized for retrenching workers if they do not have significant profits against which capital expenditure can be deducted or if they balance these retrenchments with the hiring of new workers.

The second possible approach is to remove or reduce the deductibility of capital expenditure on a specified group of items, regardless of whether capital expenditure on these assets coincides with any changes in overall firm employment levels. These specified group of items are those capital assets that are deemed as significant contributors to the problem of automation-induced job displacement, based on the conjunction of several criteria: 1) their propensity to be deployed in ways that displace existing jobs, 2) their potential for widespread adoption and by extension for widespread job displacement, and 3) their low likelihood of contributing to the creation of new job opportunities.

This approach relies on the classification of capital assets into categories based on their propensity to substitute or complement employment, as well as the subsequent promulgation of this classification as schedules detailing the deductibility rates of capital expenditure on different assets. The use of schedules here as an administrative mechanism is well-established. Indeed, schedules that specify the depreciation are used in the administration of the existing capital allowance system. Modifying and expanding these schedules to also specify different rates of deductibility for different capital assets is a natural extension of the existing capital allowance system. A panel of independent experts can be convened to draft the schedules and perform the categorization of assets.

This approach avoids labour market rigidities because the deductibility of capital expenditure is tied to the nature of the asset, instead of the firm's decisions on employment. Hence, the firm is not subject to any new restrictions or policies on hiring or firing workers, preserving labour market flexibility.

Leakage

Our proposal of reducing the deductibility of capital expenditure for some classes of capital assets faces the major challenge of avoiding leakage through outsourcing and offshoring. Companies can respond to our proposal by offshoring or outsourcing production that intensively uses employment-substituting capital to foreign contractors or suppliers that are not affected by our proposed changes to the tax regime.⁷⁹

The EU faced a similar issue, termed “carbon leakage”, in its implementation and enforcement of the EU Emissions Trading Scheme (“ETS”), a cap-and-trade policy on

⁷⁹ Muûls (n 38), 5-8.

greenhouse gas emissions.⁸⁰ While the ETS was introduced to encourage companies to reduce emissions, there was concern that companies would respond by outsourcing emissions-intensive production to countries with less stringent environmental laws in order to ensure that their own emissions remained within their allowances.⁸¹ The extent to which carbon leakage has actually blunted the impact of EU emissions policy is unclear.⁸²

The upshot is that the possibility of “automation leakage” should be seriously considered by policymakers when designing and implementing changes to the existing system of capital allowances. Notwithstanding this, the extent of any automation leakage is likely to be less than that of carbon leakage. The suppliers and contractors in the case of carbon leakage are likely to be based in less-developed economies with a weaker incentive to implement stringent environmental regulations. In the case of automation leakage, however, the infrastructure, technological ecosystems, and highly-specialized labour required for automation-intensive production are likely to be found in highly-developed economies – the very economies with the strongest incentive to discourage excessive automation for fear of displacing domestic employment. This reduces the potential severity of the issue and creates the possibility of cooperation between developed economies to jointly tackle the related problems of automation-induced job displacement and automation leakage.

In the event that the issue of automation leakage materializes, it is likely to be concentrated in a limited number of sectors. The reason for this is that the relocation of production to avoid the incentives against using employment-substituting capital is only viable under the following limited conditions: first, the returns from employment-substituting capital in that sector must be significantly greater than the returns from employment-complementing alternatives; second, the sector, or at least crucial links in the supply chain, must be tradable; finally, it must be economically viable in that sector to relocate supply chains to an economy without similar policies to discourage employment-substituting capital.

Conclusion

In light of the research by House and Shapiro indicating that the investment supply elasticities of long-lived capital assets are very high,⁸³ the adjustments we have proposed would seem to have considerable potential to affect taxpayer behaviour. When we consider that most forms of automation are highly capital-intensive and that adjustments to the MACRS generally only affect businesses which are capital-intensive,⁸⁴ the proposed adjustments seem reasonably fit for purpose. These effects, when combined with the revenue generation function of the adjustments, make them a very viable proposal for the “automation tax”. In fact, the schedular

⁸⁰ Fiona Harvey, “Britain merely 'outsourcing' carbon emissions to China, say MPs” (18 April 2012) (available at <https://www.theguardian.com/environment/2012/apr/18/britain-outsourcing-carbon-emissions-china>) (accessed on 28 September 2018); and Borkent et. al. (n 34).

⁸¹ Harvey (n 80); Muûls (n 38), 5-8.

⁸² Sander de Bruyn, Dagmar Nelissen and Marnix Koopman, “Carbon leakage and the future of the EU ETS market” (April 2013), 1-60; and Muûls (n 38), at p 5-8.

⁸³ House and Shapiro (n 48), 762; Loutzenhiser (n 58), para 6.12.2.

⁸⁴ Review of Business Taxation Vol 1 (n 64), 118.

nature of the MACRS and capital allowances system provide the Government with much needed flexibility to successfully regulate the rapidly developing field of automation. We therefore conclude that the MACRS is likely to be the most appropriate candidate for the proposed “automation tax”.

G. Conclusion

This article has argued that by using reverse depreciation/ capital allowances, governments and tax authorities are able to make use of an existing and well-established system that is uniquely well-suited to deal with the problems of lack of precision and slowness of response to change. An automation tax could practically be implemented using reverse depreciation/ capital allowances as a mechanism. As a useful tool for governments to have on hand, an automation tax can be quickly implemented by building on the existing depreciation/ capital allowances framework where necessary. It can be used to manage the balance between the positive and negative externalities of automation and artificial intelligence by calibrating the level of their adoption through the use of these tax incentives. As the benefits from the efficiency savings from automation and artificial intelligence continue to be attractive to the majority of states, we do not anticipate that the robot tax will be adopted widely. However, it remains a useful policy tool in those select situations where social considerations may need to be prioritized.