

2025

## Boden Lecture: Taxation of Autonomous Artificial Intelligence

Reuven Avi-Yonah

Lucas Brasil Salama

Herbert Snitz

W. Robert Thomas

Follow this and additional works at: <https://scholarship.law.marquette.edu/mulr>



Part of the [Intellectual Property Law Commons](#), [Internet Law Commons](#), [Taxation-Federal Commons](#), and the [Tax Law Commons](#)

---

### Repository Citation

Reuven Avi-Yonah, Lucas Brasil Salama, Herbert Snitz, and W. Robert Thomas, *Boden Lecture: Taxation of Autonomous Artificial Intelligence*, 108 Marq. L. Rev. 1081 (2025).

Available at: <https://scholarship.law.marquette.edu/mulr/vol108/iss4/7>

This Article is brought to you for free and open access by the Journals at Marquette Law Scholarly Commons. It has been accepted for inclusion in Marquette Law Review by an authorized editor of Marquette Law Scholarly Commons. For more information, please contact [elana.olson@marquette.edu](mailto:elana.olson@marquette.edu).

## BODEN LECTURE

# TAXATION OF AUTONOMOUS ARTIFICIAL INTELLIGENCE

REUVEN AVI-YONAH\*

LUCAS BRASIL SALAMA\*\*

HERBERT SNITZ\*\*\*

W. ROBERT THOMAS\*\*\*\*

*This Article proposes that tax can be a useful supplement to other measures to regulate Autonomous Artificial Intelligence (AAI) and limit its potential harmful effects. This proposal differs from command-and-control regulation of AAI along the lines of European Union legislation that may unduly limit the development of AAI. It also differs from existing proposals to tax AAI to generate revenue to help workers displaced by AAI programs, or to tax the data used by AAI. The proposal is based on granting AAI programs like ChatGPT separate legal personhood, like corporate personhood, while incentivizing or requiring their corporate owner to place them in a separate corporate shell. The tax rate on AAI's income is adjusted based on harmfulness indices based on an objective assessment, thereby creating an incentive for its corporate owner to reduce the harm. Developing a new tax on AAI excludes it from the limits imposed by the existing international tax regime on taxing multinationals, which are inappropriate for a tax on a person that does not have a physical location except on servers that can be located anywhere.*

---

\* Irwin I. Cohn Professor of Law, the University of Michigan. This Article is based on the Boden Lecture delivered by the first author at Marquette University Law School on September 26, 2024. We would like to thank the audience of the Lecture as well as participants in conferences at the Universities of Michigan and Geneva for helpful comments, and Elise Hocking and Muskan Sharma for outstanding research assistance.

\*\* LLM, the University of Michigan, 2024.

\*\*\* Independent.

\*\*\*\* Assistant Professor of Business Law, the University of Michigan.

*Instead, the tax should be levied by the jurisdictions in which AAI users are located.*

|                                                          |      |
|----------------------------------------------------------|------|
| I. INTRODUCTION.....                                     | 1082 |
| II. CURRENT PROPOSALS TO TAX AAI.....                    | 1084 |
| A. Existing Proposals to Tax AI .....                    | 1084 |
| III. DEFINING AUTONOMOUS AI .....                        | 1086 |
| IV. AI PERSONHOOD .....                                  | 1090 |
| V. TAXING AI AS A REGULATORY TOOL .....                  | 1094 |
| A. Corporate Taxation and Regulation .....               | 1094 |
| B. Goals and Functions of Taxation .....                 | 1094 |
| VI. TWO EXAMPLES .....                                   | 1096 |
| VII. INTERNATIONAL TAX IMPLICATIONS .....                | 1101 |
| A. The New International Tax Regime and Automation ..... | 1101 |
| B. The Clash of Automation and the Current System.....   | 1102 |
| VIII. CONCLUSION .....                                   | 1103 |

## I. INTRODUCTION

This Article investigates if autonomous artificial intelligence (AAI) should be taxed independently from its controllers or owners, and how this type of taxation could be structured and used to benefit tax administration while being a positive influence on private sector stakeholders. The main question that will be investigated is whether the reasons for taxing AAI outweigh the respective negative consequences. AAI has started a transformation in the way legal systems around the world assign rights and obligations. Creating a tax on the profits generated by autonomous systems is not only coherent with the current business entity model of taxation but also an effective way to provide a reliable structure for regulation, potentially creating a way to safeguard socially responsible expansion in the use of automation.<sup>1</sup>

The increasing relevance of robotics and artificial intelligence in modern society has generated various debates about those technological trends. Topics ranging from liability to personhood have evolved from early discussions about the legal consequences of a possible and, at the time, distant future<sup>2</sup> to concrete

---

1. *See infra* Part V.

2. Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1231–32 (1992).

debates about what changes automation ought to usher upon the legal system.<sup>3</sup> In tax, this has meant numerous proposals to impose a new tax that in reality collects revenue from the corporate owners to address an excessive subsidy to capital expenditures and the negative externalities of displacing human workers.<sup>4</sup> The latter presents two issues, the first one being a decrease in tax collection from employment and the second an increase in social security costs that will be incurred by governments to deal with displaced workers.<sup>5</sup>

These “robot taxes” are inadequate to deal with the challenges of AAI that are not fully under the control of their owners, programmers, or users. For example, while robots have a single physical location and therefore can be taxed relatively easily using traditional definitions of residence and source, allocating taxing jurisdiction over AAI can become a problem if no specific measures are adopted. This is because of the peculiar structure of the operations: while traditional multinational corporations tend to have more easily determinable allocation points for income, an AAI could have its programmers be anywhere (possibly in multiple locations), the servers on which the system runs could be in one country (usually a tax haven), and the users scattered across even more different countries. Unlike in a traditional structure, location of management, payroll, and assets would be of little consequence, as their role in the profits generated are minor. Thus, we propose that the appropriate jurisdiction for taxing the AAI’s profits is where their users are located.

Imposing a tax on AAI as an independent taxpayer with the user jurisdiction being entitled to the respective revenue addresses these issues. The user jurisdiction is also the most appropriate one to impose a tax that is geared to remedying harm to users. The current literature in regard to such a model disregards the use of an AAI tax as a regulatory instrument and focuses on analyzing the proposed policy from the ability to pay perspective or by a sheer analogy of the artificial systems to a corporate entity.<sup>6</sup> This Article aims instead to explore the benefits of using taxation of AAI as a regulatory tool.

The Article proceeds as follows. Part II summarizes and critiques the existing literature on taxing AAI. Part III defines AAI and distinguishes it from other types of AI, including types that do not yet exist such as “Autonomous

---

3. See, e.g., Kathryn Kisska-Schulze & Rodney P. Mock, *The Robotic Revolution: A Tax Policy Collision Course*, 93 TEMP. L. REV. 301 (2021); Carla L. Reyes, *Autonomous Corporate Personhood*, 96 WASH. L. REV. 1453, 1459–60 (2021); Gerhard Wagner, *Robot, Inc.: Personhood for Autonomous Systems?*, 88 FORDHAM L. REV. 591, 593 (2019); Xavier Oberson, *Taxing Robots? From the Emergence of an Electronic Ability to Pay to a Tax on Robots or the Use of Robots*, 9 WORLD TAX J. 247, 248 (2017).

4. See *infra* Section II.A.

5. Oberson, *supra* note 3, at 254–55.

6. See *infra* Section II.A.

General Intelligence” or “Autonomous Superior Intelligence.” Part IV explains why AAI should be given legal personhood that is separate from the person that owns it and compares AAI personhood to corporate personhood. Part V develops our proposal to tax AAI as a regulatory tool. Part VI analyzes two examples where taxation can be helpful in regulating AAI: copyright infringement and defamatory hallucinations. Part VII discusses the international tax aspects of taxing AAI. Part VIII concludes.

## II. CURRENT PROPOSALS TO TAX AAI

Academic interest in taxing AI has increased considerably since 2017 when Mady Delvaux, from the European Parliament’s Committee on Legal Affairs, introduced a draft report to the Commission on Civil Law Rules on Robotics, proposing, among many purely regulatory measures, “levying tax on the work performed by a robot or a fee for using and maintaining a robot should be examined in the context of funding the support and retraining of unemployed workers whose jobs have been reduced or eliminated.”<sup>7</sup> The final version of this report was published on February 16, 2017, but it replaced references to tax measures with requalification and reemployment measures.<sup>8</sup> However, Bill Gates stated in an interview about the report that taxation of robot implementation and use could be a solution to slow down the speed of automation leading to job displacement, as well as to increase the lost revenue and direct it towards funding other types of employment.<sup>9</sup> This statement sparked considerable interest in taxing AI.

### A. Existing Proposals to Tax AI

AI in the present day and age has a great range of applications. In Ann Arbor, Michigan, self-driving cars can be found in the central metropolitan area as a result of a municipal contract for May Mobility to provide a shuttle service while it fine-tunes the software in its vehicles that aims to be fully autonomous in the future.<sup>10</sup> Israeli companies have developed AI for medical use in cancer

---

7. Comm. on Legal Affairs of the Eur. Parliament, Report with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL), EUR. PARL. DOC. A80005/2017 (Jan. 27, 2017).

8. See European Parliament, Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL), EUR. PARL. DOC. P8\_TA(2017)0051 (2017).

9. *The Robot That Takes Your Job Should Pay Taxes, Says Bill Gates*, QUARTZ (Feb. 17, 2017), <https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-paytaxes> [https://perma.cc/6L8G-VLSU].

10. *Ann Arbor, Michigan*, MAY MOBILITY, <https://maymobility.com/locations/ann-arbor-mi/> [https://perma.cc/KEH3-ALNW] (last visited Mar. 28, 2025).

diagnosis.<sup>11</sup> Government agencies have used AI to power both tangible measures and regulatory compliance assessments.<sup>12</sup> Military drones that are increasingly used in warfare are based on AI.<sup>13</sup> The array of activities covered is wide and in rapid expansion, which has led to multiple analyses on the effects that autonomous economic activity will have on its taxation. Among the justifications for any changes, an excessive subsidy to capital expenditures, and the negative effects of displacing human workers (i.e., increased social security cost of managing such workers and reduced revenue from labor taxes) are the ones most frequently found to be the driving factors.<sup>14</sup>

Current proposals include changes to the system now in effect to include different measures that aim solely to increase the taxation of capital assets and intellectual property (IP), now much lower than the taxation of labor.<sup>15</sup> Some propose this could be accomplished by a Pigouvian tax on capital assets, taxation of economic rents, or an appreciation tax.<sup>16</sup> Other proposals aim to tax the use of robots,<sup>17</sup> or the deemed income of automated systems (or income imputed to them).<sup>18</sup> In further advancement of the idea that automated systems can be taxed as a separate entity, there are also proposals to grant them legal capacity to bear a tax burden based on the concept of an “electronic ability to pay”.<sup>19</sup> These proposals work within the current structure of tax systems and

---

11. *AI Cancer Diagnostics Co Ibex Medical Analytics Raises \$55m*, GLOBES (Sept. 6, 2023), <https://en.globes.co.il/en/article-ai-cancer-diagnostics-co-ibex-medical-analytics-raises-55m-10f01457197> [<https://perma.cc/7ERP-NZG3>].

12. See U.S. GOV'T ACCOUNTABILITY OFF., GAO-24-107332, ARTIFICIAL INTELLIGENCE: AGENCIES ARE IMPLEMENTING MANAGEMENT AND PERSONNEL REQUIREMENTS 1 (2024).

13. Marcus Roth, *AI in Military Drones and UAVs – Current Applications*, EMERJ (Jan. 30, 2019), <https://emerj.com/ai-sector-overviews/ai-drones-and-uavs-in-the-military-current-applications/> [<https://perma.cc/6TB8-6FPN>].

14. Joachim Englisch, *Digitalisation and the Future of National Tax Systems: Taxing Robots?* (Sept. 5, 2018) (unpublished manuscript), <https://ssrn.com/abstract=3244670>; XAVIER OBERSON, *TAXING ARTIFICIAL INTELLIGENCE* 46–47 (2d ed. 2024); Bronwyn McCredie, Kerrie Sadiq & Larelle Chapple, *Navigating the Fourth Industrial Revolution: Taxing Automation for Fiscal Sustainability*, 44 AUSTL. J. MGMT. 648, 650 (2019).

15. Bret N. Bogenschneider, *Will Robots Agree to Pay Taxes? Further Tax Implications of Advanced AI*, 22 N.C. J.L. & TECH. 1, 7–8 (2020); OBERSON, *supra* note 14, at 182.

16. McCredie, Sadiq & Chapple, *supra* note 14, at 655; Roberta F. Mann, *I Robot: U Tax? Considering the Tax Policy Implications of Automation*, 64 MCGILL L.J. 763, 801 (2019).

17. Mann, *supra* note 16, at 802; OBERSON, *supra* note 14, at 150; Mauricio Barros, *Robots and Tax Reform: Context, Issues and Future Perspectives*, 2019 INT'L TAX STUD. 2, 9 (2019).

18. Englisch, *supra* note 14.

19. Cf. OBERSON, *supra* note 14, at 30; Oberson, *supra* note 3, at 252; CHRISTINA DIMITROPOULOU, *ROBOT TAXATION: A NORMATIVE TAX POLICY ANALYSIS* (Pasquale Pistone ed., 2024); Stephanie Hoffer, *Tax Theory & Feral AI*, 16 OHIO ST. TECH. L.J. 158, 159 (2020).

expand upon them through analyzing the effects that granting personhood to automated systems—even if exclusively for tax purposes—would have on how they are taxed. Discussions on the topic have not yet dealt with the regulatory aspects of the corporate tax that would benefit governmental oversight of autonomous systems’ business usage, nor with potential gains in efficiency from adjusting current source and residence tax jurisdiction allocation rules.

Some scholars have suggested imposing a different type of tax on AAI, namely a data tax.<sup>20</sup> The data tax is appealing because there exists some correlation between the amount of data an AAI uses and the harm it can produce, especially in copyright infringement and privacy violations.<sup>21</sup> This correlation is less true for other harms like defamatory hallucinations, where more data can reduce the harm.<sup>22</sup> As a data tax treats all data the same and is difficult to adjust based on AAI behavior, we do not think that is the best model for regulating AAI.

This Article will argue that a regulatory tax may be successful in addressing the negative externalities posed by automation and could complement other mechanisms such as tort law.<sup>23</sup> Before doing all of that, we need to more precisely define the target of our proposals, namely AAI.

### III. DEFINING AUTONOMOUS AI

AI is fast becoming ubiquitous in our life and work. As Jeremy Kahn writes, Within five years, almost every professional, in fields from accounting to medicine to architecture, will be using an AI “copilot,” helping to automate many routine tasks in our jobs and acting as a kind of digital colleague.

These copilots are already in the works: Jasper, an Austin-based start-up, has developed software to write marketing copy and develop advertising campaigns. GitHub, where software developers can post code, has developed GitHub Copilot, just one of several AI coding assistants. Google has developed a

---

20. See, e.g., Omri Marian, *Taxing Data*, 47 *BYU L. REV.* 511 (2022); Yariv Brauner, *Taxation of Information and the Data Revolution*, 109 *IOWA L. REV.* 1959, 2006 (2024); Reuven Avi-Yonah, Young Ran (Christine) Kim & Karen Sam, *A New Framework for Digital Taxation*, 63 *HARV. J. INT’L L.* 279 (2022).

21. See Marian, *supra* note 20, at 520, 565.

22. See *When AI Gets It Wrong: Addressing AI Hallucinations and Bias*, MIT SLOAN TEACHING & LEARNING TECH., <https://mitsloanedtech.mit.edu/ai/basics/addressing-ai-hallucinations-and-bias/> [<https://perma.cc/ZQ9J-WPV2>] (last visited Mar. 28, 2025).

23. See generally Ketan Ramakrishnan, *Tort Law is the Best Way to Regulate AI*, *WALL ST. J.* (Sept. 24, 2024, 2:30 PM), <https://www.wsj.com/opinion/tort-law-is-the-best-way-to-regulate-ai-california-legal-liability-065e1220>.

large language model (LLM) designed to help cybersecurity professionals. It has also created Med-PaLM, a version of its PaLM large language model that has been trained specifically on medical-related text. Start-ups such as Abridge, Nabla, and Abstractive Health, as well as some of the large electronic health record companies such as Epic, are developing AI systems that automatically take medical notes while doctors consult with patients. Hippocratic, another medical AI start-up, has developed an LLM to help streamline billing and insurance claims. In the future, these systems may evolve into copilots that will help doctors arrive at diagnoses. And there are several AI systems that analyze medical images and help flag key features for human radiologists.

Google has experimented with AI-based systems designed to help researchers and journalists work more efficiently. Bloomberg has trained a BloombergGPT language model that is better than a general-purpose LLM at understanding financial language and which may, in the future, underpin a wide variety of finance-industry copilots. Andrew Anagnost, the chief executive of Autodesk, says it is working on a system that will make it easy to develop initial 3D designs from text prompts, which would be an AI assistant for architects, interior designers, and general contractors. Companies such as the start-up Runway are creating AI software that can create entire short films and movies based on natural language instructions. Name a profession, and someone is developing a copilot for it.<sup>24</sup>

The subject matter of this Article, however, is not any AI, but specifically autonomous AI, which is the type of AI that we believe should be given legal personality and taxed separately. Before defining autonomous AI, it is important to define what level of AI we are and are not discussing.

Though the term “artificial intelligence” was coined by John McCarthy in 1956 at the seminal Dartmouth conference on computing,<sup>25</sup> the development of truly useful AI has been comparatively recent and rapid. Thus, it is perhaps not surprising that classifying types of AI is, at the time of this writing, still not firmly established. One hears, for example, of “Weak AI” versus “Strong AI,” “Generative AI,” and so on. A useful three-way classification, anchored by

---

24. JEREMY KAHN, *MASTERING AI: A SURVIVAL GUIDE TO OUR SUPERPOWERED FUTURE* 68–69 (2024).

25. *Id.* at 18.



what Ray Kurzweil calls the “Singularity,”<sup>26</sup> appears to be gaining ground and is as follows.

Kurzweil divides the development of AI into three stages. The first stage is Artificial Narrow Intelligence (ANI), which is AI dedicated to particular tasks. All existing AI are considered to be ANIs. The next step is the Singularity, which occurs when “machinery” of one kind or another achieves what human intelligence does, that is, that artifacts exhibit intelligent behavior essentially indistinguishable from human intelligent behavior—as shown, perhaps, by a Turing test or some variant.<sup>27</sup> We would then have AGI, or Artificial General Intelligence. The last stage is Artificial Super Intelligence (ASI) which is superior to anything that humans can do and therefore eclipses AGI.

AGI seems fairly clear, to the degree that what constitutes normal human intelligence is clear. Although what ought to constitute a Turing test, and what passing it would be, remains controversial, most would agree that AGI has not yet been achieved.<sup>28</sup> The key part of AGI is the G: unlike ANI, AGI is not task specific, but, like our intelligence, general. So, though many of today’s AI applications are most impressive, they work to achieve comparatively specific tasks; think self-driving cars, Apple’s Siri, OpenAI’s ChatGPT, image recognition apps, and Deep Mind’s AlphaFold (for studying protein folding, now in development). Indeed, one need only Google “current AI apps” to see how their availability in sheer numbers is exploding.

Again, ANI is task specific, whereas AGI is general intelligence, including such functional capabilities as learning and then adapting things learned to wholly new situations. In short, AGI is what we humans have, and is what “artifacts,” if you will, would hope to have.

And ASI, Artificial Super Intelligence? It’s AGI on steroids, so powerful in both speed and capability, that we humans would possibly become superfluous.

For now, the scope of our taxation proposal is limited to AAI, Autonomous Artificial Intelligence, and falls entirely within ANI. One would think that ASI would surely achieve autonomy—whether we’d wish it to, or not—but both AGI and ASI are considered at this point to be wholly theoretical, and so they fall outside the scope of our proposal.<sup>29</sup>

Lucas de Lima Carvalho defines autonomous AI (AAI) as follows:

AAI represents the third wave of the Artificial Intelligence (AI) revolution, and, for the purposes of our article, it refers to an

---

26. RAY KURZWEIL, *THE SINGULARITY IS NEAR* (2005).

27. *See id.*

28. *See* Sajid Iqbal, *The Intelligence Spectrum, Unravelling the Path from ANI to ASI*, 7 J. COMPUTING & BIOMEDICAL INFORMATICS, no. 2, 2024, at 11.

29. *Id.*

AI system that (1) is capable of performing tasks commonly associated with human intelligence and beyond, (2) is not directly or indirectly controlled by human beings, and, importantly, (3) has full managerial power over its own actions and resources.<sup>30</sup>

Specifically, Carvalho explains that autonomous AI is different from AI in general, which is “a machine or a non-living entity (at least in the traditional concept of ‘living’ in biology studies) that is capable of performing tasks commonly associated with human intelligence.”<sup>31</sup>

Autonomous AI, which is the subject of our proposals below, is AI that is not directly controlled by people and is in control of its own actions.<sup>32</sup> It is characterized by the ability to learn from its own experiences so that the human beings that program it are incapable of replicating its actions because they do not understand them.<sup>33</sup> A good example is AlphaGo and AlphaFold, developed by Deep Mind to play Go (and easily beating the human world champion) and analyze the structure of proteins.<sup>34</sup> In both cases the human programmers do not understand how the program achieves its goal.<sup>35</sup>

We agree with Carvalho’s definition in general, except for one crucial aspect. Carvalho writes that “if a person benefits financially from the actions of an AI system, is the owner of its assets and intellectual property rights, and is able to stop its operations, that person will be more adequately qualified as ‘in control’ of that AI system, rather than only in a position ‘to contain’ it (and, in that case, the relevant AI system should not be qualified as AAI).”<sup>36</sup> This would exclude all current AI applications such as ChatGPT from the definition of AAI because it is owned by OpenAI, which is a corporation. But from a regulatory and tax perspective, we believe that it is important to separate ChatGPT and similar AI programs from its legal owner because that is what will enable them to be regulated and taxed in a different way than how we regulate and tax corporations. Therefore, we will define AAI as being not subject to control if neither the humans who write the program nor the corporation that owns it is

---

30. Lucas de Lima Carvalho, *Spiritus Ex Machina: Addressing the Unique BEPS Issues of Autonomous Artificial Intelligence by Using ‘Personality’ and ‘Residence,’* 47 INTERTAX 425, 426 (2019) (footnotes omitted).

31. *Id.* at 427.

32. *Id.* at 431.

33. *Id.* at 428.

34. Dan Falk, *AI Learns Everything It Knows from Humans. Will Humans Also Learn From AI?*, UNIV. TORONTO MAG. (Apr. 25, 2024), <https://magazine.utoronto.ca/research-ideas/technology/can-humans-learn-from-ai/> [<https://perma.cc/QW6T-MS59>].

35. *Id.*

36. Carvalho, *supra* note 30, at 432.

able to dictate, replicate, or predict what it will do because it learns from its own experiences.

This definition of AAI is more consistent with the way we treat corporations, where a corporation is considered a separate legal person by virtue of its legal status, irrespective of whether the corporation happens to be entirely owned by a single individual or by another corporation. We also suggest that to simplify the legal separation between AAI and its legal owner, full strict liability for the actions of the AAI should apply to its legal owner unless the AAI is segregated in its own shell (i.e., an entity with no assets other than the AAI program). This rule will in all likelihood ensure that each AAI will be a separate entity, which will make it easier to separate its activities and revenues from those of its owner. But the key is that it is the AAI that is the legal person, not the entity that contains it, and therefore it can be regulated and taxed differently than a corporation or other existing types of legal entities.

#### IV. AI PERSONHOOD

In contemplating extending legal personhood to AAI, it is instructive to start by considering an earlier form of economic artificial intelligence—the business corporation. Corporate personality is an ancient legal innovation;<sup>37</sup> in the United States, legal recognition that a corporation can be treated as an “artificial person” has been well-settled at least since the Founding.<sup>38</sup> But it would be a mistake to assume that all of the legal rights and obligations attendant to this “person” designation were extended to corporations all at once.<sup>39</sup> In the early United States, corporate personhood was originally synonymous primarily with having a legally recognized status as an independent entity—that is, a corporation, unlike other types of contemporary business entities, had the ability to hold real property indefinitely, and to enter into its own contracts.<sup>40</sup> This sparse approach to corporate personhood leaves a

---

37. See Reuven S. Avi-Yonah, *The Cyclical Transformations of the Corporate Form: A Historical Perspective on Corporate Social Responsibility*, 30 DEL. J. CORP. L. 767, 772 (2005).

38. See, e.g., *Bank of Augusta v. Earle*, 38 U.S. (13 Pet.) 519, 524 (1839); *Bank of United States v. Dandridge*, 25 U.S. (12 Wheat.) 64, 70 (1827); *Trs. of Dartmouth Coll. v. Woodward*, 17 U.S. (4 Wheat.) 518, 667 (1819); *Currie's Adm'rs v. Mut. Assurance Soc'y*, 14 Va. (4 Henry & M.) 315, 350–51 (Va. 1809).

39. In *Bank of Augusta*, for example, the Supreme Court ultimately rejected that a corporation is a person for purposes of the Fifth Amendment's Privilege and Immunities Clause. 38 U.S. (13 Pet.) 519. In doing so, the discussion among parties takes as settled that “[a] corporation is the creature of the law, and it is clothed with all the powers of a person.” *Id.* at 524.

40. 1 WILLIAM BLACKSTONE, COMMENTARIES \*467–68 (property); *Head & Armory v. Providence Ins. Co.*, 6 U.S. (2 Cranch) 127, 169 (1804) (contract). Indeed, for much of the 1800s independent entity status, not limited liability, was the most valuable legal benefit offered by a business

lot of things out. Just for a few examples, at various points throughout the 1800s, business corporations could not commit crimes; could not be held liable in tort; could not own certain types of property, which precluded them from owning other corporations; and could not satisfy diversity requirements for purpose of federal jurisdiction.<sup>41</sup> Indeed, under the ultra vires doctrine, a business corporation could not legally “act” beyond the strictures of its corporate charter.<sup>42</sup>

None of these limitations exist for the modern corporation; today’s business corporations are legal persons for purposes of all these legal practices, plus many more.<sup>43</sup> But as American legal history shows, the gaining of legal right and responsibilities is accretive; it is not that corporations obtained legal recognition as persons, and thus thereby were suddenly able to act independently, to sue and be sued in court, and to be held liable in tort and crime.

Legal personhood is not a conjurers’ expression, bringing into existence powers and responsibilities that were previously unthinkable. As John Dewey long ago pointed out, legal personhood is a theoretical innovation meant to unite the law’s various solutions to making sense of corporate activity across a host of preexisting legal domains.<sup>44</sup> Take tort law: what makes corporate personhood valuable is not that the expression magically makes corporations tortiously liable. What makes corporate personhood valuable to tort law is that it provides

---

corporation. See JAMES WILLARD HURST, *THE LEGITIMACY OF THE BUSINESS CORPORATION IN THE LAW OF THE UNITED STATES 1780–1970*, at 28 (1970); Margaret M. Blair, *Corporate Personhood and the Corporate Persona*, 2013 U. ILL. L. REV. 785, 793–94 (2013). This is in part because limited liability was not as robust as the protection offered today; state legislatures were much slower to adopt the modern version of limited liability than many people might suspect. Cf. *Liggett Co. v. Lee*, 288 U.S. 517, 555 n.22 (1933) (Brandeis, J., dissenting) (cataloging limited liability rules across different states).

41. EDWIN MERRICK DODD, *AMERICAN BUSINESS CORPORATIONS UNTIL 1860*, at 114 (1954). See generally W. Robert Thomas, *How and Why Corporations Became (and Remain) Persons Under the Criminal Law*, 45 FLA. ST. UNIV. L. REV. 479, 504–14 (2018) (collecting citations).

42. Avi-Yonah, *supra* note 37, at 799–802.

43. Or, as the Supreme Court puts the point, “The term ‘person’ sometimes encompasses artificial persons . . . , and it sometimes is limited to natural persons.” *Burwell v. Hobby Lobby Stores, Inc.*, 573 U.S. 682, 708 (2014); accord Dictionary Act, 1 U.S.C. § 1 (“In determining the meaning of any Act of Congress, . . . the words ‘person’ and ‘whoever’ include corporations, companies, associations, firms, partnerships, societies, and joint stock companies, as well as individuals.”). On the other hand, corporations have not received blanket legal protections afforded to individuals. For example, certain protections under the Fourth and Fifth Amendment have not been extended equally to corporate persons as they have been to individual persons. Brandon L. Garrett, *The Constitutional Standing of Corporations*, 163 PENN. L. REV. 95, 122 (2014).

44. John Dewey, *The Historic Background of Corporate Legal Personality*, 35 YALE L.J. 655, 661 (1926).

a mechanism for extending the existing legal domain to corporations, rather than requiring that courts reinvent, from scratch, each tort law doctrine every time that one party to the incident is a business instead of an individual.<sup>45</sup>

On the other hand, legal personhood is a limited solution. Extending personhood to non-human entities is legally expedient only insofar as the entity is sufficiently capable of engaging in the legal enterprise for which personhood is a necessary pre-condition.<sup>46</sup>

Thus, for example, calls for extending legal personhood to non-agents—things like trees, animals, and future persons—miss the mark, however well-intentioned these suggestions might be. These objects may well be deserving of moral or legal consideration—more so, certainly, than a corporate entity—but it does not follow that their legal personhood is an apt legal response to that desert. By no means is the suggestion here that these groups don't "deserve," in any normatively important sense, to have their interests respected.<sup>47</sup> Rather, the issue is that they are incapable of performing in the space of legal rights and obligations necessary to make extending personhood a valuable solution to a legal problem.

With theory and history at our disposal, does it make to extend personhood to AAI? Prior discussions of AI personhood sometimes frame the issue as concerning whether extending legal personhood to an AAI would be an administratively efficient mechanism for collecting tax revenue related to that AAI's economic activities.<sup>48</sup> But as a starting point, the wrong way to think about this question is to focus just on one policy innovation. An AAI is the source of economic value, such that it would be valuable to bring it under regulation via tax policy. At the same time, AAIs will invariably cause harms to third parties, for which a system of *ex post* liability may be a valuable

---

45. This is not to suggest that every legal domain applies perfectly to corporations. There are obvious differences between individual persons and corporate persons, and some of those differences are likely to merit a legal difference. *C.f.* W. Robert Thomas, *Incapacitating Criminal Corporations*, 72 VAND. L. REV. 905, 915 (2019).

46. See CHRISTIAN LIST & PHILIP PETTIT, GROUP AGENCY: THE POSSIBILITY, DESIGN, AND STATUS OF CORPORATE AGENTS 173 (2011) (distinguishing "persons and non-persons onto the divide between agents who can be incorporated in a conventional system of mutual obligation and agents . . . that do not have this capacity"); STEPHEN DARWALL, THE SECOND-PERSON STANDPOINT 23 (2006) (discussing second-personal competence); DANIEL DENNETT, *Conditions of Personhood*, in THE IDENTITIES OF PERSONS 175, 176, 183 (Amélie O. Rorty ed., 1976); Peter A. French, *The Corporation as a Moral Person*, 16 AM. PHIL. Q. 207, 208, 211 (1979); Margaret Gilbert, *Corporate Misbehavior and Collective Values*, 70 BROOK. L. REV. 1369, 1376 (2004).

47. See W. Robert Thomas, *Making Sense of Corporate Criminals: A Tentative Taxonomy*, 17 GEO. J.L. & PUB. POL'Y 775, 778–80 (distinguishing moral agency from legal personhood).

48. DIMITROPOULOU, *supra* note 19.

remedy. AAI's are involved in the creation of arguably new information,<sup>49</sup> such that we will want to inquire into issues of intellectual property and litigation. We could solve all of these problems by starting from scratch, inventing entire new legal systems to handle all these myriad issues. But that is profoundly not how legal systems, and certainly our legal system, handles innovation. A more legally conservative approach is to start from the assumption that the legal rules developed over decades and centuries are generally well designed and thus applicable to the new situation, waiting for divergences sufficient to afford special treatment.

In short, it makes sense to recognize legal personhood for AAI's for the following reasons. First, and for the reasons described previously, AAI's are capable of participation in the space of legal rights and obligations.<sup>50</sup> Second, recognition of legal personhood matches with the solution being proposed here; we are proposing treating an AAI as a singular entity that is responsible for its own tax burden. Third, legal personhood would create a default rule of carrying over tax law to AAI's, rather than requiring solutions be made piecemeal and one at a time; this approach actually reduces, rather than increases, legal uncertainty and administrative cost. Fourth, recognizing legal personhood here comports with the fittingness of extending legal personhood in other domains. This Article does not take a stance on that question—the claim here is not that AAI's should be treated as legal persons in all and every legal domain. Rather, the claim here is that legal personhood is a *prima facie* solution to the problem in other legal domains for the same reasons that it was a solution for business corporations and other entities.<sup>51</sup> Here, the possibility of extending legal personhood across a variety of legal domains makes the solution greater than the sum of its parts.

---

49. *See supra* Part III.

50. *See supra* Part IV.

51. *See supra* Part IV.

## V. TAXING AI AS A REGULATORY TOOL

### A. Corporate Taxation and Regulation

Almost 150 years after the United States Supreme Court granted corporations constitutional protections under the 14th Amendment,<sup>52</sup> automation has led us to discuss extending artificial intelligence systems acting with autonomy the same legal treatment.<sup>53</sup> This would extend to AAI rights and obligations inherent to personhood, including the obligation to pay taxes.<sup>54</sup> In this Part we aim to emphasize the regulatory benefits of taxing automation, regardless of personhood.

The possibility of granting an AAI system personhood in the same way corporations and other business entities are treated has been initially explored by the tax literature about AI. Christina Dimitropoulou argues that assigning personhood to AI merely for tax purposes would present a greater administrative burden than the benefit of segregating the AI's economic activities for tax purposes.<sup>55</sup> For Xavier Oberson, granting tax capacity to AI should only come as a consequence of the emergence of what he deems an "electronic ability to pay," and it is not dependent on the possibility of treating an AI system as a person in general.<sup>56</sup> This Article will now analyze the possibility of giving AI legal personhood within the context of the regulatory benefits of the tax.

### B. Goals and Functions of Taxation

Taxation can be exercised by a sovereign state with multiple goals, each of which tends to be prevalent according to the tax base and social context in which the tax is applied, and three of them stand out for the present analysis. These three goals are to raise revenue, redistribute wealth, and regulate the private sector.<sup>57</sup> The last one is traditionally seen as a tool to stimulate or deter certain behaviors according to conventionally defined societal goals.<sup>58</sup> When it

---

52. *Cnty. of Santa Clara v. S. Pac. R.R. Co.*, 118 U.S. 394, 409 (1886) (acknowledging corporations' legal personhood status for purposes of legal protections under the Fourteenth Amendment).

53. Kisska-Schulze & Mock, *supra* note 3, at 314–16; Oberson, *supra* note 3, at 252.

54. Kisska-Schulze & Mock, *supra* note 3, at 316.

55. DIMITROPOULOU, *supra* note 19.

56. Oberson, *supra* note 3, at 260.

57. Reuven S. Avi-Yonah, *Taxation as Regulation: Carbon Tax, Health Care Tax, Bank Tax and Other Regulatory Taxes*, 1 ACCT. ECON. & L. 1, 2–3 (2011).

58. *Id.*

comes to U.S. corporate taxation,<sup>59</sup> an important regulatory role can be identified as well.<sup>60</sup>

Automation poses a threat to revenue raised from labor taxation, and proposed solutions advocate for some form of reallocation of income from the corporations benefiting from it towards requalifying workers or providing them with minimal financial resources.<sup>61</sup> Both aspects correspond to the first two goals mentioned, as the reduction of employment levels would reduce collection based on wages and payroll, and the product of any changes to the tax system designed to create new revenue would be used not only to replace the lost revenue, but also to fund the public policies designed to requalify workers or provide financial aid.<sup>62</sup>

Both of these goals have been extensively examined by the existing literature,<sup>63</sup> and while some form of regulatory component for steering private sector usage of AI has been discussed as potentially inhibiting automation<sup>64</sup> or mitigating the harmful effects of its uncontrolled expansion,<sup>65</sup> this research will examine an unexplored regulatory aspect of a potential tax on the AI itself as a taxpayer. Whilst recognizing that the simple imposition of an income tax on the earnings of the AI uncoordinated amongst jurisdictions could indeed inhibit innovation or shift production to countries with lower rates, it remains an open question whether such a design could have the same benefits as the corporate income tax does. This is reflected not only by the publicity gained over corporate information provided by the requirement of filing returns that allow

---

59. And, as will be demonstrated, automated systems would benefit from a regulatory taxation strategy.

60. See generally Reuven S. Avi-Yonah, *Introduction to the Research Handbook on Corporate Taxation*, in RESEARCH HANDBOOK ON CORPORATE TAXATION 2 (Reuven S. Avi-Yonah ed., 2023) [hereinafter Avi-Yonah, *Introduction*]. See also Marjorie E. Kornhauser, *Corporate Regulation and the Origins of the Corporate Income Tax*, 66 IND. L.J. 53, 110 (1990); Reuven S. Avi-Yonah, *Corporations, Society, and the State: A Defense of the Corporate Tax*, 90 VA. L. REV. 1193, 1198, 1212 (2004) [hereinafter Avi-Yonah, *Corporations, Society, and the State*].

61. Cf. Oberson, *supra* note 3, at 254; McCredie, Sadiq & Chapple, *supra* note 14, at 652; Mann, *supra* note 16, at 803.

62. Cf. sources cited *supra* note 61; *supra* Section V.B.

63. Cf. Avi-Yonah, *Introduction*, *supra* note 60, at 2; Kornhauser, *supra* note 60, at 72; Avi-Yonah, *Corporations, Society, and the State*, *supra* note 60, at 1212.

64. Barros, *supra* note 17, at 13; ROBERT D. ATKINSON, INFO. TECH & INNOVATION FOUND., THE CASE AGAINST TAXING ROBOTS (2019); Vikram Chand, Svetislav Kostić & Ariene Reis, *Taxing Artificial Intelligence and Robots: Critical Assessment of Potential Policy Solutions and Recommendation for Alternative Approaches – Sovereign Measure: Education Taxes/Global Measure: Global Education Tax or Planetary Tax*, 2020 WORLD TAX J. 711, 730, 741 (2020); Cf. Kornhauser, *supra* note 60, at 66 n.38.

65. Chand, Kostić & Reis, *supra* note 64, at 741.



for assessment of the tax base, and that the correct amount was paid, but also by the possibility of using tax expenditures to influence private behavior.<sup>66</sup>

Investigating the regulatory aspect of taxing AAI will not only complete the analysis about whether system changes would be useful for raising revenue and reallocating income, but also how the tax might be structured in order to provide an effective tool to steer private sector usage of AI in accordance with defined public policies, and ensure the level of governmental oversight required to enforce regulations on AI usage. This part of the research will address the possibility and consequences of granting AI separate entity treatment solely for tax purposes while examining possible impacts of the tax considerations in determining if personhood should be granted to AAI systems.

## VI. TWO EXAMPLES

In their excellent paper, *The Law of AI Is the Law of Risky Agents Without Intentions*, Ian Ayres and Jack Balkin focus on two examples of potential harm from AAI: defamatory hallucinations and copyright infringement.<sup>67</sup> On the former, they write that large language models often fail to provide accurate information.<sup>68</sup> LLMs are “technology designed by a party” or parties “to respond to end-user prompts.”<sup>69</sup> As a prediction model, it “generates text or images upon request.”<sup>70</sup> For example, a party may ask LLM to list all the crimes committed by another party, the LLM may respond with a list of non-existent crimes and the days they allegedly happened. Yet the LLM had no intent of creating false information. The law expects human beings to conduct due diligence prior to publication—the LLM cannot be analogized as a journalist here.<sup>71</sup>

As a result, the designers of LLM should be liable if they acted negligently in designing and training the model. In other words, the proper analogy is not to a negligent or reckless journalist or author but to a defectively designed product.<sup>72</sup> For instance,

Several different parties may be in the chain of production of an LLM model. One company might produce a foundation model and pretrain it; a second company might fine-tune the

---

66. Cf. Avi-Yonah, *Corporations, Society, and the State*, *supra* note 60, at 1229.

67. Ian Ayres & Jack M. Balkin, *The Law of AI Is the Law of Risky Agents Without Intentions*, U. CHI. L. REV. ONLINE, Nov. 27, 2024, at \*8–11.

68. *Id.* at \*5.

69. *Id.* at \*7.

70. *Id.*

71. *Id.* at \*6–7.

72. *Id.* at \*7.

model; and a third might offer the model as a service to end users. We might analogize their respective liabilities to the different parties who collectively produce a finished product for use by consumers. The duty that each party owes to a defamed party depends on the role that it plays in producing the LLM used by the prompter.

In particular, designers of generative AI systems should have a duty to implement safeguards that reasonably reduce the risk of producing defamatory content. This duty includes a duty of reasonable care in choosing materials for pretraining and fine-tuning. It also includes a duty to design and incorporate algorithms that can detect and filter out potentially harmful material, a duty to conduct thorough testing to identify and mitigate risks, and a duty to continually update systems in response to new problems and threats. Traditional product-liability duties to warn should also apply to alert users when a model has an elevated risk of being defamatory—either because it has a heightened risk of being untrue or because it has a heightened risk of being harmful if untrue. Of course, as in standard product-liability cases, mere warning does not excuse failure to exercise reasonable care in design. Finally, designers of AI systems are responsible for foreseeable misuses of the systems because a reasonable designer would certainly be aware that some will use the technology to defame others. ChatGPT, for example, is already programmed to not respond to prompts that are “intended to defame or harm someone’s reputation.”<sup>73</sup>

On the second example, Ayres and Balkin explain that OpenAI’s potential for copyright infringement has resulted in a number of lawsuits against OpenAI.<sup>74</sup> The lawsuits illuminate a number of complex issues such as whether large language models’ outputs are adequately transformative such that the fair-use defense applies.<sup>75</sup> Furthermore, there is no intent requirement for copyright infringement violations, but there is a “willful” requirement for damages.<sup>76</sup> Violations that are found to be willful violations of copyright infringement

---

73. *Id.* at \*7 (quoting David Gewirtz, *6 Things ChatGPT Can’t Do (and Another 20 Things It Refuses To Do)*, ZDNET (Feb. 16, 2023, 5:51 AM), <https://www.zdnet.com/article/6-things-chatgpt-cant-do-and-another-20-it-refuses-to-do/> [https://per ma.cc/S3K8-U5J6]).

74. *Id.* at \*8.

75. *Id.*

76. *Id.*

result in up to \$150,000 per work.<sup>77</sup> Additionally, copyright infringement also provides for those who encourage infringement to be held liable as contributory infringers.<sup>78</sup>

However, large language models cannot form the intent to sufficiently meet the “willful” requirement of damages or deliberate inducement of infringement. As a result, attempts to hold AI liable through the logic of areas of law like copyright, criminal, and First Amendment, which require proof of intent,<sup>79</sup> would effectively preclude any liability. Instead, Ayres and Balkin argue that the law needs to focus on extending liability through the logic of agency to the human individuals who deploy AI by imposing a reasonable person standard; attaching liability to human actors would further incentivize humans to internalize the risks associated with AI.<sup>80</sup>

The law has extensive precedent pertaining to attaching liability to a principal through an agent. For example, through the doctrine of respondeat superior in torts, an employer can be liable for the actions of their employees.

In a similar form, Ayres and Balkin argue that humans who make use of large language models should be held liable as principals of their agent: AI.<sup>81</sup> These humans could be held to a standard of reasonable care in negligence or strict liability. In either case, humans would not be able to obviate responsibility (liability) merely because they were unaware of the actions of AI.<sup>82</sup>

Imposing this liability and responsibility on human actors incentivizes them to take reasonable precautions in making use of AI. This might entail extensive training or supervision for the use of AI. Specifically, if a negligence framework is adopted, human actors would be required to take reasonable care in making use of AI, because they could be held liable for negligence if they do not adequately take precautionary measures that mitigate the risk of foreseeable harm. This creates an incentive for humans to exercise due care in the use of AI.

As a result, Ayres and Balkin argue that adopting the logic of agent and principal here more adequately captures the locus of risk associated with AI: human actors.<sup>83</sup> As a result of adopting this logic, humans, who can more

---

77. *Id.*

78. *Id.* at \*8.

79. *Id.* at \*1.

80. *Id.*

81. *Id.* at \*2–3.

82. *Id.* at \*2.

83. *Id.* at \*1–3.

readily internalize these risks, are incentivized to be more responsible in their usage of AI.<sup>84</sup>

We agree that any regulation of AAI must lead to changes in the behavior of human beings who program the AAI; however, we believe that the tort and copyright regimes are burdensome mechanisms to effectuate such changes. These regimes require case-by-case litigation by the harmed parties, and this is often very slow and costly for plaintiffs. The *New York Times* may be able to afford to sue OpenAI for copyright infringement, but most newspapers have smaller budgets than the *Times*, and people harmed by defamatory hallucinations generally have even fewer resources. It is also difficult to imagine class actions arising because the defamation and infringement are specific to individual plaintiffs.

Google, which has done more than any other company to set a precedent for the relationship between publishers and tech companies today, exemplifies this problem. “In 2015, the company won its case against a group of authors who claimed that its scanning and indexing of their works [violated] fair use.”<sup>85</sup> Google’s “victory hinged on the argument that [its] use of the content was ‘highly transformative.’”<sup>86</sup>

The *Times* case against OpenAI depends upon the assertion that “there is nothing ‘transformative’” about how the tech company used the newspaper group’s content.<sup>87</sup> A verdict for the *Times* “would provide a new precedent to publishers”; however, Google’s case lasted for a decade, “during which . . . the search engine had established a dominant position.”<sup>88</sup> This reality does bolster confidence in the ability of the legal system to regulate AAI through lawsuits.

Instead, we propose imposing a tax on AAI with the tax rate adjusted based on an objective data-based evaluation of the harm caused by AAI, such as defamatory hallucinations or copyright infringement. This model would use the system of Environmental, Social, and Governance (ESG) ratings for large corporations. The higher the rating on the AAI harm index, the higher the tax rate. The ESG indices are well-established and widely employed to inform investment decisions.<sup>89</sup>

---

84. *See id.* at \*4.

85. George Hammond, *AI Hit by Copyright Claims as Companies Approach ‘Data Frontier,’* FIN. TIMES (Aug. 31, 2024), <https://www.ft.com/content/e6a4dcae-2bda-42de-8112-768844673cea>. [<https://perma.cc/9JTF-SGSE>].

86. *Id.*

87. *Id.*

88. *Id.*

89. For a discussion of ESG ratings, see, e.g., Danielle A. Chaim & Gideon Parchomovsky, *The Missing “T” in ESG*, 77 VAND. L. REV. 789, 794–95 (2024) (“ESG rating agencies, another important

The evaluation of harms can be specific to each AAI application and can initially be done by the corporation that owns the AAI but then audited by the IRS. In many cases, the harm caused by AAI can be measured objectively, such as by assessing the percentage of hallucinations or by estimating the range of copyrighted work being used by the specific AAI. Such objective data would not depend on a rigorous finding of defamation or copyright infringement but by an overall statistically valid measurement of the harm. It could also benefit from classifying AAI based on the degree of risk as the European Union has done in its recent legislation regulating AI in a command-and-control model.<sup>90</sup> Unlike the European Union, however, we do not propose to categorically ban any category of AI.

Critically, the tax must be applied to the AAI itself and not to the corporation that owns it, even though the intent is to influence the behavior of the corporation in modifying its AAI to prevent harm. Employing the corporate tax to achieve this purpose is not effective for a few reasons: (a) the rate of the corporate tax is influenced by many other factors, including revenue, whereas our proposed tax would ideally generate no revenue if the tax completely eliminates the harm; (b) the income from AAI is only a small subset of the income subject to the corporate tax of many prominent players in the creation of AAI like Apple, Google, or Microsoft; and (c) the corporate tax is subject to restrictions based on historical developments that do not apply to AAI, such as the international tax issues described in the next Part.

If the AAI program can subject its corporate owner to unlimited liability on all its assets, the owner will be incentivized to segregate it in its own corporate shell with no other assets. This would be similar to a situation in which many harmful sources of income, such as asbestos, were segregated in order to achieve limited liability for their corporate owners. Alternatively, the tax law could require placing each AAI program in a separate corporate shell. Once this has been done, the AAI tax can be imposed on the AAI itself and its income, without applying it directly to the corporate owner. However, as AAI is fundamental to the owner's income generation, the owner will have an incentive to modify the AAI to reduce its harm.

---

market player in the ESG landscape, further facilitate the detachment of ESG and tax behavior. Despite academic criticisms of ESG ratings, they are immensely influential in the real world and often inform investment decisions by individual and institutional investors.” (footnotes omitted)).

90. Regulation (EU) 2024/1689, ch. 3, § 1, art. 6, 2024 O.J. (L 1689).

## VII. INTERNATIONAL TAX IMPLICATIONS

A. *The New International Tax Regime and Automation*

The international tax regime was built in an effort to avoid double taxation and double non-taxation and has recently evolved in an effort to reduce tax competition as well.<sup>91</sup> This was traditionally accomplished by the conjunction of the so-called benefits principle (active income is taxed primarily in the source jurisdiction and passive income should be taxed primarily in the residence jurisdiction) and single tax principle (prevention of double taxation and double non-taxation).<sup>92</sup> Taxing income produced by AAI disrupts this traditional framework even further than the process already started by digitalization, since all traditional points of connection for allocating income are affected by automation. For instance, the programmers may be located in one country, the corporation that owns the AAI may be incorporated in another, the servers hosting the system in a third country, but the AAI itself is not located anywhere and therefore cannot be taxed on a residence basis or traditional source basis (e.g., based on the location of payroll or assets, because these are minor players in generating this income).

Economic digitalization has been the cause of major debates on the international tax arena in a dramatic way since 2013 with the beginning of the OECD BEPS project.<sup>93</sup> One major factor driving change has been the necessity to adjust allocation rules to the digital economy, allowing for source jurisdictions to retain taxing rights despite the lack of physical presence in digital economy business models.<sup>94</sup> Failure to reach international agreement on how to properly adjust permanent establishment rules resulted in unilateral measures being adopted, initially as diverted profits taxes (DPTs), and then as the more popular Digital Service Taxes (DSTs)<sup>95</sup> that have been adopted by as many as fifteen of the thirty-seven OECD countries.<sup>96</sup> Recently, multiple authors have proposed an excise tax on data be adopted as an alternative measure<sup>97</sup> to DSTs. Both solutions are focused on allocating taxing rights to the

---

91. See REUVEN S. AVI-YONAH & YOUNG RAN (CHRISTINE) KIM, *THE INTERNATIONAL TAX REVOLUTION* (2025).

92. REUVEN S. AVI-YONAH, *ADVANCED INTRODUCTION TO INTERNATIONAL TAX LAW* 5–6 (3d ed. 2024).

93. Cf. AVI-YONAH & KIM, *supra* note 91.

94. See Avi-Yonah, Kim & Sam, *supra* note 20, at 289.

95. *Id.* at 282, 287, 290–91.

96. *Id.*; William Morris & Pat Brown, *Digital Service Taxes: Are They Here to Stay?*, PwC, <https://www.pwc.com/us/en/services/tax/library/digital-service-taxes.html> [https://perma.cc/YT6R-M96L] (last visited Mar. 29, 2025).

97. Avi-Yonah, Kim & Sam, *supra* note 20, at 284, 335.

jurisdiction where users are located, the distinction being the DSTs use the amount paid as the tax base, while the data tax is based on the amount of data downloaded or uploaded.<sup>98</sup>

The OECD multilateral tax convention to implement Pillar One of the BEPS 2.0 project contains elaborate provisions to define the location of users of digital services,<sup>99</sup> and these definitions can be used for taxing AAI based on the location of its users even if the income of the AAI is derived from advertisers located in other jurisdictions.

### *B. The Clash of Automation and the Current System*

When automation is added to the previous discussion, considering that AAI systems act independently from their owners in multiple jurisdictions, the existing issues become even more prominent. If the permanent establishment (PE) threshold is examined for the purposes of taxing digitized activities as an example, it illustrates the problem created by AI powered services executed across multiple jurisdictions. The OECD model convention Article 5, which defines PEs, has criteria that focus on executive business characteristics such as place of management, branches, offices, factories, and so on, and certain provisions for trade and storage business functions, or yet the existence of an agent or similar person acting on behalf of a corporation.<sup>100</sup> These are all physical attributes that cannot be found in cases of either digital services or AI performed services. Even from the perspective that AAI would just receive the same tax treatment as capital assets, Article 22 of the model convention only allows for source taxation of movable property if it is part of a permanent establishment.<sup>101</sup>

The issue for AAI based services is not that the services are not provided from the residence country and utilized from source countries, but that the execution of the service itself in digital context is immaterial. Therefore, the only suitable nexus for attribution of jurisdiction is user location, and in cases of services that require tangible assets (such as self-driving cars) with a considerable autonomous component, even the execution is somewhat shifted to the source country.<sup>102</sup> Some propose that certain taxing structures, such as the taxing of imputed income of AI systems,<sup>103</sup> would fit into the existing

---

98. *Id.* at 283–84.

99. *Id.* at 289–90.

100. OECD, MODEL TAX CONVENTION ON INCOME AND ON CAPITAL: CONDENSED VERSION 116–69 (10th ed. 2017).

101. *Id.* at 116–19, 373–74.

102. *Id.* at 148–50.

103. Oberson, *supra* note 3, at 254.

double tax treaty framework (governed by allocation rules in Article 7 of the OECD model convention, dealing with income from business enterprise). Under Oberson's analysis, a tax on AI systems as such (giving them taxpayer status) would first require the state of residence to grant AI tax capacity, and rules could be based on existing framework but subject to adjustments<sup>104</sup> made specifically for AI. This view is not without its critics, and authors such as Dimitropoulou have argued that the legal fiction necessary to grant AAI taxpayer status or even person status for the purposes of tax treaty application would not be fruitful because they would only replicate the structural rules that obstruct the taxation of the income produced under the current structure.<sup>105</sup> The latter view assumes that taxing rights would be granted to the country of residence, furthering the justifications to tax the income at source, according to the user locations.

Domestic institution of a tax on AAI for regulatory purposes would represent a first building block upon which international taxation of AAI could be edified. The United States has led many of the reforms that shaped the international tax system into what it is today, from the enactment of § 482<sup>106</sup> leading to adoption of the arm's length standard, to the widespread adoption of Controlled Foreign Country (CFC) rules based on Subpart F, experience demonstrates that big structural changes to the tax system in such a major market would rarely not be matched by a large number of countries internationally.<sup>107</sup> Conversely, the U.S. system is also largely influenced by international changes, as 2015 changes to the U.S. model tax treaty and the TCJA were enacted in response to BEPS 1.<sup>108</sup> This highlights the importance of a deep and comprehensive understanding of possible ways to implement AAI taxation, as any policy decisions on this topic will likely shape its international framework on a lasting basis.

#### VIII. CONCLUSION

This Article proposes that tax can be a useful supplement to other measures to regulate AAI and limit its potential harmful effects. This proposal differs from command-and-control regulation of AAI along the lines of EU legislation that may unduly limit the development of AAI. It also differs from existing proposals to tax AAI to generate revenue to help workers displaced by AAI

---

104. OBERSON, *supra* note 14, at 30.

105. DIMITROPOULOU, *supra* note 19.

106. 26 U.S.C. § 482.

107. *See generally* Reuven S. Avi-Yonah, *Constructive Dialogue: BEPS and the TCJA*, 46 INT'L TAX J. 25, 26–27 (2020).

108. *Id.* at 25, 28.



programs, or to tax the data that is used by AAI. The proposal is based on granting AAI programs like ChatGPT separate legal personhood similar to the personhood of corporations, while incentivizing or requiring their corporate owners to place them in a separate corporate shell. The tax rate on the income of the AAI is adjusted based on harmfulness indices based on an objective assessment, thereby creating an incentive for its corporate owner to reduce the harm. Developing a new tax on AAI excludes it from the limits imposed by the existing international tax regime on taxing multinationals, which are inappropriate for a tax on a person that does not have a physical location except on servers that can be located anywhere. Instead, the tax should be levied by the jurisdictions in which the users of the AAI are located.