

Towards eXplainable Artificial Intelligence (XAI) in Tax Law: The Need for a Minimum Legal Standard

Tax administrations globally increasingly rely on artificial intelligence (AI) systems for automation. However, automation has a huge potential impact on the rights of taxpayers subject to algorithmic assessments, which is compounded by the opacity of complex AI systems. This article argues that adequate protection of taxpayers' rights demands the use of eXplainable AI (XAI) technologies that can render the functioning and decisions of tax AI systems understandable for taxpayers, administrative appeal bodies and the courts. This demand follows from the constitutional principles that guide taxation. Still, it is insufficiently addressed by soft and hard law instruments on AI, which do not address the particular information needs of the tax domain. To address this gap, the authors conclude the article by mapping technical and legal challenges for the proper application of explanation techniques to tax AI to ensure that automation does not come at the expense of taxpayers' rights.

Contents

1. Introduction	574
2. The Use of AI in Taxation and the Need for Explanations	578
2.1. Motivation for implementing XAI systems in tax law	581
2.1.1. The strong and compulsory requirement to explain tax decisions	582
2.1.2. No taxation without representation – A constitutional requirement for XAI in tax law	584
2.2. Tax secrecy and trade secrecy as challenges for explanation	585
2.3. Technical challenges and requirements for XAI in the tax sphere	589

* Błażej Kuźniacki: Assistant Professor in Tax & Technology at the Amsterdam Centre for Tax Law, School of Law, University of Amsterdam (Amsterdam, The Netherlands), PwC Netherlands (Global Tax Policy & International Tax Services).
Marco Almada: Researcher at Department of Law, European University Institute (Florence, Italy).
Kamil Tyliński: Data scientist at Mishcon de Reya LLP (London, United Kingdom).
Łukasz Górski: Assistant Professor at Interdisciplinary Centre for Mathematical and Computational Modelling, University of Warsaw (Warsaw, Poland).
Beata Winogradska: Researcher at Faculty of Information Technology at Polish Japanese Academy of Information Technology (Warsaw, Poland) and Kozminski University (Warsaw, Poland).
Reza Zeldenrust: Researcher at the Amsterdam Centre for Tax Law, School of Law, University of Amsterdam (Amsterdam, The Netherlands).
B. Kuźniacki, M. Almada, K. Tyliński and Ł. Górski acknowledge that this article is based upon work supported in whole by the Notre Dame-IBM Tech Ethics Lab. Such support does not constitute an endorsement by the sponsor of the authors' views expressed in this publication. B. Kuźniacki and Reza Zeldenrust further acknowledge that their work on this chapter has been developed within the framework of the Amsterdam Centre for Tax Law (ACTL) research project "Designing the tax system for a cashless, platform-based and technology-driven society" (CPT project). The CPT project is financed with university funding and funds provided by external stakeholders (i.e. businesses and governments) interested in supporting academic research to design fair, efficient and fraud-proof tax systems. For more information about the CPT project and its partners, please visit its website <https://actl.uva.nl/cpt-project/cpt-project.html>. The support received by the authors within the framework of the CPT project does not constitute an endorsement by the sponsors of the views expressed in this publication by the authors. Marco Almada would also like to thank Fundacion Carolina for granting him a doctoral scholarship, under which his work on this research was partially funded.

3.	Constitutional Principles and Human Rights – A Current Path to a Minimum Legal Standard on Tax XAI?	591
3.1.	Constitutional principles and rights related to tax XAI	593
3.2.	XAI in the context of the ECHR	595
4.	Hard and Soft Law Instruments – A Prospective Path to a Minimum Legal Standard on Tax XAI?	600
4.1.	EU law on AI	600
4.1.1.	Data protection law and the right to an explanation	601
4.1.2.	The AI Act	607
4.2.	Non-legally binding guidelines	610
4.2.1.	Recommendation of the OECD Council on Artificial Intelligence	611
4.2.2.	Council of Europe Recommendation on the human rights impacts of algorithmic systems	612
5.	No Minimum Legal Standard for XAI in Tax Law and Fragmentation of the AI Legal Environment	613
6.	Concluding Remarks	615

1. Introduction

Artificial intelligence (AI) systems¹ are increasingly important for modern taxation. Some of these systems are deployed to support tax compliance activities,² for example by automating invoice generation.³ Others are used by tax authorities in their enforcement work: to obtain information about taxpayer behaviour,⁴ guide tax inspection practices,⁵ assess fraud risks⁶ or even automate the procedures involved in seizing assets in cases of fraud.⁷ As these systems become capable of performing tasks that would require specialized human labour,

1. For the purposes of this study, the authors adopt a broad definition of AI: a computer system is an AI system if it uses computing techniques to interact with the world in a way that responds to the opportunities and challenges offered by its context: J.J. Bryson & A. Theodorou, *How Society Can Maintain Human-Centric Artificial Intelligence*, in *Human-Centered Digitalization and Services* p. 306 (M. Toivonen & E. Saari eds., Translational Systems Sciences, Springer Singapore 2019), available at https://doi.org/10.1007/978-981-13-7725-9_16. This definition encompasses machine learning and knowledge representation systems but also leaves room for other computing techniques that might generate outputs such as content, predictions or decisions in a context-sensitive way.

2. See, e.g., J. Alm et al., *New Technologies and the Evolution of Tax Compliance*, 39 *Virginia Tax Review* 3 (2020).

3. L. Di Puglia Pugliese et al., *A Natural Language Processing Tool to Support the Electronic Invoicing Process in Italy*, in *2021 11th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)* (2021), available at <https://doi.org/10.1109/IDAACS53288.2021.9660987>.

4. See, e.g., J. Butler, *Analytical Challenges in Modern Tax Administration: A Brief History of Analytics at the IRS Symposium on Artificial Intelligence & the Future of Tax Law: AI in Tax Compliance and Enforcement*, 16 *Ohio St. Tech. L. J.* 1 (2020); J. Lismont et al., *Predicting Tax Avoidance by Means of Social Network Analytics*, 108 *Decision Support Systems*, pp. 13-24 (2018), available at <https://doi.org/10.1016/j.dss.2018.02.001>.

5. See, e.g., F.S. Antón, *Artificial Intelligence and Tax Administration: Strategy, Applications and Implications, with Special Reference to the Tax Inspection Procedure*, 13 *World Tax J.* 4 (2021), *Journal Articles & Opinion Pieces IBFD*.

6. See, e.g., the use of automated risk assessment by various European tax authorities: D. Hadwick & S. Lan, *Lessons to Be Learned from the Dutch Childcare Allowance Scandal: A Comparative Review of Algorithmic Governance by Tax Administrations in the Netherlands, France and Germany*, 13 *World Tax J.* 4 (2021), *Journal Articles & Opinion Pieces IBFD*.

7. See, e.g., the MiDAS system adopted by the Michigan Unemployment Agency in 2013: R. Calo & D.K. Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 *Emory Law Journal* 4, pp. 827-830 (2021).

they have the potential to save time and money for both taxpayers and governments.⁸ But automation also introduces several risks: AI systems are not error-proof, which means that they can produce biased decisions;⁹ they may be used for purposes beyond the legitimate scope that motivated their introduction,¹⁰ or they may be used in ways that deprive taxpayers of their right to contest potentially wrongful decisions.¹¹ All these risks are compounded by the various forms of opacity that surround AI systems, which may preclude taxpayers from learning about the tax decision-making procedure or even about the existence of a decision based on an AI system in the first place.¹² To ensure successful outcomes, lawyers and computer scientists should work together on approaches that address the potential harm of using AI technologies in taxation.

This article focuses on one such approach: explainable artificial intelligence (XAI) – that is, the development of techniques that make the functioning of an AI system understandable for a given audience.¹³ Such techniques become necessary because applications such as those described above are often powered by machine learning methods, deep learning in particular, which make decisions based on applying complex mathematical models to large volumes of data.¹⁴ This combination of *big data* and mathematical complexity renders AI systems opaque in various ways: the mathematical models that underpin such systems are scientifically backed but difficult to understand for humans, even those with expert training – the sheer scale of the data and computing operations involved in these models is hard to grasp, and trade and industrial secrets prevent disclosure of detailed model explanations.¹⁵ Thus, it is not always obvious how the system’s input affects its output – in other words, what is the exact causal link between the data ingested by the system and the decision it makes.

XAI methods aim to uncover this causal link, thus facilitating several aims. In general, explaining how a system works strengthens the trust in its result. From the perspective of a system developer, those methods make it easier to optimize, debug – i.e. find mistakes – or compare the models. From the user perspective, those methods aim to appease their need to get a deeper understanding of a system’s inner workings. This understanding, in turn, allows end-users to contest the system’s decision or – in some circumstances – to affect the system’s output by altering their own behaviour. For example, a bank customer can alter

8. See, e.g., OECD, *Tax Administration 3.0: The Digital Transformation of Tax Administration* (Forum on Tax Administration, OECD 2020).
9. N. Braun Binder, *Artificial Intelligence and Taxation: Risk Management in Fully Automated Taxation Procedures*, in *Regulating Artificial Intelligence* (T. Wischmeyer & T. Rademacher eds., Springer International Publishing 2020), available at https://doi.org/10.1007/978-3-030-32361-5_13.
10. B.-J. Koops, *The Concept of Function Creep*, 13 Law, Innovation and Technology 1 (2021), available at <https://doi.org/10.1080/17579961.2021.1898299>; L. Scarcella, *Tax Compliance and Privacy Rights in Profiling and Automated Decision Making*, 8 Internet Policy Review 4 (2019).
11. C. Sarra, *Put Dialectics into the Machine: Protection against Automatic-Decision-Making through a Deeper Understanding of Contestability by Design*, 20 Global Jurist 3 (2020), available at <https://doi.org/10.1515/gj-2020-0003>.
12. On opacity as a challenge to AI oversight, see M. Busuioc, *AI Algorithmic Oversight: New Frontiers in Regulation*, in *The Handbook on Regulatory Authorities* (M. Maggetti, F. Di Mascio & A. Natalini eds., Edward Elgar Publishing 2022).
13. A. Barredo Arrieta et al., *Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI*, 58 Information Fusion, p. 85 (2020), available at <https://doi.org/10.1016/j.inffus.2019.12.012>.
14. For a legal introduction to current AI technologies and their dependence on data, see F. Lagioia & G. Sartor, *Artificial Intelligence in the Big Data Era: Risks and Opportunities*, in *Legal Challenges of Big Data* (J. Cannatacci, V. Falce & O. Pollicino eds., Edward Elgar 2020).
15. J. Burrell, *How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms*, 3 Big Data & Society 1 (2016), available at <https://doi.org/10.1177/2053951715622512>.

their spending habits to get a good credit ranking. From the regulator’s perspective, knowing how a system works allows them to exert a larger degree of control over it and – in general – to design better-informed policies regarding the placement of such systems in society.

In tax contexts, the challenges of opacity are further compounded by the high degree of secrecy surrounding the activities of tax authorities.¹⁶ However, legality also requires authorities to justify their decisions with legal arguments, which can be assessed and even questioned by citizens.¹⁷ Yet, current XAI techniques were developed mainly for technical purposes, which are not necessarily aligned with legal justification standards. Even if a machine learning method exhibits high accuracy and detects relevant features (at the explanation level), the legal rationale for a decision might still be unsound.¹⁸ This disconnect between explanation and justification becomes salient as AI-based systems are deployed for various tasks in the tax domain. However, it is not alien to legal thinking, as it appears in the work of legal realists.¹⁹ This suggests the possibility of establishing dialogues between technical and legal approaches to the challenges of AI.

The ultimate goal, in the authors’ opinion, should be to bring technical and legal approaches together to produce a more robust understanding of explanation, which balances transparency requirements and the utility of AI models in the context of their high accuracy and potential to revolutionize the future of public services. Public administration – as well as the private sector – should embark on a journey to become more agile and leverage the current technological advances while respecting the rights of citizens. One of the milestones for this end is establishing a public research agenda to realize the above goal.²⁰

XAI methods, from this perspective, provide the decision-makers with an account of how a given AI-based system works. From that point, it is the obligation of the decision-maker to provide a decision’s addressee with a detailed and understandable account of how the system arrived at a given result, how the system’s assessment was accounted for in the final decision and the validity of this assessment when contrasted with the relevant legal background. In other words, a technically guided explanation is to be transformed into justification by the public authority. Thus, there is a need for a relevant legal and technical investigation to assess how XAI techniques could eventually contribute to the discharge of such duties of *justification* by providing information about how an AI system arrives at its decision in each particular case.²¹ In this article, we focus on how XAI techniques, as expla-

16. See sec. 2.2.
17. A. Bibal et al., *Legal Requirements on Explainability in Machine Learning*, 29 Artificial Intelligence and Law 2 (2021), available at <https://doi.org/10.1007/s10506-020-09270-4>.
18. C. Steging, S. Renooij & B. Verheij, *Rationale Discovery and Explainable AI*, in *Legal Knowledge and Information Systems* (Frontiers in Artificial Intelligence and Applications, IOS Press Dec. 2021), available at <https://doi.org/10.3233/FAIA210341>.
19. As McCormick put it, “[w]hat prompts a judge to think of one side rather than the other is quite a different matter from the question whether there are on consideration good justifying reasons in favor of that rather than the other side”. N. McCormick, *Legal Reasoning and Legal Theory* p. 16 (Clarendon Law Series, Oxford University Press 1994), available at <https://doi.org/10.1093/acprof:oso/9780198763840.001.0001>. Cf. J. Mumford, K. Atkinson & T. Bench-Capon, *Machine Learning and Legal Argument*, in *Legal Knowledge and Information Systems* (E. Schweighofer ed., Frontiers in Artificial Intelligence and Applications 346, IOS Press 2021).
20. A. Zuiderwijk, Y. Chen & F. Salem, *F. Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda*, 38 Government Information Quarterly 3 (2021).
21. Cf. Ł. Górski & S. Ramakrishna, *Explainable Artificial Intelligence, Lawyer’s Perspective*, in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law* (ICAAIL ‘21, Association for Computing Machinery 2021), available at <https://doi.org/10.1145/3462757.3466145>.

nation methods, impact the interpretation and application of tax law *vis-à-vis* the relevant legal background.

Over the past few years, the field of XAI has developed various techniques to explain AI decision-making processes in terms humans can understand. Some of these approaches focus on identifying the factors that an AI system takes into account when making decisions: local post hoc explanations clarify which factors were considered and their relevance for each particular decision,²² whereas global post hoc approaches observe a large range of decisions and build simplified models of how the original AI system operates.²³ This idea of simplification also appears in ex ante approaches to explanation, which attempt to build the AI system itself so that humans can assess its inner workings.²⁴ Each of these approaches can be useful for different purposes.²⁵ For example, taxpayers need to know how they have been individually affected by decisions made by tax authorities. Local post hoc explanations appear to be best suited to such needs. Courts, in turn, might need both local and global post hoc tools in order to assess the lawfulness of tax decisions stemming from – or facilitated by – AI systems. Consequently, XAI techniques are developing in ways that address the needs of multiple stakeholders.

Two challenges could hamper the adoption of XAI techniques for tax AI systems.²⁶ From a technical perspective, explaining decisions in the legal domain can pose difficult problems. Although facts are relevant for the lawfulness of a decision, a decision based on statistically relevant correlations might nevertheless be unacceptable due to a lack of legal grounds.²⁷ Hence, tax explanations also need to account for the normative dimension of tax-related decisions, a task that requires new techniques still under development.²⁸ Even if these techniques are successfully developed, they might have little practical effect if they are not adopted by tax authorities and other stakeholders using AI for tax purposes. While economic factors and social pressure could prompt the adoption of such techniques, this article examines one potential factor that could push tax authorities and other stakeholders towards using XAI techniques: the need to comply with existing legal requirements for justifying tax-related decisions.

-
22. M.T. Keane & E.M. Kenny, *How Case-Based Reasoning Explains Neural Networks: A Theoretical Analysis of XAI Using Post-Hoc Explanation-by-Example from a Survey of ANN-CBR Twin-Systems*, in *Case-Based Reasoning Research and Development* (K. Bach & C. Marling eds., Lecture Notes in Computer Science, Springer International Publishing 2019), available at https://doi.org/10.1007/978-3-030-29249-2_11.
23. N. Mehdiyev & P. Fettke, *Prescriptive Process Analytics with Deep Learning and Explainable Artificial Intelligence*, in *Proceedings of the 28th European Conference on Information Systems (ECIS)* sec. 4.4 (AIS 2020).
24. See, e.g., C. Rudin, *Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead*, 1 *Nature Machine Intelligence* 5 (2019), available at <https://doi.org/10.1038/s42256-019-0048-x>.
25. On the idea of multiple forms of explanation, see C. Zednik, *Solving the Black Box Problem: A Normative Framework for Explainable Artificial Intelligence*, 34 *Philosophy & Technology* 2 (2021), available at <https://doi.org/10.1007/s13347-019-00382-7>.
26. By “tax AI system” (or “AI in tax”, “AI in taxation” etc.) the authors understand any AI system that is used solely or partly for purposes of delivering or facilitating the decisions of tax authorities imposing taxation, granting or denying tax relief and exemptions, selecting taxpayers for tax inspections or other decisions directly or indirectly affecting taxpayers, i.e., tax-related decisions.
27. Bibal et al., *supra* n. 17, at sec. 2.2.
28. For an introduction to explanation techniques in the legal domain, see *id.*, at sec. 3.3.

The remainder of this article is structured as follows. Section 2. examines the roles that AI plays in taxation and the rationales that underpin the decision on whether or not to explain such systems, highlighting the need to strike a balance between the constitutional and legal requirements of transparency, on the one hand, and the limits to disclosure imposed by trade and industry secrets relating to AI systems and state secrecy regarding the operations of tax authorities (tax secrecy) on the other. Section 3. provides an initial assessment of this balance at the constitutional level, looking at how constitutional principles and the European Convention on Human Rights (ECHR) address the risks stemming from tax AI. Section 4. then analyses the effect of hard and soft law instruments on AI technologies by evaluating the transparency requirements such instruments establish and their adequacy for the particular informational needs of the tax domain. Based on these analyses, section 5. argues that, currently, binding and non-binding instruments do not provide a minimum legal standard for XAI in tax law, especially as they have little to say about tax-specific issues, leading to a fragmented regulatory scenario in which taxpayer rights are trampled in the name of administrative efficiency. Deeming this scenario unacceptable, section 6. concludes with a few remarks on the need for a minimum legal XAI standard to ensure that automation in the tax domain does not come at the expense of the constitutional and legal safeguards that ensure the legitimacy of taxation in the modern state.

2. The Use of AI in Taxation and the Need for Explanations

One of the big advantages tax administrations have over private sector stakeholders is that the former can rely on having access to huge volumes of tax-sensitive data as they implement their AI systems.²⁹ This big data grows substantially by the day, drawing from a variety of sources, many of which are high in quality and can only be obtained by tax authorities.³⁰ In this data-rich environment, tax administrations enjoy conditions that are very favourable to the use of AI technologies:³¹ (i) the volume of data is so great that its direct examination is not feasible through human labour;³² and (ii) the large amounts of data mean that any “noise” stemming from random errors is likely to be cancelled,³³ and what remains is a clear signal that an AI model can learn. Consequently, it is not surprising that tax administrations use AI systems for various purposes.

The application of AI by tax authorities can take various forms. For instance, 80% of the 12,000 complaints that the Dutch Tax and Customs Administration receives annually are answered automatically through natural language processing (NLP) algorithms. This algorithmic system costs the equivalent of two data analyst technicians working for two

29. On the various uses of AI in the tax sector, see C. Milner & B. Berg, *Tax Analytics: Artificial Intelligence and Machine Learning—Level 5* p. 5 (PwC 2017), available at <https://www.pwc.no/no/publikasjoner/Digitalisering/artificial-intelligence-and-machine-learning-final1.pdf>.
30. I.e., data from real time VAT reporting, VAT payment and re-payment claims, automatic exchange of tax information among countries, financial accounting, asset accounting and payroll accounting, and tax returns.
31. A. Collosa, *Use of Big Data in Tax Administrations*, CIAT (1 Sept. 2021), available at <https://www.ciat.org/use-of-big-data-in-tax-administrations/?lang=en> (accessed 3 Mar. 2022).
32. OECD, *Advanced Analytics for Better Tax Administration: Putting Data to Work* 21 (OECD 2016), available at <https://doi.org/10.1787/9789264256453-en>.
33. I.e., unwanted data items, features or records that are detrimental to explaining the feature, or the relationship between feature and target. J.C. Andrade, *Dealing with Noisy Data in Data Science*, Judith Chao Andrade blog (29 Apr. 2019), available at <https://yusout.com/2019/04/29/dealing-with-noisy-data-in-data-science/> (accessed 3 Mar. 2022).

weeks. Normally, responding to these complaints would take employees of the Dutch Tax and Customs Administration hundreds of working days to read and respond to hundreds of thousands of unstructured messages.³⁴ Similarly, the Finnish tax administration relies on AI-based software to analyse and monitor the tax collection email inbox, which receives messages from between 1,500 and 3,000 monthly tax collection stakeholders, such as courts of law and inheritance administrators. The AI algorithms transfer these emails and their attachments directly to the work queue of the GenTax programme for processing, which saves almost two years of manual work by an employee of the Finnish tax administration.³⁵ It is worth noting that the accuracy of NLP models has consistently increased over the past few years, with some of the models exceeding the capabilities of human specialists, such as lawyers. The LawGeex AI system has an impressive 94% average accuracy rate in a task intended to check a set of NDAs for potential legal issues, surpassing the 85% accuracy rate achieved by lawyers on average.³⁶ Although this is a very narrow example of how well these models can perform, it is important to note that there are and will be more and more instances of models exceeding human performance in certain low-level tasks.

Apart from supporting the employees of tax administrations in such mundane tasks, one of the biggest contributions made by AI systems is in the detection of tax fraud, risk profiling and auditing (deciding where and when to conduct tax inspections).³⁷ The Norwegian Tax Authorities use data analytics and machine learning techniques to improve efficiency when selecting cases for VAT inspection. The percentage of successful audits is almost double that achieved with manual processes.³⁸ In Japan, an AI system was designed to analyse financial results and report on documents as well as voice recordings of company managers explaining their business performance. Since its implementation, the number of companies suspected of engaging in tax evasion has dropped,³⁹ and close to a quarter of the tax audits carried out in France in 2019 were supported by AI data mining systems. Merging this technology with human tax inspectors resulted in the collection of EUR 11 billion, a 30% increase over 2018.⁴⁰ This success prompted French regulators to issue a decree on 13 February 2021 allowing French tax authorities to search various online platforms for evidence of tax fraud. They will be able to use publicly available personal data gleaned from social media, platforms like Airbnb or classifieds ads on the website *Le Bon Coin* to feed AI

-
34. OECD, *Tax Administration 2019: Comparative Information on OECD and Other Advanced and Emerging Economies* 55 (Tax Administration, OECD 2019), available at <https://doi.org/10.1787/74d162b6-en>.
 35. PwC, *Tax Disruption Report 2021/2022*, p. 36, available at https://www.pwc.ch/en/publications/2022/tax-disruption-report-2021-2022_EN_web.pdf (accessed 17 Aug. 2022).
 36. *Comparing the Performance of Artificial Intelligence to Human Lawyers in the Review of Standard Business Contracts*, LawGeex Artificial Intelligence (Feb. 2018), available at <https://images.law.com/contrib/content/uploads/documents/397/5408/lawgeex.pdf>.
 37. A survey of the OECD's Forum on Tax Administration (FTA) of 16 tax administrations (Australia, Canada, China, Finland, France, Ireland, Malaysia, Mexico, Netherlands, New Zealand, Norway, Singapore, Sweden, Switzerland, United Kingdom, United States) shows that audit case selection is the principal application of advanced analytics techniques. OECD, *supra* n. 32, at 20. See more: M. Zackrisson, A. Bakker & J. Hagelin, *AI and Tax Administrations: A Good Match*, 74 Bull. Intl. Taxn. 10, sec. 3 (2020), Journal Articles & Opinion Pieces IBFD.
 38. F. Barraza Luengo et al., *Manual Sobre Gestión de Riesgos de Incumplimiento Para Administraciones Tributarias* p. 541 (R. Zambrano & I.G. Arias Esteban eds., Centro Interamericano de Administraciones Tributarias (CIAT) 2020).
 39. T. Narumi, *AI'll Have that: Japan Tax Agency to Use Automated System to Catch out Evading Firms*, Mainichi Daily News (19 Feb. 2020), available at <https://mainichi.jp/english/articles/20200219/p2a/00m/0na/028000c> (accessed 3 Mar. 2022).
 40. A. Collosa, *Artificial Intelligence Applied to Auditing*, CIATblog (13 Oct. 2020), available at <https://www.ciat.org/ciatblog-artificial-intelligence-applied-to-auditing/?lang=en> (accessed 3 Mar. 2022).

data mining algorithms for their investigations, most notably when people lie about their tax domicile.⁴¹ In the United Kingdom, the HMRC has already been relying on data from social media since 2010, cross-checking it with data from a plethora of other databases⁴² to detect fraudulent activity. This task is performed by an AI system known as Connect, which was introduced in 2010 and developed with the help of BAE Systems.⁴³ Connect searches for correlations between declared income (e.g. the tax records of businesses and individuals) and lifestyle (e.g. social media, the Charity Commission, flight sales and passenger information), comparing it with multivariate statistical models using AI. Enquiries are currently triggered in the vast majority of tax compliance risk cases (over 90%) by information and analyses generated by Connect. Although Connect allegedly cost GBP 100 million, it is worth that price insofar as it has helped recover more than GBP 3 billion in taxes since its launch date.⁴⁴

These examples showcase the wide-ranging applicability of AI systems in taxation. A recent trend – the rapid and steadily increasing global spread of AI systems in public administrations, including tax administrations – would seem to be ongoing.⁴⁵ International and supra-national organizations, notably the OECD, explicitly encourage tax administrations from all over the world to implement AI systems in support of their tasks as part of the transformation from Tax Administration 2.0 to 3.0.⁴⁶ A recent survey conducted by the OECD’s Forum for Tax Administration (FTA), involving tax administrations in 59 advanced and emerging economies, reveals the increasing use of AI systems: almost 75% of the surveyed tax administrations report that “they are using or that they are in the implementation phase for the future use of cutting-edge techniques to exploit data in ways that reduce the need for human intervention”.⁴⁷

The maximization of the use of AI systems by tax administrations around the world would seem to be a necessity rather than a choice due to personnel shortages in that public domain relative to the volume of work⁴⁸ and the intertwined digital-physical reality of the tax

41. J. Darmanin, *France Starts Scouring Social Media to Catch Tax Fraudsters*, POLITICO (19 Feb. 2021), available at <https://www.politico.eu/article/france-starts-scraping-social-media-to-catch-tax-fraudsters/> (accessed 3 Mar. 2022).

42. The databases include: tax returns (including VAT, PAYE, income tax and corporation tax returns), bank accounts and pensions, credit reference agencies, credit and debit card accounts, online payment providers such as PayPal, foreign tax jurisdictions (including treaties and automatic exchange agreements) and the common reporting standard, government agencies such as companies house, the land registry and the border agency, Property websites such as Zoopla and Rightmove, Amazon, eBay, Gumtree and similar sales websites, Google Street View, council tax records, DVLA records, DWP records, electoral roll, insurance companies, charities commission, flight sales and passenger information.

43. BAE Systems plc is a British multinational arms, security and aerospace company based in London, England. The company is the largest defence contractor in Europe and among the world’s largest defence companies. See <https://www.baesystems.com/en/our-company/about-us>.

44. J. Sanghrajka, *HMRC’s Connect Computer and Investigations*, Taxation 4752 (14 July 2020), available at <https://www.taxation.co.uk/articles/hmrc-s-connect-computer-and-investigations> (accessed 3 Mar. 2022).

45. G. Misuraca & C. van Noordt, *Artificial Intelligence in Public Services: Overview of the Use and Impact of AI in Public Services in the EU* (AI Watch 2020), available at <https://data.europa.eu/doi/10.2760/039619>.

46. A recent OECD report used the term “Tax Administration 3.0” to mark a new stage of digitalization, in which taxation is moved closer to taxable events through built-in, automated compliance mechanisms and the interconnection between tax authority systems and the systems taxpayers use to run their businesses. See OECD, *supra* n. 8, at p. 7.

47. OECD, *Tax Administration 2021: Comparative Information on OECD and Other Advanced and Emerging Economies* 27 (Tax Administration, OECD 2021), available at <https://doi.org/10.1787/cef472b9-en>.

48. N. Mehdiyev et al., *Explainable Artificial Intelligence (XAI) Supporting Public Administration Processes – On the Potential of XAI in Tax Audit Processes*, Wirtschaftsinformatik 2021 Proceedings sec. 2.2 (2021).

world.⁴⁹ At the same time, the significant increase in the use of AI by tax administrations dramatically disrupts the tax profession. Tax advisers and taxpayers must adapt to the AI capabilities of tax administrations by becoming AI savvy enough to implement relevant technologies properly in due course.⁵⁰ So far, however, the attention of people engaged in the development and use of AI systems in tax law appears to be predominantly focused on the efficiency and accuracy of such systems. As sections 2.1. to 2.3. show, their attention needs to be broadened to avoid the negative consequences of the use of AI systems in tax law.

2.1. Motivation for implementing XAI systems in tax law

The strongest attention to AI systems appears to be given to their contribution to increasing the efficiency and accuracy of tax people in both the private and public sectors. Nevertheless, developers and users should not limit their focus to these features.⁵¹ Even the most efficient and accurate tax law AI system will still be plagued by the occasional prediction error, which could trigger serious consequences for the taxpayer. AI systems could also bring about harm due to various other reasons, such as inadequate design, reliance on biased data or the application thereof for an illegitimate purpose.⁵² To prevent AI from having such undesirable consequences – and to provide some redress should those consequences come to pass – scholar and practitioners have developed various approaches to ensure algorithmic accountability.⁵³

XAI can contribute to accountability in the tax domain by providing *explainability* to the AI systems in the tax sector, that is, by ensuring that the functioning and the outputs of these systems can be explained in a way that a human observer can understand. Without such an explanation, taxpayers might be concerned about whether automated decisions yield fair and non-discriminatory results.⁵⁴ And such a concern would be warranted, as scholars and civil society groups have uncovered various situations in which AI systems introduce or compound existing harm, not just in the tax domain but in other public activ-

-
49. In the modern economy, the digital and real worlds are interconnected. Therefore, it is natural in such an economy that human reality combines with the virtuality of the algorithm, which should be taken into account by tax policymakers and tax administrations across the globe while designing and executing tax law: I. Cugusi, *Prospects for Taxation of the Digital Economy between “Tax Law and New Economy” and “Tax Law of the New Economy”*, 12 *World Tax J.* 4, sec. 9 (2020), Journal Articles & Opinion Pieces IBFD.
50. A survey of the EY Quantitative Economics and Statistics (QUEST) group and Wakefield Research on a group of 100 of the largest multinational companies with median USD 12 billion+ in revenue shows that more than 73% of the tax functions surveyed intend to hire more data and technology professionals; correspondingly, only 25% or less anticipate hiring more accountants and lawyers. K. Neidhardt, R. Doll & A. Lee, *The Intelligent Tax Function. 2020 Global Tax Technology and Transformation Survey Highlights* 8 (EY 2020).
51. B. Kuźniacki, *How constitutional principles pave the way to eXplainable AI in tax law*, Digital Constitutionalism: The Future of Constitutionalism (23 Jan. 2022), available at <https://digi-con.org/how-constitutional-principles-pave-the-way-to-explainable-ai-in-tax-law/>.
52. On the various kinds of harm stemming from AI technologies, see N.A. Smuha, *Beyond the Individual: Governing AI’s Societal Harm*, 10 *Internet Policy Review* 3 (2021).
53. For an overview of such approaches, see M. Wieringa, *What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability*, in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* ’20)*, Association for Computing Machinery 27 Jan. 2020), available at <https://doi.org/10.1145/3351095.3372833>.
54. *Algorithmic Impact Assessment Tool*, Canada.Ca (22 Mar. 2021), available at <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html> (accessed 3 Mar. 2022).

ities such as health care, criminal justice or public education.⁵⁵ By rendering tax AI systems explainable, tax authorities ensure that the decisions made with their involvement can be contestable through the same channels used to contest human decisions, thus affording taxpayers the same legal protections they enjoy in respect of a decision made by a human.⁵⁶

This protection is needed because automated decisions can have a considerable impact on taxpayers' rights. A strong example comes from the Netherlands, in which the so-called childcare benefits scandal (*toeslagenaffaire*) has unveiled systematic discrimination against tens of thousands of welfare recipients.⁵⁷ In 2021, it came to light that Dutch tax authorities were using a risk assessment system that infringed the fundamental rights of an estimated 35,000 welfare recipients.⁵⁸ By wrongfully accusing these recipients of welfare fraud, the system inflicted considerable financial harm on innocent people, sometimes even causing them to suffer health issues.⁵⁹ Once these effects were revealed, the backlash prompted Mark Rutte, Prime Minister of the Netherlands, to resign.⁶⁰ Problems of this sort will only become more frequent unless proper actions are taken promptly by entities involved in designing, deploying, applying and supervising AI systems to ensure their explainability in ways relevant to their stakeholders.⁶¹

The European Union has shown considerable effort in guiding the creation of AI systems in the draft AI regulation recently published by the European Commission.⁶² It is not the final version and will require significant changes before its implementation, but it is a step in the right direction, which should provide greater legal certainty for the implementation of AI systems.

2.1.1. *The strong and compulsory requirement to explain tax decisions*

In the private sector, where the AI-supported decision-making process takes place in business-to-consumers (B2C) and business-to-business (B2B) relationships, decisions require explanation only when a specific law requires it (weak/narrow XAI motivation). In B2C relations, explainability is required under data protection law⁶³ and consumer protection

-
- 55. See, inter alia, Hadwick & Lan, *supra* n. 6; M. Whittaker et al., *AI Now Report 2018* 12–24 (AI Now Institute Dec. 2018); R. Xenidis, *Tuning EU equality law to algorithmic discrimination: Three pathways to resilience*, 27 *Maastricht Journal of European and Comparative Law* 6. (2020).
 - 56. R. Leenes, E. Kosta & I. Kamara, *Introduction to Research Handbook on EU Data Protection Law*, in *Research Handbook on EU Data Protection Law* (E. Kosta & R. Leenes eds., Edward Elgar Publishing 2022); Sarra, *supra* n. 11.
 - 57. Amnesty International, *Xenophobic Machines: Discrimination through Unregulated Use of Algorithms in the Dutch Childcare Benefits Scandal* (Amnesty International 2021) [hereinafter *Xenophobic Machines* (2021)]; Kuźniacki, *supra* n. 51.
 - 58. Tweede Kamer der Staten-Generaal [Netherlands' 2nd Parliamentary Chamber], *Eindverslag – Parlementaire ondervragingscommissie Kinderopvangtoeslag "Ongekend Onrecht"*, pp. 3-7 (The Hague, 17 Dec. 2020).
 - 59. Here, again, the Dutch Childcare Benefits Scandal is so far the most illustrative. O. Akinici, *Sibel werd ziek van de Toeslagenaffaire en kan en wil er niets meer over horen*, *Trouw* (27 Nov. 2020), available at <https://www.trouw.nl/gs-b0b63a37> (accessed 3 Mar. 2022).
 - 60. J. Henley, *Dutch government resigns over child benefits scandal*, *The Guardian* (15 Jan. 2021), available at <https://www.theguardian.com/world/2021/jan/15/dutch-government-resigns-over-child-benefits-scandal> (accessed 17 Aug. 2022).
 - 61. Mehdiyev et al., *supra* n. 48, at sec. 2.2; Hadwick & Lan, *supra* n. 6, at sec. 2.2.
 - 62. European Commission, Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative Acts, COM(2021) 206 final (21 Apr. 2021) [hereinafter *AI Act* (Commission Proposal) (2021)].
 - 63. On explanation in data protection law, see sec. 4.1.1.

law,⁶⁴ while, in the B2B sphere, this obligation ensues from laws on promoting fairness and transparency for business users of online services.⁶⁵ There are also legal explainability requirements related to the use of algorithmic trading tools in the financial and insurance sectors.⁶⁶

As long as AI tax-oriented systems are used only in B2C or B2B relations, explainability requirements will remain weak and narrow. However, the vast majority of relationships in the tax domain are between tax authorities and taxpayers. These relationships are thus in the government-to-citizen sphere (G2C, or more precisely, the tax authority-to-taxpayer (TA2T) sphere), in which tax authorities are obliged to justify their decisions. In consequence, using AI in G2C relations creates stronger legal obligations of explainability than in B2C and B2B relationships.⁶⁷ In particular, the decisions of tax authorities must comply with the principle of formal motivation. To comply with this principle, all factual and legal grounds on which the decision is based should be mentioned and explained by the tax authorities.⁶⁸ The justification for such decisions must be clear and precise and reflect the real motives behind the decision.⁶⁹ Such justification is impossible if tax decisions are not explainable, e.g. when they follow the recommendations of non-explainable AI systems. Explainable AI systems are considered more trustworthy and reliable, which leads to their wider adoption regardless of the environment in which they are deployed.⁷⁰ Furthermore, there is a fundamental relationship between XAI and transparency, fairness and accountability.⁷¹ As trust increases, it allows model creators to gather more data, which yields better and more reliable results for the predictive systems. Thus, explainability should be considered along error metrics when deploying AI models, especially in human-in-the-loop architecture. An increasing number of professionals advocate using XAI to facilitate AI readiness and higher AI utilization in professional services sectors.⁷²

The bar for explanation is raised proportionally to the level of discretionary power enjoyed by the tax authorities.⁷³ In simple cases in which decisions are made based on objective conditions – for instance, the participation exemption after meeting the conditions of holding

-
- 64. In the European Union, *see* art. 6(a) of Directive 2011/83 of the European Parliament and of the Council of 25 October 2011 on consumer rights, OJ L 304, 22 Nov. 2011, in connection with recital 22 of Directive 2019/2161 of the European Parliament and of the Council of 27 November 2019 on better enforcement and modernization of EU consumer protection rules, OJ L 328, 18 Dec. 2019.
 - 65. In the European Union, *see* art. 5 in connection with recital 24 of Regulation 2019/1150 of the European Parliament and of the Council of 20 June 2019 on promoting fairness and transparency for business users of online intermediation services, OJ L 186, 11 July 2019.
 - 66. In the European Union, *see* art. 17(2) Directive 2014/65 of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments, OJ L 173, 12 June 2014.
 - 67. Bibal et al., *supra* n. 17, at sec. 2.2.
 - 68. *Id.*, at sec. 2.2.1.
 - 69. Art. 41 Charter of Fundamental Rights of the European Union (the EU Charter), OJ C 326/391, 26 Oct. 2012. For national law, *see*, for example, the BE: Law of 29 July 1991 on the formal motivation of administrative decisions.
 - 70. D.B. David, Y.S. Resheff & T. Tron, *Explainable AI and Adoption of Financial Algorithmic Advisors: An Experimental Study*, Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (2021).
 - 71. D.D. Shin. *The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI*, 146 International Journal of Human-Computer Studies (2021).
 - 72. L. Kimbell et al., *AI Readiness: A collaborative design toolkit for professional services firms* (Oxford Brookes University & Practice Management International LLP 2021).
 - 73. J.-L. Autin, *La motivation des actes administratifs unilatéraux, entre tradition nationale et évolution des droits européens*, Revue française d'administration publique, 137-138 (2011), available at <https://doi.org/10.3917/rfap.137.0085>.

10% of shares in another company for two years⁷⁴ – the required explanation is weaker. The tax authority only has to explain in its decision that the conditions required by the applicable legal text are fulfilled. By contrast, whenever tax authorities have wide discretion in rendering decisions, as in cases regarding tax avoidance regulated by general anti-abuse clauses,⁷⁵ they must meticulously explain their choices and legal reasoning.

Furthermore, additional explainability requirements may apply whenever the decision-making process of tax authorities is automated.⁷⁶ This shows that explanation is always required when AI systems are used to render decisions in TA2T relationships. If such decisions are automated and stem from a significant discretionary power of the tax authorities, the highest possible level of XAI should be ensured.

In the context of the above, a specific distinction should be made between the various explainability levels that can be enforced whenever public administration organizations rely on models in their decision-making. Scholars have suggested four levels of machine learning explainability.⁷⁷ The first is a weaker level that reveals only the model’s main features. The second level is focused on providing all the model’s features. The third level intends to provide the set of features that led to the specific outcome, whereas the fourth level involves sharing the whole model. The last level could serve the purpose of explainability as it would allow practitioners to evaluate predictions and thoroughly test the models deployed in the legal environment, thus allowing the scientific and business communities to improve the models. In addition to the above, there is a need to adjust the legal procedures before implementing any specific AI solution. The authors advocate for transparency in the use of models so the subjects of AI-supported decisions are fully aware of who or what is making the decisions. This includes designing specific procedures to appeal decisions more quickly when there are no humans in the loop.

2.1.2. *No taxation without representation – A constitutional requirement for XAI in tax law*

In section 3.1., the authors argue that the limits for constitutionally compliant use of AI systems in tax law are set by the constitutional requirement that taxes only be levied employing a statutory law in a precise and predictable manner (the principle of “no taxation without representation”), as arises from the rule of law in the area of taxation.⁷⁸ Notably, a sufficient

74. In the European Union, see art. 3 Council Directive 2011/96/EU of 30 November 2011 on the common system of taxation applicable in the case of parent companies and subsidiaries of different Member States, OJ L 345, 29 Dec. 2011.

75. In the European Union, see art. 6 (general anti-avoidance rule, GAAR) Council Directive (EU) 2016/1164 of 12 July 2016 laying down rules against tax avoidance practices that directly affect the functioning of the internal market, OJ L 193, 19 July 2016, and art. 7(1) (principal purposes test, PPT) Multilateral Convention to Implement Tax Treaty Related Measures to Prevent BEPS (MLI). For the potential and risks related to an application of AI systems to such rules, see B. Kuźniacki & K. Tyliński, *Identifying the Potential and Risks of Integration of AI to Taxation: The Case of General Anti-Avoidance Rule*, in *Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law* (G. D’Agostino, A. Gaon & C. Piovesan eds., Carswell 2021).

76. E.g., FR: Code on the relationships between the public and the administration, art. R. 311-3-1-2.

77. A. Bibal et al., *Impact of Legal Requirements on Explainability in Machine Learning*, in *ICML Workshop on Law and Machine Learning* (10 July 2020).

78. J. Hattingh, *The Multilateral Instrument from a Legal Perspective: What May Be the Challenges?*, 71 Bull. Intl. Taxn. 3/4, sec. 2 (2017), Journal Articles & Opinion Pieces IBFD, with reference to the late Lord Bingham’s articulation of the tenets of the rule of law as depicted in T.H. Bingham, *The Rule of Law* (Allen Lane 2010).

level of AI system explainability constitutes a building block towards the compatibility of its use with the constitutional principles of legal certainty and statutory taxation that embodies the principle of legal precision and predictability of taxation. The explainability of AI systems is a step towards the constitutionality of its use, but it cannot ensure such compatibility on its own because even a perfectly explainable AI system could nevertheless be arbitrarily used by tax authorities, thus rendering taxation unpredictable. However, without an explainable AI system – either due to tax or trade secrecy and/or technological complexity – it would not be possible to determine whether the use of AI systems contradicts constitutional principles, e.g. due to its arbitrary or discriminatory character.

The constitutional principle of “no taxation without representation” can be seen as a constitutional requirement for XAI in tax law. Tax law and its execution must at least be subject to express and clear legal safeguards for protecting taxpayer rights, meaning that the execution of tax law must be maximally transparent and precise for taxpayers. Accordingly, there is no room for using unexplainable AI systems for taxation purposes, including tax-related practices such as identifying which taxpayers will be subjected to a tax inspection. In other words, the essential feature of an AI system in tax matters must, in light of constitutional principles, be its explainability. If AI systems in tax law are deprived of explainability, they do not pass the proportionality test when it comes to ensuring the effective prevention of tax fraud, insofar as they are not necessary to achieve that purpose and remain too invasive to the fundamental rights of taxpayers.⁷⁹

2.2. Tax secrecy and trade secrecy as challenges for explanation

The confidentiality of information obtained in the course of a tax official’s duties (tax secrecy) is a crucial issue for tax administrations worldwide.⁸⁰ Although the scope of tax secrecy varies from country to country,⁸¹ there are many alarming examples involving the design and functioning of AI systems used by tax administrations. Scholars from South America have identified a lack of procedures for taxpayers to obtain information about the functioning of AI systems in tax audits to detect tax fraud and measure the risks of compliance with tax obligations by Latin American tax administrations, such as those of Chile, Peru, Brazil, Colombia, Ecuador, Argentina and Mexico.⁸² In Poland, the Clearing House System (*System Teleinformatyczny Izby Rozliczeniowej*, STIR) establishes a risk indicator for each entrepreneur, which is calculated using secret AI algorithms based on criteria used by the financial sector to combat VAT fraud.⁸³ The National Revenue Authority (NRA) can block an entrepreneur’s bank account or cancel an entrepreneur’s VAT registration based on a risk indicator provided by the STIR. The AI algorithms used to determine such risk are never disclosed to the taxpayer or the courts, even during litigation between the taxpayer and the NRA, which can trigger serious issues with the compatibility of STIR with fundamental

79. See sec. 3.

80. *Confidentiality of Tax Information*, Tax Inspectors Without Borders, available at <http://www.tiwb.org/get-involved/experts/confidentiality-of-tax-information/> (accessed 3 Mar. 2022).

81. See E. Kristofferson et al. eds., *Tax Secrecy and Tax Transparency: The Relevance of Confidentiality in Tax Law* (Peter Lang 2013).

82. The most frequently used AI tools have been clustering algorithms of the self-organizing feature map (SOFM) type, K-means, neural networks, and Bayesian networks: A. Faúndez-Ugalde, R. Mellado-Silva & E. Aldunate-Lizana, *Use of Artificial Intelligence by Tax Administrations: An Analysis Regarding Taxpayers’ Rights in Latin American Countries*, 38 Computer Law & Security Review, sec. 6 (Sept. 2020), available at <https://doi.org/10.1016/j.clsr.2020.105441>.

83. PL: Tax Ordinance Act of 1997, sec. III.B (art. 119zg and others).

taxpayer rights.⁸⁴ Similarly, AI systems used by the French⁸⁵ and German⁸⁶ tax administrations for automation-based risk assessment purposes (e.g., risk-scoring tax returns for audits) are hidden from the taxpayer behind a veil of tax secrecy.⁸⁷ As a result, taxpayers cannot ascertain the circumstances under which a case will be classified as high or low risk, whether before or after the tax assessment. Nor are they afforded an even marginal description of the basic functioning of the AI algorithms.⁸⁸

Tax secrecy in connection with AI systems is justified by the need to ensure the uniformity and effectiveness of tax inspections. Presumably, if taxpayers were aware of the risk-assessment criteria or the algorithms' modes of functioning, they could circumvent these filters and thus escape tax inspection despite being involved in tax fraud or non-compliance.⁸⁹ In other words, by knowing the details of AI systems used by tax administrations, taxpayers could reverse engineer their functioning and know exactly how to prevent themselves from being targeted for tax inspection. Tax secrecy does not prevent large tax consultancy firms, before such reverse engineering takes place, from systematically evaluating cases in which they have been involved to provide insights into the system's functioning in exchange for monetary remuneration.⁹⁰ This shows that tax secrecy does not affect all taxpayers equally. Taxpayers who are clients of large tax consultancy firms can bear the costs of anticipating the functioning of AI systems used by tax administrations, while other taxpayers cannot. This compromises both the rationale of tax secrecy and equality of access to information, thereby exacerbating the distrust in tax authorities. The implementation of AI systems by tax administrations should not be an excuse for extending tax secrecy in a way that reduces the bargaining power of individual taxpayers in the social contract between them and the administrations,⁹¹ thus rendering the use of XAI in tax law an impossible goal to achieve broadly and inclusively.

To some extent, trade secrecy plays a similar role to tax secrecy, i.e. preventing the discovery of the inner functioning of AI systems, thus potentially allowing minimal changes to the input data to bring about different outcomes, for example, different classifications in the

-
84. M. Papis-Almansa, *The Polish Clearing House System: A 'Stir'ring Example of the Use of New Technologies in Ensuring VAT Compliance in Poland and Selected Legal Challenges*, 28 EC Tax Review 1 (1 Feb. 2019); A. Bal, *Ruled by Algorithms: The Use of 'Black Box' Models in Taxation*, 95 Tax Notes International 12 (2019); M. Rojszczak, *Compliance of Automatic Tax Fraud Detection Systems With the Right to Privacy Standards Based on the Polish Experience of the STIR System*, 49 Intertax 1, pp. 46-49 (2021).
85. *Fraud Targeting and Valorisation of Requests [Ciblage de la Fraude et Valorisations des Requêtes, CFVR]. Arrêté du 12 novembre 2019 modifiant Arrêté du 21 Février 2014 portant création par la direction générale des finances publiques d'un traitement automatisé de lutte contre la fraude dénommé "ciblage de la fraude et valorisation des requêtes"* [Arrest amending the system of fraud targeting and valorization of requests, amended 2021].
86. *Gesetz zur Modernisierung des Besteuerungsverfahrens* [Taxation Modernization Act], BGBl I, p. 1679 T. (18 July 2016), sec. 88(5) *Abgabenordnung* (amended 2021).
87. A. Léchenet, *French Tax Authority Pushes for Automated Controls despite Mixed Results*, AlgorithmWatch (23 Nov. 2020), available at <https://algorithmwatch.org/en/france-tax-automated-dgfp/> (accessed 3 Mar. 2022); F.J. Marx, *Der Einsatz von Risikomanagementsystemen nach § 88 Abs. 5 AO als Kernelement der Modernisierung des Besteuerungsverfahrens*, 9 UBG 6, 363 (1 June 2016), available at <https://doi.org/10.9785/ubg-2016-090608>.
88. Hadwick & Lan, *supra* n. 6, at sec. 3.3.
89. Braun Binder, *supra* n. 9, at pp. 298-299.
90. N. Braun Binder, *Ausschließlich automationsgestützt erlassene Steuerbescheide und Bekanntgabe durch Bereitstellung zum Datenabruf*, Deutsche Steuer-Zeitung (DStZ) pp. 526-535 (2016).
91. Hadwick & Lan, *supra* n. 6, at sec. 3.3.

insurance, loan or tax risk assessment systems used by various firms.⁹² Beyond that, trade secrecy's goal is to protect the data model functionality of AI systems from being stolen.⁹³ Considering that such models may have been developed by companies through many years of research in a specific field, trade secrecy can be justified. Yet, just as in the case of tax secrecy, a reduction of the vulnerability to the reverse-engineering of AI systems used in the field of tax law under the guise of trade secrecy cannot be allowed to incapacitate such systems of explainability. It becomes necessary to strike an adequate balance between the various rights and values at stake.⁹⁴

To strike the right balance between the values and rights of tax administrations and developers of AI systems protected by tax and trade secrecy, on the one hand, and those of taxpayers, on the other, a new set of rules and practices needs to be developed. In that regard, it has been argued that algorithm auditing can be useful for evaluating the rules and practices adopted in any given AI system.⁹⁵ These audits would provide knowledge about the inner workings of the system under analysis, creating the conditions for the democratic scrutiny of how tax authorities automate their operations. However, the literature points out that algorithm audits have several limitations, such as being resource-intensive – and thus infeasible for low and middle-income countries that might otherwise benefit from AI⁹⁶ – or the risks of regulatory capture.⁹⁷ Although this article cannot hope to exhaustively debate the proper contours of algorithm audits in the tax domain, it is important to highlight a few essential conditions for any approach to auditing.

In the case of litigation between a taxpayer and the tax authorities, judges will need information about the relevant AI system. A balanced approach in that regard should be able to differentiate between the information received by the court and the taxpayer. The former might receive more granular information, including the data used to train the AI model

-
- 92. Such attempts are called adversarial attacks and they have happened, for example, regarding the computer vision systems of autonomous vehicles. A minimal change in a stop signal, imperceptible to the human eye, led vehicles to misconstrue it as a 45-mph signal. R.S.S. Kumar et al., *Legal Risks of Adversarial Machine Learning Research* (ICML 2020).
 - 93. R.S.S. Kumar et al., *Threat Taxonomy - Failure Modes in Machine Learning*, Security documentation, Microsoft (11 Nov. 2019), available at <https://docs.microsoft.com/en-us/security/engineering/failure-modes-in-machine-learning> (accessed 9 Jan. 2021).
 - 94. It is important to strike a pragmatic balance between disclosure obligations and legitimate commercial interests: H. Bloch-Wehba, *Transparency's AI Problem* 8 (Knight First Amendment Institute at Columbia University 2021), available at <https://knightcolumbia.org/content/transparencys-ai-problem> (accessed 26 Feb. 2022). Trade secrecy should protect but not overprotect intellectual property, thereby failing to serve the purposes of either trade secret law or privilege law. "Withholding information from the accused because it is a trade secret mischaracterizes defense advocacy as a business competition": R. Wexler, *Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 Stanford Law Review 5, p. 1429 (2018).
 - 95. I.D. Raji et al., *Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing*, in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (FAT* '20, Association for Computing Machinery 2020), available at <https://doi.org/10.1145/3351095.3372873>; A.E. Waldman, *Power, Process, and Automated Decision-Making*, 88 Fordham Law Review 2 (2019), as they demand financial and technical resources that might not be available to low and middle-income countries that might otherwise benefit from AI: C. Abungu, *Algorithmic Decision-Making and Discrimination in Developing Countries*, 13 Journal of Law, Technology, & the Internet 1 (2022). There is also the risk of regulatory capture: J. Laux, S. Wachter & B. Mittelstadt, *Taming the Few: Platform Regulation, Independent Audits, and the Risks of Capture Created by the DMA and DSA*, 43 Computer Law and Security Review (2021). Consequently, further study will be necessary to design suitable mechanisms for auditing tax AI systems.
 - 96. Abungu, *supra* n. 95.
 - 97. Laux, Wachter & Mittelstadt, *supra* n. 95.

and the weight assigned to each input factor⁹⁸ essential to providing the output that led to the decision of tax authorities in the taxpayer's case⁹⁹ to enable a guided assessment of the system's performance, probably with the use of expert testimony. This judicial oversight would accordingly implement a kind of human-in-the-loop AI control model. The information obtained by taxpayers does not need to be quite granular. Instead, the taxpayer might receive, upon request, an explanation from the tax authority detailing the decision and why it took precedence over plausible alternatives. This differentiation in the scope and content of the explanation of AI systems is justified by the obligation of the court to deliver a judgment backed up by reasonable and persuasive arguments in support thereof. Taxpayers bear no such burden. For them, it is sufficient to understand why the AI system led to outcome X rather than outcome Y (counterfactual) and that such an outcome is or is not predicated on unfair or discriminatory elements.

The main task of XAI in tax law vis-à-vis tax and trade secrecy should be to ensure that both judges and taxpayers receive information that will clearly and precisely shed light on all elements of the functioning of an AI system decisive, not merely tangential, to the case of the taxpayer (or other addressees). Focusing on only some, rather than all, vital elements of an AI system may lead both judges and taxpayers astray.¹⁰⁰ For example, and by way of analogy regarding the use of AI systems in criminal law, the legal literature has rightly criticized the Wisconsin Supreme Court in the United States for confirming a conviction based on a risk assessment report derived from the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) in the judgment in *Loomis v. Wisconsin* (13 July 2016).¹⁰¹ Critics have pointed out that the court emphasized the data quality standard, which set a very low bar for understanding COMPAS's predictive risk assessment.¹⁰² Data accuracy is an important quality for purposes of data science, but it is still only one among many various qualities an algorithmic system must have. In particular, data accuracy does not provide us with information about how the algorithm delivers results or deals with input data.¹⁰³ In fact, the court failed to assess various aspects of COMPAS's predictive algorithm, such as

98. This information would appear to be crucial to understanding the reason for a given decision recommended by the outcome of an AI system: A. Chander, *The Racist Algorithm?*, 115 Michigan Law Review 6, pp. 1023-1024 (2017).

99. For example, the Council for Transparency in Chile, in *Zubizarreta v. Servicio de Impuestos Internos*, ordered the tax entity to provide the taxpayer with information on the Integrated Taxpayer Information System (SIIC, currently replaced by the Comprehensive Tax Compliance System – SICT) on aspects that affected its characterization. Faúndez-Ugalde, Mellado-Silva & Aldunate-Lizana, *supra* n. 82, at sec. 4.

100. For an example, see A.L. Washington, *How to Argue with an Algorithm: Lessons from the Compas-ProPublica Debate*, 17 Colorado Technology Law Journal 1 (2018).

101. US: Wis. SC, 13 July 2016, *Loomis v. Wisconsin*, 881 N.W.2d 749 (Wis. 2016), cert. denied, 137 S.Ct. 2290 (2017).

102. Id. See also H.-W. Liu, C.-F. Lin & Y.-J. Chen, *Beyond State v Loomis: Artificial Intelligence, Government Algorithmization and Accountability*, 27 International Journal of Law and Information Technology 2 (2019), available at <https://doi.org/10.1093/ijlit/eaz001>; Washington, *supra* n. 100; M.B. Zafar et al., *Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment*, in *Proceedings of the 26th International Conference on World Wide Web (WWW '17, IWWWC Steering Committee 3 Apr. 2017)*, available at <https://doi.org/10.1145/3038912.3052660>; K. Freeman, *Algorithmic Injustice: How the Wisconsin Supreme Court Failed to Protect Due Process Rights in State v. Loomis*, 18 North Carolina Journal of Law & Technology 5 (2016).

103. E.g., S. Corbett-Davies et al., *Algorithmic Decision Making and the Cost of Fairness*, in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (ACM Press 2017)*, available at <https://doi.org/10.1145/3097983.3098095>; Y. Wand & R.Y. Wang, *Anchoring Data Quality Dimensions in Ontological Foundations*, 39 Communications of the ACM 11 (1996), available at <https://doi.org/10.1145/240455.240479>.

how it combines sources, weighs variables, establishes ranks and sets category boundaries. Thus, one of the main roles of XAI in tax law is to ensure that judges in tax cases know and understand the entire set of features of the AI system that led to a particular outcome (tax decision). A compressed and higher-level version of such an explanation should be provided to taxpayers. Under no circumstances, however, should judges or taxpayers receive incomplete information concerning an AI system unless the omitted information was of negligible value to the judgment and the taxpayer's understanding of that system. The relevance of information should be duly evaluated by the tax authorities and AI experts (if needed).

Receiving an appropriate explanation from the tax authorities whenever they rely on AI systems would contribute to the taxpayer's trust in the functioning of such systems. It would also be informative for judges when they resolve disputes between stakeholders that use AI systems and the "victims" of the lack of their explainability. In the long run, disputes of that nature will decrease because the number of disputes in that sphere is inversely proportional to the AI systems' degree of explainability.

2.3. Technical challenges and requirements for XAI in the tax sphere

Considering AI systems mainly from a technological angle, explainability contributes to the easier, cheaper, faster and more accurate functionality of such systems by enhancing the identification and elimination of bugs in the software used to run these systems.¹⁰⁴ Although developing XAI can severely impact the pace early on in the process, as it requires additional software development, the explanation may save a considerable amount of time by identifying bugs that otherwise would not be identified until a later stage.¹⁰⁵ The lack of such identification at the mentioned late stages typically causes severe and sometimes even irreversible consequences to various stakeholders, including society as a whole.¹⁰⁶

A proper approach to addressing the challenges stemming from technological constraints requires acknowledging a clear distinction between the various fundamental types of AI and XAI solutions. One approach to XAI is to construct interpretable (transparent) models, i.e. AI systems that operate in a way that an external observer can understand.¹⁰⁷ For example, expert systems in the tax domain make decisions based on explicit rules obtained by modelling the knowledge of tax professionals into a form that a computer can process.¹⁰⁸ Since these rules are derived directly from human knowledge and written in formal language, they can – at least in theory – be assessed by human observers, making it easier to identify errors and, if need be, contest them in court. Consequently, it has been argued that

104. On this point, see Kuźniacki, *supra* n. 51.

105. Cf. in respect of continuous delivery, which involves keeping software ready for release to production, always and at all times: E. Laukkanen, J. Itkonen & C. Lassenius, *Problems, Causes and Solutions When Adopting Continuous Delivery - A Systematic Literature Review*, 82 *Information and Software Technology*, p. 61 (2017), available at <https://doi.org/10.1016/j.infsof.2016.10.001>.

106. For example, Amnesty International reported on the use of an AI system (the risk classification model, SyRI) by the Dutch tax authorities that led to discriminatory results, leaving many people "with mental health issues and stress on their personal relationships, leading to divorces and broken homes": *Xenophobic Machines* (2021), *supra* n. 57, at p. 13. See also Kuźniacki, *supra* n. 51.

107. For an overview of approaches to interpretable AI, see Barredo Arrieta et al., *supra* n. 13, at sec. 3.

108. I. Faina, B. Alturas & F. Almeida, *Expert Systems Within Taxation: Competencies for Employees*, in *2020 15th Iberian Conference on Information Systems and Technologies (CISTI)* (June 2020), available at <https://doi.org/10.23919/CISTI49556.2020.9141022>.

interpretable models should be the only ones used in high-stakes situations¹⁰⁹ such as those involving taxpayer rights.

However, direct interpretability is not always feasible for tax AI. In many cases, there is a direct trade-off between interpretability and the scale on which a model operates, as interpretable models tend to perform significantly worse when dealing with large datasets and become complex to the point that even experts struggle to make sense of them.¹¹⁰ As a result of these limitations, applications of AI to the tax domain – which need to deal with large volumes of taxpayer data and, in applications such as tax fraud detection, achieve high levels of accuracy – tend to rely on machine learning models that are not transparent by design.¹¹¹

Faced with black-box models, XAI adopts post hoc approaches to explanation,¹¹² which build techniques to explain how opaque models such as neural networks arrive at their decisions. In a democratic society, these decisions must reach a variety of stakeholders: tax authorities need to know what is going on within their systems; taxpayers have the right to know how they have been affected by specific decisions; courts must know the relevant factors for judging a case; and so forth. Few, if any, of these stakeholders have a high level of technical knowledge about an AI system’s inner workings, meaning that explanations must provide information in an accessible format.¹¹³ However, each type of stakeholder needs information about different aspects of an AI system. Therefore, explanation models need to be tailored to the needs of the recipients of such explanations.¹¹⁴

Taxpayers – as the subjects/addressees of decisions rendered fully or in part by AI systems – and their tax advisers are primarily concerned with *why* questions: why did a system calculate this value of due tax and not something smaller (or bigger)? Why was a taxpayer classified into a particular category? To answer these questions, a system’s behaviour must be interpreted to let the stakeholder know which information is being used and which factors contribute to a particular outcome.¹¹⁵ For these stakeholders, therefore, local post hoc explanation methods – such as individual conditional explanation (ICE) plots, Shapley values, local interpretable model-agnostic explanations (LIME) or other alternative local surrogate methods, case-based explanations or counterfactual explanations – appear to be most suitable, as they help clarify whether the AI decisions provided in particular cases involving taxpayers are reliable and reasonable.¹¹⁶

Tax authorities are both operators and executors of AI systems for tax purposes. In such roles, they are responsible for making and justifying individual decisions based solely or in part on outcomes rendered by AI systems. They mainly inquire whether the provided AI decisions are reliable and reasonable. Accordingly, their principal questions involve: (i) *what* the AI system is doing, i.e. whether the system is transparent enough to describe the

109. See, e.g., Rudin, *supra* n. 24.
110. Barredo Arrieta et al., *supra* n. 13, at sec. 3.
111. Braun Binder, *supra* n. 9, at pp. 300-301.
112. Barredo Arrieta et al., *supra* n. 13, at sec. 4.
113. Explainable AI systems are preferable in general and although this preference is context-agnostic, it gains special importance in high-stakes decisions: D. Leslie, *Understanding Artificial Intelligence Ethics and Safety: A Guide for the Responsible Design and Implementation of AI Systems in the Public Sector* (The Alan Turing Institute 2019), available at <https://doi.org/10.5281/zenodo.3240529>; Rudin, *supra* n. 24.
114. On the plurality of explanation needs, see Zednik, *supra* n. 25, at sec. 2.2.
115. Id., at sec. 4.2.
116. Mehdiyev et al., *supra* n. 48, at sec. 3.3.2.

input that must be entered and the generated output; and (ii) *why* AI is being used for this purpose. For instance, we can assume that the tax risk assessment AI system has computed a risk value of 0.853. In such a case, the main responsibility of a tax officer is to interpret that value as an indicator of a major risk of fraudulent tax behaviour. Similarly, a tax officer is responsible for determining whether and why and to which extent a risk assessment in a particular case has been generated in line with relevant legal provisions and principles or whether the AI risk assessment system fails to do so. In the latter case, the assessment may discriminate by various illegitimate data points associations, for example, by linking a foreign place of birth with a high level of tax risk.¹¹⁷

For these authorities, the local post hoc explanation models mentioned above may also work well to answer the *why* questions.¹¹⁸ In contrast, *what* questions are most suitably answered using global post hoc explanation models – such as parallel distributed processing (PDP), global surrogate models, Shapley dependent plots, Shapley summary plots or accumulated local effects (ALE) plots – that provide decision-makers with a general view of all decisions made by the AI system rather than explanations for individual decisions.¹¹⁹

Finally, judges need information to resolve disputes between taxpayers and tax authorities, including those triggered by AI systems. In such disputes, courts are interested in carefully identifying *what* AI systems are doing and *why* they are doing it to determine whether their use by the tax authorities was in line with legal provisions, i.e. fair and non-discriminatory as opposed to biased and arbitrary. For the most part, this information is likely to come from global post hoc explanations, which might be consumed by the judges themselves or by expert witnesses. However, it is not inconceivable that certain judicial cases might hinge upon how a system arrived at a decision in a particular case. In these cases, judges might also benefit from the outputs of local post hoc explanations.

3. Constitutional Principles and Human Rights – A Current Path to a Minimum Legal Standard on Tax XAI?

AI is a global phenomenon. As the previous sections show, tax authorities worldwide increasingly rely on AI technologies for their operation, and, in doing so, they often rely on systems and technical solutions produced by a handful of large corporations.¹²⁰ The aforementioned global character of AI has also impacted attempts to regulate its use, as international organizations such as the OECD and the Council of Europe propose guidelines and legal standards for global AI governance.¹²¹ In the taxation domain, such an international approach to regulation is also expected, given the global character of many taxable activities and various successful attempts to implement minimum standards via international

117. By analogy to financial risk in loan risk assessment systems, see Zednik, *supra* n. 25, at secs. 3.1 and 4.2.

118. Cf. Mehdiyev et al., *supra* n. 48, at sec. 3.3.2.

119. Id.

120. On the global supply chains of AI technologies, see K. Crawford, *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (Yale University Press 2021), available at <https://doi.org/10.2307/j.ctv1ghv45t>; J.E. Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism* (Oxford University Press 2019); S. Gürses & J. van Hoboken, *Privacy after the Agile Turn*, in *The Cambridge Handbook of Consumer Privacy* (E. Selinger, J. Polonetsky & O. Tene eds., Cambridge Law Handbooks, Cambridge University Press 2018), available at <https://doi.org/10.1017/9781316831960.032>.

121. Even national regulators, such as the United States, make use of direct and indirect mechanisms to exert influence over AI governance beyond their borders. On the global landscape of AI regulation, see M. Hildebrandt, *Global Competition and Convergence of AI Law*, in *Elgar Encyclopedia for Comparative Law* (J.M. Smits et al. eds., Edward Elgar 2022), available at <https://doi.org/10.31235/osf.io/j36ke>.

treaties, particularly to prevent tax avoidance.¹²² Accordingly, this section and the next investigate the feasibility of a minimum legal standard for XAI in tax law.

The question the authors attempt to answer in this section and the next is whether current and prospective law provides a minimum legal standard of justification that AI systems deployed in tax applications must meet. The authors explore the need for a minimum legal standard for XAI in tax law at the global level, as implementing and enforcing such a standard on a country-by-country basis could lead to extreme divergences across the globe. These divergences, in turn, could severely undermine explainability as a safeguard regarding AI systems and even lead to unfair AI competition, in which actors deploy and use explainable AI systems only in some jurisdictions while ignoring them in others, to the detriment of competitors bounded to stricter requirements.

To prevent unfair tax competition and abuses of tax regimes globally, the OECD – with political support from the G20 – has pushed nearly 100 countries and jurisdictions to implement a minimum standard to prevent tax treaty abuse.¹²³ This successful example shows the potential of minimum standards as a tool for global governance. Given the various threats posed by non-explainable AI systems in tax law and the global character of these risks, the authors see no good reason to abandon such ambitions regarding XAI in tax law for the foreseeable future. The threats posed by non-explainable AI systems in tax law are no less real and serious than those stemming from tax avoidance,¹²⁴ even if they remain unseen and under-analysed by tax policymakers.

This section maps the potential constitutional grounds for such a minimum legal standard. Since the idea of “no taxation without representation” was one of the rallying calls behind modern constitutionalism,¹²⁵ most constitutions establish the legal contours of taxation to a considerable level of detail. Those constitutions – and the international legal texts on human rights that were adopted starting in the second half of the 20th century – also include several provisions on the need for due process regarding decisions that restrict the rights of natural and legal persons. In consequence, the power of the state to establish and collect taxes is subject to various constitutional limits.

As shown in section 2.1., such limits could be eroded by the unchecked use of AI. In particular, the opacity of the inner workings of AI systems and the organizational contexts in which they are embedded means that taxpayers could lack information about how their data

.....

122. “Given the risk to revenues posed by treaty shopping, countries have committed to ensure a minimum level of protection against treaty shopping (the “minimum standard”).” This statement regards the prevention of international tax avoidance through the abuse of tax treaties. See OECD, *Preventing the Granting of Treaty Benefits in Inappropriate Circumstances – Action 6: 2015 Final Report*, at p. 10 (OECD 2015), Primary Sources IBFD, also available at <https://doi.org/10.1787/9789264241695-en>. This minimum standard may be implemented first of all via arts. 6 and 7 of the *Multilateral Convention to Implement Tax Treaty Related Measures to Prevent Base Erosion and Profit Shifting* (MLI) (7 June 2017). As of 21 December 2021 (the newest update by the OECD), 96 countries and jurisdictions signed that multilateral treaty; see <https://www.oecd.org/tax/treaties/beps-ml-signatories-and-parties.pdf>. Another example of the minimum standard to prevent tax avoidance via a multilateral treaty (legally binding on all EU Member States) is Council Directive (EU) 2016/1164 of 12 July 2016 Laying Down Rules against Tax Avoidance Practices that Directly Affect the Functioning of the Internal Market, OJ L 193/1 (19 July 2016).

123. See BEPS Action 6 and MLI, at <https://www.oecd.org/ctp/treaties/beps-ml-signatories-and-parties.pdf>.

124. See sec. 2.

125. See, e.g., A. Sajó & R. Uitz, *The Constitution of Freedom: An Introduction to Legal Constitutionalism* p. 290 (Oxford University Press 2017).

was processed and incorporated into the decision-making process or even about the use of AI in the first place. Even if such information is provided, the technical skills required to understand the operation of AI systems mean that the mere disclosure of the information is not enough.¹²⁶ individuals and small businesses – in fact, anyone but the very largest companies – need some form of explanation of how AI systems operate. Consequently, XAI technologies might be de facto necessary to ensure the effectiveness of constitutional rights when AI is used.

In this section, the authors build upon a longer analysis of constitutional and human rights sources for XAI requirements in tax law, carried out elsewhere by some of the authors of this article.¹²⁷ Section 3.1. engages with the general constitutional features of tax systems worldwide and how AI systems impact them. Since the boom in tax AI technologies described above is a relatively recent development, not all legal systems have dealt with constitutional challenges to tax automation. However, such challenges have already been lodged in a few European countries, which is why section 3.2. examines the legal framework of the ECHR as a potential source of explanation requirements applicable to the use of XAI in taxation. While that convention is not a global document, and its applicability to tax issues is somewhat limited, it has been invoked in various cases involving AI and related technologies. As such, a synthesis of the analysis of the legal requirements for XAI in the ECHR system provides important insights for developing a minimum global standard of explanation for tax AI technologies.¹²⁸

3.1. Constitutional principles and rights related to tax XAI

In principle, constitutions in democratic countries ensure that the calculation and levying of any tax must be precisely and comprehensively regulated by statutory law.¹²⁹ As a result, it is assumed that only legislators have the authority to determine the general principles of tax law and govern every detail of tax collection. The executive power and its agencies, such as tax administrations, have, at most, minor authority over the content of tax law, while still

126. On the limits to information disclosure as a form of accountability, see M. Ananny & K. Crawford, *Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability*, 20 New Media & Society 3 (2018), available at <https://doi.org/10.1177/1461444816676645>.

127. B. Kuzniacki et al., *Requirements for Tax XAI under Constitutional Principles and Human Rights*, in *Explainable and Transparent AI and Multi-Agent Systems* (D. Calvarese et al. eds., Springer 2022 forthcoming).

128. Please note that the considerations in this section do not analyse the changes in legal procedures crucial to the successful implementation of any AI system into the workflow of tax administrations, which fall outside the scope of this research.

129. To illustrate this point, we give four examples. Art. 34 of the French Constitution states that “Statutes shall determine the rules concerning the base, rates and methods of collection of all types of taxes; the issuing of currency”; see *Constitution de la République française* (4 Oct. 1958), available at <https://www.assemblee-nationale.fr/connaissance/constitution.asp> (accessed 9 Aug. 2022). Similarly, art. 217 of the Polish Constitution stipulates that “The imposition of taxes, other public levies, the determination of entities, subjects of taxation and tax rates, as well as the rules for granting reliefs and remissions and the categories of entities exempted from taxes shall be effected by statute”. *Konstytucja Rzeczypospolitej Polskiej*, 2 April 1997 (Journal of Laws of 1997, No. 78, item 483). In the same vein, art. 265 of the Indian Constitution says that: “No tax shall be levied or collected except by authority of law.” Finally, the Brazilian Constitution forbids the Union, States, the Federal District, and Municipalities from imposing or increasing taxes by means other than statutes (art. 150, I, with a few exceptions presented in the same article), a prohibition that encompasses any changes to the constitutive elements of a tax – such as the tax rate, the base, the triggering event or the applicable penalties (art. 97 National Tax Code). See also Hattingh, *supra* n. 78, at sec. 2.

being empowered to execute it as issued by legislators.¹³⁰ In particular, tax authorities must be accountable.¹³¹ For that purpose, any decisions by a tax authority that have an impact on taxpayers must be issued in a sufficiently clear, precise and predictable manner.¹³² The principle of legal certainty, as recognized by EU law in Court of Justice of the European Union (ECJ) case law, sets similar requirements.¹³³

As these are general provisions of tax law, the mere introduction of AI technologies into the fold is not enough – at least from a normative perspective – to void the constraints to which tax authorities are subject.¹³⁴ Consequently, the use of opaque AI technologies in a tax context might introduce obstacles to compliance with the constitutional requirements of legal certainty and accountability surrounding the activities of tax authorities.¹³⁵ In this context, explaining AI decisions becomes necessary¹³⁶ to assess whether tax authorities comply with the constitutional principles that bind their activities, thus ensuring that automated taxation remains within the domain of the rule of law.

The above observations materialized already to a large extent in at least one important judicial precedent. In *eKasa* (17 Dec. 2021),¹³⁷ the Slovak Constitutional Court implicitly underscored the need for the explainability of AI systems in light of constitutional principles.¹³⁸ It stated that “[t]he law restricting fundamental rights must be specific enough to make its application predictable” (paragraph 122) and that “[t]he application of technological progress in public administration cannot result in an impersonal state whose decisions are inexplicable, unexplained and at the same time no one is responsible for them” (paragraph 127). The Slovak Constitutional Court was also vocal about the incompatibility with the constitution of legislation that permits the use of AI systems without effective supervision, thereby failing to ensure the proportional application of the technology (paragraph 129).¹³⁹ Such proportionality is not possible in the absence of access to fundamental building blocks

-
- 130. See, inter alia, Kuźniacki et al., *supra* n. 127; H. Ávila, *Certainty in Law* (Springer 2016).
 - 131. L. Huttner & D. Merigoux, *Catala: Moving Towards the Future of Legal Expert Systems* (2022), available at <https://hal.inria.fr/hal-02936606> (accessed 26 Feb. 2022).
 - 132. Kuźniacki, *supra* n. 51.
 - 133. See, for example, PT: European Court of Justice (ECJ), 3 Oct. 2013, Case C-282/12, *Itelcar – Automóveis de Aluguer Lda v. Fazenda Pública*, para. 44, Case Law IBFD; BE: ECJ, 5 July 2012, Case C-318/10, *Société d’investissement pour l’agriculture tropicale SA (SIAT) v. État belge*, paras. 58–59, Case Law IBFD.
 - 134. See, for example, the discussion of the Slovak constitutional case law in this subsection.
 - 135. This failure may stem from the technical sources of opacity this article engages with, but also from non-technical sources of opacity such as the secrecy requirements discussed in sec. 2.2.
 - 136. Though potentially insufficient for that purpose: G. Buchholtz, *Artificial Intelligence and Legal Tech: Challenges to the Rule of Law*, in *Regulating Artificial Intelligence* (T. Wischmeyer & T. Rademacher eds., Springer International Publishing 2020), available at https://doi.org/10.1007/978-3-030-32361-5_8; E. Bayamlioglu & R. Leenes, *The ‘Rule of Law’ Implications of Data-Driven Decision-Making: A Techno-Regulatory Perspective*, 10 Law, Innovation and Technology 2 (2018), available at <https://doi.org/10.1080/17579961.2018.1527475>.
 - 137. Slovenian Constitutional Court, 17 Dec. 2021, *Ústavného súdu Slovenskej republiky PL. ÚS 25/2019-117 V mene Slovenskej republiky* (accessed 7 Mar. 2022).
 - 138. Kuźniacki, *supra* n. 51.
 - 139. In that regard, the Slovak Constitutional Court referred to the Advocate General Opinion in Opinion 1/15 of ECJ: SK: ECJ, 8 Sept. 2016, *Opinion of Advocate General Mengozzi delivered on 8 September 2016. Opinion 1/15 (Request for an opinion submitted by the European Parliament)*, para. 176: “[...] no one can fail to be aware that the transfer of voluminous quantities of personal data of air passengers, which includes sensitive data, requiring, by definition, automated processing, and the retention of that data for a period of five years, is intended to permit a comparison, which will be retroactive where appropriate, of that data with pre-established patterns of behaviour that is ‘at risk’ or ‘of concern’, in connection with terrorist activities and/or serious transnational crime, in order to identify persons not hitherto known to the police or not suspected”.

of the AI systems used by tax administrations. These blocks include, inter alia, the system inputs or assessment criteria and the logic of the decision or individual assessment. Also, the mentioned proportionality requires the determination of patterns used for automated assessment and models or other databases that lead to a particular decision (paragraphs 137-138). In light of these requirements, non-explainable AI systems cannot be effectively supervised in the sense specified by the Slovak Constitutional Court.¹⁴⁰

Accordingly, sufficient explainability appears to be a pivotal feature of all constitutionally compliant AI systems in the tax domain.¹⁴¹ This conclusion, exclusively from a legal perspective, stems from the fact that the execution of tax law shall be maximally transparent and precise for taxpayers. Of course, transparency and precision of the execution of tax law are not absolute values and must be weighed with effectiveness in preventing tax fraud. However, any compromise between these desiderata has to be well-justified, not just attending to the public interest in combating fraud but also taking into account in a proportional way the various other interests of taxpayers and authorities.¹⁴²

3.2. XAI in the context of the ECHR

Several provisions of the ECHR could be seen as requirements for the explainability of AI systems. This section discusses three such provisions: (i) the requirement of proper justification, which is part of the right to a fair trial (article 6 of the ECHR); (ii) the right to respect for private and family life (article 8 of the ECHR); and (iii) the prohibition of discrimination in conjunction with protection of property (article 14 and Protocol 12 of the ECHR).¹⁴³ Given their general character, none of these provisions explicitly mentions XAI, and the case law of the European Court of Human Rights (ECtHR) is yet to engage with matters of AI in the tax domain. Nevertheless, requirements for explanation can be read into these provisions, and national courts of parties to the ECHR increasingly do so.

Regarding the right to a fair trial in the ECHR, its applicability to tax law cases was significantly restricted by the ECtHR (Grand Chamber) in *Ferrazzini v. Italy* (12 July 2001).¹⁴⁴ In the judgment, the ECtHR (by a majority of 11 to 6) decided that article 6 of the ECHR does not apply to ordinary tax proceedings because taxation and tax-related disputes fall “outside

140. Interestingly, in the Court’s view, the Slovak Constitution sets a much higher legal standard for the legality of AI systems used by public institutions than the GDPR, which the authors discuss in sec. 4.1.1. This higher standard appears, according to the court, not least because the GDPR is just an incomplete harmonization of the processing of personal data by public authorities for the purpose of exercising their tasks, whereas the Constitutional principles require AI systems to be subjected to registries, individual disclaimers, testing of datasets or models they use, access to the logic of their decisions and otherwise more holistic assessment of their impact on fundamental rights (secs. 133-135, 138). See also Kuźniacki, *supra* n. 51.

141. Kuźniacki, *supra* n. 51

142. Kuźniacki et al., *supra* n. 127, sec. 2.

143. We do not distinguish here between taxpayers who are individuals (natural persons) and companies (legal persons), since art. 34 of the ECHR stipulates that the European Court of Human Rights (ECtHR) “may receive applications from any person, nongovernmental organisation or group of individuals claiming to be the victim of a violation” of the ECHR by the person’s or entity’s Member State. Thus, not only individual taxpayers but also companies have standing before the ECtHR if they are not “governmental organisations” within the meaning of art. 34 of the ECHR. For the relevant explanation and case law of the ECtHR, see ECtHR, *Practical Guide on Admissibility Criteria* pp. 9-10 (European Court of Human Rights 1 Feb. 2022). However, claims of companies under the ECHR are not unproblematic for the ECtHR. See more in M. Emberland, *The Human Rights of Companies: Exploring the Structure of ECHR Protection* (Oxford University Press 2006).

144. ECtHR, 12 July 2001, Application no. 44759/98, *Case of Ferrazzini v. Italy*.

the scope of civil rights and obligations, despite the pecuniary effects which they necessarily produce for the taxpayer”, since taxation is deemed to form part of the hard core of public authority prerogatives (paragraph 29).¹⁴⁵ Consequently, article 6 of the ECHR would only be applicable in situations in which state action in the tax domain has a “criminal connotation” (paragraph 20); in that case, protection is offered by the criminal limb of article 6 of the ECHR. Tax proceedings with criminal connotations concern the imposition of penalties (tax surcharges/criminal tax sanctions) on citizens in their capacity as taxpayers, essentially as punishment to discourage them from reoffending (a deterrent and a punitive purpose).¹⁴⁶ In other words, according to the ECtHR, article 6 applies to tax disputes only insofar as they are triggered by tax surcharges imposed on taxpayers for their behaviour of a criminal nature, typically regarding offences of lack of tax compliance. A determination of whether a tax dispute enters the ambit of article 6 is circumstantial and subject to vast discretion on the part of the ECtHR.¹⁴⁷

Some scholars propose relying on articles 47 (right to an effective remedy and fair trial) and 48 (presumption of innocence and right of defence) of the EU Charter, which partly overlap with article 13 of the ECHR, rather than article 6 of the ECHR to surmount the obstacles stemming from *Ferrazzini*.¹⁴⁸ The rationale is that the level of protection guaranteed by the ECHR in comparison to the EU Charter in terms of defence rights is similar,¹⁴⁹ but only the latter is not confined to disputes relating to civil rights and obligations. Thus, articles 47 and 48 of the EU Charter can apply to “normal” tax proceedings.¹⁵⁰ Still, the EU Charter has a limited scope of application in general. In light of the principle of conferral, it applies to the

-
145. Philip Baker has described this ruling as one of the greatest failings of the ECtHR, since, in practical terms, it means that under the ECHR, a taxpayer in a member state of the ECHR does not have the right “in an ordinary tax dispute to a fair trial by an independent and impartial tribunal”. P. Baker, *60 Years of the European Convention on Human Rights and Taxation*, 61 Eur. Taxn. 12, secs. 2 and 4 (2021). Journal Articles & Opinion Pieces.
146. See, for example, ECtHR, 4 May 2017, Application no. 15485/09, *Case of Chap Ltd v. Armenia case*; ECtHR, 17 Apr. 2012, Application no. 21539/07, *Case of Steininger v. Austria case*; ECtHR, 24 Nov. 2006, Application no. 73053/01, *Case of Jussila v. Finland*.
147. This can also lead to bizarre outcomes: if a tax surcharge (penalty) was imposed on a taxpayer, even a little one, the taxpayer is then fully protected under the criminal limb of article 6. If, by contrast, a taxpayer is not faced with a tax surcharge but is liable to an additional amount of tax with no deterrent/penalization purpose, even if the amount is enormous, the taxpayer does not enjoy the panoply of criminal charge rights in this provision. See Baker, *supra* n. 145, at sec. 2.
148. Hadwick & Lan, *supra* n. 6, at sec. 2.2. For the potential overlap between arts 6 and/or 13 ECHR (right to an effective remedy) with art. 47 the EU Charter, see F. Casarosa et al., *Module 3 – Right to an Effective Remedy*, in *ACTIONES Handbook on the Techniques of Judicial Interactions in the Application of the EU Charter* sec. 2.2. (2019) and the ECJ case law cited therein, available at <https://cjc.eui.eu/wp-content/uploads/2019/03/D1.1.c-Module-3.pdf> (accessed 30 June 2022).
149. As scholars have observed, the absence of justification for a tax decision constitutes a violation of art. 47(1) of the EU Charter and of the EU law principle of observance of the rights of the defence. See P. Pistone, *The EU Charter of Fundamental Rights, General Principles of EU Law and Taxation*, in *European Tax Law* 99-100 (P.J. Wattel et al. eds., Wolters Kluwer 2018). A mere statement by the tax authorities that their decision follows from the output of an AI system would not seem to differ from the complete absence of justification and thus appears to violate the aforementioned principles. The taxpayers are the addressees of tax decisions and whenever such decisions significantly affect their interests, “they must be placed in a position in which they can effectively make known their views as regards the information on which the authorities intend to base their decision”. See PT: ECJ, 18 Dec. 2008, Case C-349/07, *Sopropé - Organizações de Calçado Lda v. Fazenda Pública*, para. 37, Case Law IBFD.
150. Pistone, *supra* n. 149. See also K. Perrou, *Fundamental Rights in EU Tax Law*, in *Research Handbook on European Union Taxation Law* 515-517 (C. HJI Panayi, W. Haslehner & E. Traversa eds., Edward Elgar Publishing 2020), available at <https://doi.org/10.4337/9781788110846.00034>.

EU Member States only when they are implementing EU law.¹⁵¹ Consequently, as confirmed by the ECJ in *Fransson* (C-617/10), the EU Charter does not cover purely domestic situations or situations that are not caused by an application of EU law by the EU Member States.¹⁵²

Accordingly, article 6 of the ECHR constitutes a relevant legal source that ensures the right to a fair trial in disputes triggered by applying AI systems to tax law only when they arise out of the imposition of tax surcharges on taxpayers. In such cases, the two following elements of the right to a fair trial seem to be most relevant in respect of tax AI systems: (i) the minimum guarantees of equality of arms; and (ii) the right of defence.¹⁵³ In these cases, the right to a fair trial ensures taxpayers have an effective review of the information on which the tax authorities base their decisions.¹⁵⁴ Hence, the right to a fair trial under the ECHR is of limited assistance in ensuring the explainability of AI systems in tax law, even within its geographical scope of application. The reason for this does not lie in the inadequate merits of that right to XAI in tax law but in its limited personal, substantive and territorial scope of application. However, parties to the convention still retain the power to extend their understanding of the right to a fair trial beyond the minimum required by the ECHR. Doing so with regard to tax AI might avoid the constraints described above.¹⁵⁵

Another ECHR right connected to explanation requirements in the ECHR is the right to respect for private and family life, enshrined in article 8 of the Convention. In fact, it has become one of the key fundamental rights affecting the legality of the use of AI systems in tax law due to the seminal and widely debated¹⁵⁶ judgment of the Hague District Court in the Netherlands in *SyRI* (*Systeem Risico Indicatie*) (5 Feb. 2020).¹⁵⁷ In that case, the focal point of the arguments of the claimants (NJCM et al.) was the alleged violation of article 8 of the ECHR. Accordingly, the court focused its analysis on whether the *SyRI* legislation met the conditions under article 8(2) of the ECHR, which lays down restrictions on the right to respect for private life.¹⁵⁸

151. Art. 51 EU Charter. See also K. Lenaerts & J.A. Gutiérrez-Fons, *The Place of the Charter in the European Legal Space*, in *The EU Charter of Fundamental Rights: A Commentary* para. 55.89 (S. Peers et al. eds., Hart Publishing 2021).

152. See SE: ECJ, 26 Feb. 2013, Case C-617/10, *Åklagaren v Hans Åkerberg Fransson*, paras. 19-21.

153. Kuźniacki, *supra* n. 51.

154. ECtHR, 9 Oct. 2008, Application no. 62936/00, *Case of Moiseyev v. Russia*; ECtHR, 24 Sept. 2007, Application no. 38184/03, *Case of Matyjek v. Poland*.

155. As Baker put it, “[i]n some countries there will be constitutional provisions that give greater rights than under the ECHR”, including the rights to fair trial: Baker, *supra* n. 145, at sec. 2.

156. The *SyRI* case triggered hot public debate not only in Dutch, but also in foreign media and television, which presumably prompted the court to publish its judgment in an English translation (perhaps the court realized the judgment’s importance and international interest in it). For the Dutch media and television, see M. Wieringa, G. van Schie & M. van de Vinne, *De Discussie Omtrent SyRI Moet over Meer Dan Alleen Privacy Gaan*, IBestuur (10 Feb. 2020), available at <https://ibestuur.nl/podium/de-discussie-omtrent-syri-moet-over-meer-dan-alleen-privacy-gaan> (accessed 3 Mar. 2022). For foreign media and television, see T. Simonite, *Europe Limits Government by Algorithm. The US, Not So Much*, Wired (7 Feb. 2020), available at <https://www.wired.com/story/europe-limits-government-algorithm-us-not-much/> (accessed 3 Mar. 2022); J. Henley & R. Booth, *Welfare Surveillance System Violates Human Rights, Dutch Court Rules*, The Guardian (5 Feb. 2020), available at <https://www.theguardian.com/technology/2020/feb/05/welfare-surveillance-system-violates-human-rights-dutch-court-rules> (accessed 3 Mar. 2022).

157. Rechtbank Den Haag, 5 Feb. 2020, *NJCM et al. v. Netherlands*. “NJCM” stands for *Nederlands Juristen Comité voor de Mensenrechten* (the Dutch section of the International Commission of Jurists), which is an organization involved in the protection and strengthening of human rights and fundamental freedoms. NJCM, together with other civil society interest groups and two natural persons (claimants), brought the case against the Dutch state (defendant).

158. Apart from the arguments of the claimants, the reason the court paid so much attention to the ECHR rather than the GDPR was most likely that the court assumed that examining the *SyRI* legislation through

The main task of the court was to assess whether the SyRI legislation met the requirements of necessity and proportionality pursuant to article 8(2) of the ECHR in light of the aims it pursues. In that regard, the court explained, inter alia, which characteristics of the SyRI legislation failed to comply with article 8(2) of the ECHR. Notably, the objective factual data, which could lead to the justifiable conclusion that there was an increased risk of tax fraud, was fully opaque and thus non-explainable. Also salient were the risk model, the type of algorithm used in the model, and the risk analysis method (paragraphs 6.87 and 6.89). As a result, the addressees of the SyRI legislation were neither able to defend themselves against the fact that a risk report concerning them had been submitted nor could they be sure that their data had been processed on proper grounds (paragraph 6.90). This reveals that the court made an implicit connection with the right to a fair trial in article 6 of the ECHR, since the inability to lodge an effective defence stems from the same features of the SyRI legislation, i.e. its lack of transparency in respect of the functioning of SyRI, thereby rendering that system unexplainable by law.¹⁵⁹ Accordingly, in the court’s view, the SyRI legislation was not necessary in a democratic society, and it failed to strike the “fair balance” required under the ECHR between the social interest the legislation serves and the violation of private life to which the legislation gives rise (paragraphs 6.72 and 6.105-106).

Indeed, any legislation that provides for the use of an AI system without granting to its addressees (e.g. taxpayers) the right to receive a sufficient explanation of its functioning would not appear to pass the muster of legality in a democratic country.¹⁶⁰ A legislative black-box AI system is far from striking such a balance. Since the SyRI legislation permitted such a black-box system for risk profiling purposes in the Netherlands, the court was, in the authors’ view, correct to decide that this legislation violates article 8(2) of the ECHR and, consequently, the right to protection of private life enshrined in article 8(1) of the ECHR.

In addition to the non-interference requirement found by the Dutch court in this case, article 8(1) of the ECHR also establishes positive obligations to the state. In the context of an information society, such positive measures may require establishing rules that mandate the adoption of technical measures that enforce the protection of private life. As shown above, tax AI systems frequently deal with information regarding the private and family lives of taxpayers, and the outputs of these systems are used to guide decisions that interfere with taxpayers’ lives. Failure to provide adequate explanations regarding the processing of this information and the uses of the outputs of this processing might therefore violate the positive obligation of protection of private and family life.¹⁶¹

Finally, it is worth looking closer at the prohibition of discrimination (article 14) in conjunction with other provisions of the ECHR and the Protocols.¹⁶² This legal basis is rarely

.....

the lens of international human rights law would sound more convincing than doing the same from the perspective of a detailed regulation such as the GDPR. Deciding this case by legal principle rather than a specialized legal instrument legal regulation might also create a rock-solid juridical precedent. See M. van Bekkum & F.Z. Borgesius, *Digital Welfare Fraud Detection and the Dutch SyRI Judgment*, 2021 European Journal of Social Security, available at <https://doi.org/10.1177/13882627211031257>.

159. On the parallels between the right to a fair trial and the right to a private life, see, inter alia, M. Almada & M. Dymitruk, *Data Protection and Judicial Automation*, in *Research Handbook on EU Data Protection* (E. Kosta & R. Leenes eds., Edward Elgar Publishing 2022).

160. See Van Bekkum & Borgesius, *supra* n. 158, at p. 13.

161. On technical measures as positive obligations under art. 8 ECHR, see ECtHR, 17 July 2008, Application no. 20511/03, *Case of I v. Finland*.

162. The prohibition of discrimination in art. 14 ECHR enshrines the right not to be discriminated against in “the enjoyment of the rights and freedoms set out in the Convention”, thereby merely complementing the

successfully invoked by taxpayers before the ECtHR, as that court has recognized a wide margin of appreciation for States in tax matters regarding discriminatory tax measures.¹⁶³ The ECtHR has found violations of the prohibition of discrimination in quite unusual tax cases, such as those regarding (i) the discriminatory tax treatment of a non-resident in comparison to residents due to the possibility for only the latter to opt out of the payment of church tax (a violation of article 14 in conjunction with article 1 Protocol 1 of the ECHR);¹⁶⁴ (ii) a differentiation between persons declared unfit for military service and exonerated from paying a certain tax and those declared unfit for service but obliged to pay it (a violation of article 14 in conjunction with article 8 of the ECHR);¹⁶⁵ or (iii) a failure on the part of the tax authorities to take the needs of a child with disabilities into account when determining his father's eligibility for tax relief on the purchase of a suitably adapted property (a violation of article 14 in conjunction with article 1 Protocol 1 of the ECHR).¹⁶⁶ In many other tax-related cases, the ECtHR either declared that the applications were inadmissible or found no violation of the invoked article 14 of the ECHR in conjunction with other provisions of the ECHR or the Protocols.¹⁶⁷

This shows that the ECtHR generally has not recognized discriminatory tax treatment as a breach of the rights in article 14 to the non-discriminatory enjoyment of the various rights protected under the ECHR.¹⁶⁸ Thus, although XAI can help combat discriminatory tax treatment stemming from the use of AI tax-related systems by identifying the discriminatory features of such systems – such – as biased data and/or factors decisive to deliver a decision in a discriminatory way – the prohibition of discrimination in article 14 of the ECHR does not seem to provide a solid legal basis for ensuring the explainability of AI systems in tax law.¹⁶⁹

.....
other substantive provisions of the ECHR and the Protocols. This means that art. 14 does not prohibit discrimination per se, but only discrimination in the enjoyment of the rights and freedoms outlined in the ECHR; it is therefore applied in conjunction with other provisions of that Convention and the Protocols. See ECtHR, *Guide on Article 14 of the European Convention on Human Rights and on Article 1 of Protocol No. 12 to the Convention: Prohibition of Discrimination* (ECtHR 31 Aug. 2021).

163. See Baker, *supra* n. 145, at sec. 2.

164. See ECtHR, 23 Oct. 1990, Application no. 11581/85, *Case of Darby v. Sweden*, paras. 33-34. The applicant in that case was a Finnish citizen who had been working in Sweden for a while as a doctor. During that period, he could not opt out of paying a special tax to the Lutheran Church of Sweden because he was neither a Swedish citizen nor a Swedish tax resident.

165. See ECtHR, 30 Apr. 2009, Application no. 13444/04, *Case of Glor v. Switzerland*, paras. 97-98. The applicant in that case was unfit for military service as he suffered from type 1 diabetes (*diabetes mellitus*). In spite of that, he was obliged to pay the military-service exemption tax.

166. See ECtHR, 22 Mar. 2016, Application no. 23682/13, *Case of Guberina v. Croatia*, paras. 98-99. As follows from the wording of that case: the applicant in that case was a father of a 100% disabled child (incurable cerebral palsy, grave mental retardation and epilepsy). He bought the house and moved from his flat because the flat did not have lift and therefore it did not meet the needs of his disabled child and his family. Notably, it was extremely hard to take his son out of the flat to see a doctor or take him for physiotherapy and to kindergarten or school, and to meet his other social needs. Nevertheless, the tax authorities denied him the tax exemption for a person who was buying a flat or a house in order to solve his or her housing needs.

167. See, for example, the case law indicated in ECtHR, *Taxation and the European Convention on Human Rights*, Factsheets 20-22 (ECtHR Dec. 2021).

168. See Baker, *supra* n. 145, at sec. 2.

169. Considering that the principle of non-discrimination is one of the cornerstones of international human rights law and that one of the most frequently reported impacts of AI systems on human rights is the impact on the right to be shielded from discrimination, the observation in the previous sentence is unsatisfactory. See C. Muller, *The Impact of Artificial Intelligence on Human Rights, Democracy and the Rule of Law*, Ad hoc Committee on Artificial Intelligence (CAHAI) (Council of Europe 2020); B.W. Goodman, *Economic Models of (Algorithmic) Discrimination* (2016).

4. Hard and Soft Law Instruments – A Prospective Path to a Minimum Legal Standard on Tax XAI?

Constitutional requirements could provide the grounds for a minimum global standard on XAI, but its implementation is unlikely to take the form of constitutional amendments.¹⁷⁰ Instead, some of these standards are defined through legally binding (“hard law”) specific regulations on AI on a domestic or a regional level, while others amount to non-legally binding recommendations and guidelines (“soft law”) in that regard.¹⁷¹ Some of these sources of hard and soft law could have important implications for adopting XAI in a given jurisdiction, not least by sketching the contours of a prospective minimum legal standard for XAI in tax law.

Accordingly, section 4.1. analyses the current and proposed instruments for the governance of AI in the European Union. Since the European Union exercises considerable influence over the regulation of digital technologies even beyond its borders,¹⁷² an analysis of the European approach to XAI regulation provides insights on the limited extent to which existing law creates requirements for explanation in tax AI systems. Current international and national law does not establish legally binding requirements for XAI in tax,¹⁷³ but some soft law instruments engage with matters of AI and explanation, which are scrutinized in section 4.2.

4.1. EU law on AI

The impact of computer technologies on society is not a novel concern for lawmakers. Starting in the 1970s, jurisdictions worldwide developed sophisticated legal frameworks for governing the automated storage and use of information about individuals in a digital format.¹⁷⁴ In the European Union, these frameworks evolved in the direction of an auton-

170. Constitutional principles are usually flexible enough to ensure their adaptation to new technological contexts; this happens mostly through changes in interpretation and case law, rather than changes to the constitution itself: R. Leenes et al., *Constitutional Rights and New Technologies: A Comparative Study* (T.M.C. Asser Press; Distributed by Cambridge University Press 2008). This is not to say, however, that constitutional amendments in response to technology are unheard of: both the Treaty on the Functioning of the European Union and the Brazilian Constitution have been amended to include a right to the protection of personal data.

171. On soft law, see R. Calo, *Artificial Intelligence and the Carousel of Soft Law*, 2 IEEE Transactions on Technology and Society 4 (Dec. 2021), available at <https://doi.org/10.1109/TTS.2021.3113288>; G.E. Marchant, *Professional Societies as Adopters and Enforcers of AI Soft Law*, 2 IEEE Transactions on Technology and Society 4 (Dec. 2021), available at <https://doi.org/10.1109/TTS.2021.3116524>; S.L. Schwarcz, *Soft Law as Governing Law*, 104 Minn. L. Rev. 5 (2020).

172. On the global reach of EU tech regulation, see A. Bradford, *The Brussels Effect: How the European Union Rules the World* ch. 5 (Oxford University Press 2020); P.M. Schwartz, *Global Data Privacy: The EU Way*, 94 New York University Law Review 4, pp. 771-818 (2020); C. Kuner, *The Internet and the Global Reach of EU Law*, in *EU Law Beyond EU Borders* (Oxford University Press 2019), available at <https://doi.org/10.1093/oso/9780198842170.003.0004>.

173. A potential exception is the Council of Europe’s Modernised Convention 108 (Convention 108+), which includes provisions on automated decision-making that are not dissimilar to the EU framework described in sec. 4.1. While this Convention is not global in scope, it has been suggested that it could provide the foundations for a global standard: A. Mantelero, *The Future of Data Protection: Gold Standard vs. Global Standard*, 40 Computer Law & Security Review (Apr. 2021), available at <https://doi.org/10.1016/j.clsr.2020.105500>. However, the Modernised Convention itself has not yet come into force as of March 2022.

174. For a historical overview of these developments, see P. Pałka, *Data Management Law for the 2020s: The Lost Origins and the New Needs*, 68 Buffalo Law Review 2, sec. I (Apr. 2020).

omous right to data protection¹⁷⁵ governed by a handful of legal instruments defined at the EU level,¹⁷⁶ notably the General Data Protection Regulation (GDPR).¹⁷⁷ These instruments have often been touted as the foundations of AI regulation,¹⁷⁸ a claim that is consistent with the role that large datasets play in modern machine learning technologies.¹⁷⁹ Yet, data governance matters do not exhaust the role of AI regulation,¹⁸⁰ which is why the European Commission recently proposed a Regulation for a European Approach to AI.¹⁸¹

This section examines how the EU regulatory framework has sought to require some degree of explainability for the lawful use of AI systems in general. The first approach to such regulation within the EU data protection law framework is restricted to situations involving the processing of data that can be associated with natural persons, meaning that many applications in the tax domain do not fall within its scope, notably those related to corporate taxpayers. The second approach, which has its centre in the proposed AI Act, regulates AI systems as products put on the European single market, regardless of their reliance on personal data. Although we will see that many applications of AI in the tax domain fall outside the scope of either of these regulations, their examination provides useful insights for designing explanation systems in tax contexts.

4.1.1. Data protection law and the right to an explanation

The first contribution of data protection law to the governance of AI systems comes from the idea of the “right to an explanation”. This expression comes from Recital 71 of the GDPR,¹⁸²

175. Under the EU legal system, data protection has a considerable overlap with privacy, but cannot be reduced to it. There are situations in which lawful processing of personal data may lead to violations of privacy, and, in other cases, data protection rights may be violated even without an unjustified intrusion into an individual’s private or family life. On the difference between these two rights, see J. Kokott & C. Sobotta, *The Distinction between Privacy and Data Protection in the Jurisprudence of the CJEU and the ECtHR*, 3 International Data Privacy Law 4 (Nov. 2013), available at <https://doi.org/10.1093/idpl/ipt017>.

176. For an overview of the movements that led to the centralization of data regulation in the European Union, see T. Streinz, *The Evolution of European Data Law*, in *The Evolution of EU Law* (P. Craig & G. de Búrca eds., 3rd ed., Oxford University Press 2021), available at <https://doi.org/10.1093/oso/9780192846556.003.0029>. The general rules on data protection offered by the GDPR are highly specific, but there has been some pushback from EU Member States regarding specific domains of data processing, most notably when it comes to processing for law enforcement purposes: M. Leiser & B. Custers, *The Law Enforcement Directive: Conceptual Challenges of EU Directive 2016/680*, 5 European Data Protection Law Review 3 (2019). Even the GDPR, however, leaves Member States with some margin for manoeuvre, in particular when it comes to the regulation of automated decision-making: G. Maltieri, *Automated Decision-Making in the EU Member States: The Right to Explanation and Other “Suitable Safeguards” in the National Legislations*, 35 Computer Law & Security Review 5 (Oct. 2019), available at <https://doi.org/10.1016/j.clsr.2019.05.002>.

177. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1 (accessed 18 Sept. 2020).

178. See, e.g., Paul Nemitz’s claim that the GDPR “is the first piece of legislation for AI”: *Constitutional Democracy and Technology in the Age of Artificial Intelligence*, 376 Philosophical Transactions of the Royal Society A 2133, p. 8 (28 Nov. 2018), available at <https://doi.org/10.1098/rsta.2018.0089>.

179. B. Koch et al., *Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research*, in *Proceedings of NeurIPS 2021* (8 Dec. 2021).

180. For a legally oriented introduction to the various elements, other than data, that form an AI system, see D. Lehr & P. Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 University of California, Davis Law Review, pp. 653-717 (2017).

181. AI Act (Commission Proposal) (2021).

182. The GDPR is not the sole legal instrument governing data protection in the European Union. Data processing by EU institutions is governed by Regulation 2018/1725 (EUDPR), and Europol is subject to

which states that individuals subject to automated decisions should have the right “...to obtain an explanation of the decision reached...” based on the processing of their personal data. But, since a recital is not a binding instrument of EU law,¹⁸³ there has long been debate on whether such a right exists in the GDPR,¹⁸⁴ and, if so, what its legal grounds would be. The dominant position in legal scholarship has been that the right to an explanation is effectively established in articles 13-15 of the GDPR, which require that controllers of automated decision-making systems provide “meaningful information about the logic involved in the processing”.¹⁸⁵ Equipped with such information, data subjects would then be able to evaluate whether their rights, liberties and interests have been properly attended to in the processing and, if that is not the case, seek redress through instruments such as the right to human intervention in the decision-making process¹⁸⁶ or administrative and judicial proceedings.¹⁸⁷ Explanation, in this context, provides a stepping stone to ensure the contestability of automated decisions,¹⁸⁸ a rationale that found its way into case law.¹⁸⁹

Once it was established that the right to an explanation existed in the GDPR, a substantial body of literature sought to build conceptual and technological tools to ensure this right’s effectiveness. In addition to the scholarship mentioned above on the interpretation of the specific legal instruments, current scholarship on explanation has also incorporated much interdisciplinary work, which seeks to unveil what kinds of explanation are technologically feasible and whether and how they can meet different legal standards, such as the need for the justification of administrative and judicial decisions.¹⁹⁰ By mapping this interface

-
- specific rules defined in Regulation (EU) 2016/794. Furthermore, data processing by national law enforcement authorities is subject to Directive 2016/680 (the LED), and the choice of a Directive as opposed to a Regulation allows Member States to adopt less harmonized rules than the ones prescribed for general data processing. Nevertheless, all those instruments are, to a substantial extent, patterned after the GDPR, and so this section will only engage with other data protection laws if they substantially diverge from it.
183. This is not to say that recitals are deprived of all legal meaning, as their content is used to interpret ambiguities and uncertainties in the legal text itself, including by the Court of Justice itself: T. Klimas & J. Vaiciukaite, *The Law of Recitals In European Community Legislation*, 15 ILSA Journal of International & Comparative Law 1 (Jan. 2008).
184. See, e.g., B. Casey, A. Farhangi & R. Vogl, *Rethinking Explainable Machines: The GDPR’s Right to Explanation Debate and the Rise of Algorithmic Audits in Enterprise*, 34 Berkeley Technology Law Journal 1 (2019); M.E. Kaminski, *The Right to Explanation, Explained*, 34 Berkeley Technology Law Journal 1, pp. 189-218 (2019); A.D. Selbst & J. Powles, *Meaningful Information and the Right to Explanation*, 7 International Data Privacy Law 4 (Nov. 2017), available at <https://doi.org/10.1093/idpl/ix022>; S. Wachter, B. Mittelstadt & L. Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 International Data Privacy Law 2 (May 2017), available at <https://doi.org/10.1093/idpl/ix005>; L. Edwards & M. Veale, *Slave to the Algorithm: Why a Right to an Explanation Is Probably Not the Remedy You Are Looking For*, 16 Duke Law & Technology Review 1 (2017).
185. M. Brkan & G. Bonnet, *Legal and Technical Feasibility of the GDPR’s Quest for Explanation of Algorithmic Decisions: Of Black Boxes, White Boxes and Fata Morganas*, 11 European Journal of Risk Regulation 1, pp. 21-22 (Mar. 2020), available at <https://doi.org/10.1017/err.2020.10>.
186. Art. 22(3) GDPR.
187. Arts. 77 and 79 GDPR.
188. M.E. Kaminski & J.M. Urban, *The Right to Contest AI*, 121 Columbia Law Review 7 (Nov. 2021); Sarra, *supra* n. 11; M. Almada, *Human Intervention in Automated Decision-Making: Toward the Construction of Contestable Systems*, in *Proceedings of the 17th International Conference on Artificial Intelligence and Law (ICAIL 2019)* (ACM 2019), available at <https://doi.org/10.1145/3322640.3326699>.
189. See, e.g., the *Uber and Ola* judgments on the use of automatic decision-making tools for supervising workers in gig economy apps. R. Gellert, M. van Bekkum & F. Zuiderveen Borgesius, *The Ola & Uber Judgments: For the First Time a Court Recognises a GDPR Right to an Explanation for Algorithmic Decision-Making*, EU Law Analysis (28 Apr. 2021), available at <https://eulawanalysis.blogspot.com/2021/04/the-ola-uber-judgments-for-first-time.html> (accessed 30 Apr. 2021).
190. See, e.g., Bibal et al., *supra* n. 17; Brkan & Bonnet, *supra* n. 185.

between what the law needs and what technology can offer, these works stimulate the development of new explanation techniques and direct the application of existing technologies.

This article subscribes to such an interdisciplinary approach, extending it beyond the confines of data protection and toward the general requirements of explanation in tax law. This movement is necessary because many uses of AI in the tax domain are not covered by the explanation requirements present in the GDPR. A first hurdle to deriving the right to an explanation from data protection law comes from article 2(1) of the GDPR, which stipulates that the processing of personal data is a necessary condition for the application of the Regulation.¹⁹¹ Since article 4(1) further specifies that personal data is information relating to an identified – or an identifiable – natural person, any processing of personal data for tax purposes, even if the actual outcome does not amount to personal data, is covered by the GDPR.¹⁹² Even this broad construction, however, excludes some types of AI systems from the reach of the GDPR. For example, systems that only process anonymized data¹⁹³ or deal with statistical data at an aggregate level are not dealing with information that can be associated with a specific natural person.¹⁹⁴ Therefore, they are not subject to any GDPR-related explanation requirements.

More importantly for tax law purposes, it follows from this definition that EU data protection law does not offer any form of protection to non-natural entities.¹⁹⁵ Since data subject rights are directly associated with the status of a data subject – that is, of the person to whom the data pertains – the GDPR does not provide legal persons or non-personified entities with the right to obtain an explanation for decisions produced by an AI system. Consider an AI system that receives information about a corporation to evaluate how much corporate income tax it needs to pay. Whenever this system handles information about the

-
191. Art. 4(1) GDPR defines personal data as “... any information relating to an identified or identifiable natural person”. This natural person is called a “data subject”.
 192. In a recent decision, the ECJ ruled that the Latvian tax authority was bound by the GDPR as it requested that an online commerce platform provide information about people who advertised vehicles for sale: LV: ECJ (Fifth Chamber), 24 Feb. 2022, Case C-175/20, *SIA “SS” v. Valsts ieņēmumu dienests*.
 193. Exclusion from the definition of personal data requires full anonymization. Mere pseudonymization of the data is insufficient to avoid the application of EU data protection law, but the standard for what counts as anonymized data must be evaluated in light of the risks associated with data processing and the current technological capabilities for re-identifying data: M. Almada, J. Maranhão & G. Sartor, *Article 4 Para. 5. Pseudonymisation*, in *General Data Protection Regulation. Article-by-Article Commentary* (I. Spiecker gen. Döhmman et al. eds., Beck; Nomos; Hart Publishing 2022); M. Finck & F. Pallas, *They Who Must Not Be Identified – Distinguishing Personal from Non-Personal Data under the GDPR*, 10 *International Data Privacy Law* 1 (2020).
 194. The distinction between personal and non-personal decisions is further complicated by the prospect of multi-stage decisions: R. Binns & M. Veale, *Is that Your Final Decision? Multi-Stage Profiling, Selective Effects, and Article 22 of the GDPR*, 11 *International Data Privacy Law* 4 (Nov. 2021), available at <https://doi.org/10.1093/idpl/ipab020>. In the tax domain, one might think of applications that decouple the processing of personal data from the decision-making, ensuring decisions only deal with entirely anonymised data and therefore fall outside the scope of data protection law.
 195. ECJ case law only extends data protection law to legal persons to the extent that a legal person is, in fact, an official label to an identifiable individual or group – such as one-person and family-run enterprises. See, e.g., L.A. Bygrave & L. Tosoni, *Article 4(1). Personal Data*, in *The EU General Data Protection Regulation (GDPR): A Commentary* (C. Kuner, L.A. Bygrave & C. Docksey eds., Oxford University Press 2020). As Recital 14 states, this decision is not an oversight by the legislator, but a constitutive property of the EU data protection regime. However, Member States are not blocked from establishing norms to govern the processing of data referring to identified or identifiable legal persons, and indeed Austria’s Data Protection Act (DSG) (BGBl. I 165/99) extends a narrow set of data protection legal rights to legal persons. Nevertheless, the authors were not able to find any jurisdiction that extends the right to an explanation to corporations and other non-natural entities.

corporation or any other legal person, it is not processing personal data. Hence, the GDPR provides no right to an explanation to the corporation, even if this non-personal information is used to make a decision that adversely affects this legal person, such as flagging potentially unpaid tax.

Still, the right to an explanation is not *entirely* absent from tax contexts, as there are various applications in which a tax system might process personal data. In some cases, the processing of personal data may be central to the task of an AI system, especially if the system is used to assess the tax status of a natural person. In other cases, personal data might be processed as part of the operation of an AI system directed at a corporation. A tax fraud detector might check which natural persons receive payments from the corporation, while an information-providing chatbot might need to assess a user’s identity before providing information. To the extent that data about natural persons is used in decision-making, the natural persons to which this data refers have the right to request an explanation of the decision. But, since this right is held by the data subject, it cannot be directly exercised by a legal person, even if the latter is the ultimate target of data processing.

Under articles 13-15 of the GDPR, the right to an explanation applies only to situations in which there is automated decision-making under article 22 of the GDPR.¹⁹⁶ This definition has two important consequences for our analysis. First, the relevant provisions require the provision of “meaningful information about the logic involved” in the decision-making process rather than an explanation of the systems used for making the decision. Explaining the inner workings of an AI system is only relevant to the extent that it provides meaningful information about “the logic involved in the decision”.¹⁹⁷ This focus on individual decisions might, however, be insufficient to capture systemic patterns in decision-making, such as those leading to algorithmic discrimination.¹⁹⁸ As a result, the GDPR situates explanation at the level of the decision and not at the level of the system.¹⁹⁹

The second consequence is that not all AI applications qualify as automated decision-making. Article 22(1) of the GDPR states that the provisions of this article are applicable only for decisions *solely* based on the automated processing of personal data.²⁰⁰ A narrow interpretation of this provision thus means that explanation is only required in situations in which

.....

196. This article, titled “Automated individual decision-making, including profiling”, establishes a general right not to be subject to an automated decision that produces legal or otherwise significant effects upon a data subject. This right does not apply in certain cases listed by art. 22(2), but lawful decision-making in these circumstances requires the adoption of safeguards listed in art. 22(3). For an overview of art. 22 GDPR, see L.A. Bygrave, *Article 22. Automated Individual Decision-Making, Including Profiling*, in *The EU General Data Protection Regulation (GDPR): A Commentary* (C. Kuner, L.A. Bygrave & Docksey, Christopher eds., Oxford University Press 2020).

197. Arts. 13(2)(f), 14(2)(g), and 15(1)(h) GDPR.

198. On the systemic dimension of the social impact of data and their lack of treatment in the GDPR, see Palka, *supra* n. 174, at sec. II and III.B. Yet, the GDPR includes provisions that attempt to provide a view of decision-making impacts beyond individual decisions, such as those governing Data Protection Impact Assessments: E. Kosta, *Article 35. Data Protection Impact Assessment*, in *The EU General Data Protection Regulation (GDPR): A Commentary* (C. Kuner, L.A. Bygrave & Docksey, Christopher eds., Oxford University Press 2020).

199. Other jurisdictions follow a different path: for example, art. 49 of Brazil’s General Data Protection Law (LGPD) requires the adoption of technological measures that apply to the data processing system itself, in addition to those directed at specific data processing operations.

200. Since all mentions of automated decision-making refer to art. 22 GDPR, and there is no explicit definition of “automated decision-making” as a term, it follows that the contents of art. 22(1) provide an implicit definition of what an automated decision is for the purposes of the Regulation.

decisions are made without human involvement. So far, EU legal scholarship and data protection authorities have embraced a more extensive interpretation, which also includes situations in which human involvement is merely nominal: for example, having a human overseer sign off on every decision made by a machine is not enough if the overseer cannot change the outcome of the process.²⁰¹ Automated decision-making in the GDPR therefore refers to all situations in which humans are not meaningfully involved in producing a decision based on personal data.

This broader interpretation of article 22 nevertheless excludes several forms in which AI systems are used – or can be used – in the tax domain. In many applications, AI systems are embedded within complex organizational structures, which break down important decisions into multiple decisions on a smaller scale. Under such circumstances, it can be difficult to map out the decision that impacted a data subject and whether such a decision is or is not automated.²⁰² Identifying the existence or absence of meaningful human participation in a decision-making process can also be difficult in light of the conditions under which a decision happens in practice.²⁰³ Consider the case of a tax assessor that uses automation to tackle a considerable workload. As a judicious professional, this assessor is able and willing to overrule suggestions made by an AI system but will still be guided by the generated input, especially if the automation has produced reliable results in the past.²⁰⁴ Are such situations automated decision-making?

The existence of automated decision-making under the GDPR was a crucial point of contention in *SyRI*. The plaintiffs argued that the submission of a risk report qualified as automated decision-making even if that report was assessed by the recipients (paragraph 6.57), a claim that the State of the Netherlands did not accept (paragraph 6.58). The court, however, refrained from assessing whether *SyRI* constitutes automated decision-making under the GDPR (paragraph 6.60), relying instead on article 8 of the ECHR as the grounds for its judgment on the merits, as examined above. Yet, the Hague Court pointed out that the *SyRI* legislation did not afford sufficient insight into the risk model and indicators (paragraph 6.89), thus blocking the evaluation of the decision-making process (paragraph 6.90). As a

-
201. On what counts as meaningful human participation, see M. Brkan, *Do Algorithms Rule the World? Algorithmic Decision-Making and Data Protection in the Framework of the GDPR and Beyond*, 27 *International Journal of Law and Information Technology* 2 (2019), available at <https://doi.org/10.1093/ijlit/eay017>; Almada, *supra* n. 188; Article 29 WP, *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679*, wp251rev.01 (Article 29 Data Protection Working Party 2018).
202. Binns & Veale, *supra* n. 194, map five complications involved in assessing decision-making in multiple stages, which relate to the impact intermediate decisions have on the final outcome, the difficulty in pinpointing the relevant decision, and the possibility of selective automation.
203. Almada, *supra* n. 188; B. Wagner, *Liable, but Not in Control? Ensuring Meaningful Human Agency in Automated Decision-Making Systems*, 11 *Policy & Internet* 1 (2019), available at <https://doi.org/10.1002/poi3.198>.
204. On matters of trust in AI decisions and situations, this might be unwarranted; see S. Alon-Barkat & M. Busuioc, *Human-AI Interactions in Public Sector Decision-Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice*, 2022 *JPART*, available at <https://doi.org/10.1093/jopart/muac007>; C. Bartneck et al., *Trust and Fairness in AI Systems*, in *An Introduction to Ethics in Robotics and AI* (C. Bartneck et al. eds., SpringerBriefs in Ethics, Springer International Publishing 2021), available at https://doi.org/10.1007/978-3-030-51110-4_4; M. Dymitruk, *The Right to a Fair Trial in Automated Civil Proceedings*, 13 *Masaryk University Journal of Law and Technology* 1 (2019).

result, the case does not provide a clear-cut answer as to what degree of human involvement is needed to avoid the label of automated decision-making and its implications.²⁰⁵

Beyond automated decision-making, another pathway that could lead to a right to an explanation can be found in *data protection by design*. Under article 25(1) of the GDPR, data controllers must adopt technological and organizational measures to implement data protection principles and safeguards to protect the rights of data subjects. For the current distinction, the principle of *lawfulness* establishes a relevant connection with the informational requirements in tax law.²⁰⁶ Since data processing must comply with all applicable legal requirements, the design of data processing must account not just for those requirements stemming from data protection law but also for those derived from the tax domain.²⁰⁷ Any use of AI in the tax domain must therefore include measures to ensure that the introduction of AI technologies does not prevent compliance with existing duties of transparency.

Explanation of an AI system provides information about that system’s operation, as discussed above. As such, the adoption of explainable AI techniques could be a way to implement the design measures required under article 25(1) of the GDPR. They are particularly useful when it comes to challenging decisions involving AI systems – through the courts, administrative authorities, or other channels for complaint – as exercising the right to contest a decision depends on the existence of suitable information about what has been decided in the first place.²⁰⁸ Nevertheless, data protection by design does not mandate the adoption of any specific measures, adopting instead a risk-based approach in which data controllers are required to map the risks to data subject rights that a given operation produces and then adopt technological and organizational measures to address those risks.²⁰⁹ Data controllers therefore have considerable leeway to avoid explaining so long as they provide alternative mechanisms that ensure the adequate protection of data subject rights.

From the analysis above, the authors conclude that current data protection law does not establish a general duty of explanation for AI systems, even when these systems depend on personal data. Nevertheless, the vast scholarship on XAI and data protection provides two main contributions to the present study. From a conceptual perspective, research on the right to an explanation under European data protection law suggests that explanations are a useful tool for ensuring that decisions involving AI can be challenged and proposes legal and technological mechanisms that might be transposable into tax contexts. From

205. In the other cases thus far dealing with art. 22 GDPR, the automated character of a decision has not been put into question. See, e.g., the *Ola* and *Uber* cases discussed above (*supra* n. 189 and accompanying text), in which the defendant questioned not the characterization of automated decision-making, but whether the decisions produced the legal or similarly significant effect required for application of art. 22. This requirement further constrains the scope of explanation derived from art. 22 GDPR, but it will not be examined in this article. See, however, Bygrave, *supra* n. 196.

206. Art. 5(1)(a) GDPR.

207. On the interpretation of the lawfulness principle, see C. de Terwangne, *Article 5. Principles Relating to Processing of Personal Data*, in *The EU General Data Protection Regulation (GDPR): A Commentary* (C. Kuner, L.A. Bygrave & C. Docksey eds., Oxford University Press 2020).

208. On the connection between art. 22 GDPR and the broader right to judicial contestation of decisions (automated or not), see Sarra, *supra* n. 11.

209. For an overview of the risk-based approach adopted in art. 25 GDPR, see M. Almada, J. Maranhão & G. Sartor, *Article 25. Data Protection by Design and by Default*, in *General Data Protection Regulation. Article-by-Article Commentary* (I. Spiecker gen. Döhmman, V. Papakonstantinou, G. Hornung, et al. eds., Beck; Nomos; Hart Publishing 2022); L.A. Bygrave, *Article 25. Data Protection by Design and by Default*, in *The EU General Data Protection Regulation (GDPR): A Commentary* (C. Kuner, L.A. Bygrave & C. Docksey eds., Oxford University Press 2020).

an instrumental perspective, it provides natural persons in the tax domain (individual taxpayers) with mechanisms they can use to assert their rights, but only to the extent that decisions are based on their personal data. When it comes to explanations in tax, therefore, data protection law should not be seen as a direct source of law but as a fellow traveller in the mission of ensuring the lawful use of AI systems.

4.1.2. The AI Act

The AI Act was proposed by the European Commission in April 2021,²¹⁰ building upon the work developed in recent years by the High-Level Expert Group on AI.²¹¹ It adopts a tiered, risk-based classification of AI applications: some uses are proscribed,²¹² and some of the remaining applications are deemed to impose significant risks to the health, safety and fundamental rights of persons.²¹³ AI systems that fall into the latter category of high-risk systems²¹⁴ are subject to product compliance requirements that closely follow the *new legislative framework* used to provide harmonized standards for products commercialized in the European Union.²¹⁵ This approach has been described by the European Commission as a key element of the EU approach to digital regulation,²¹⁶ but it may still undergo considerable change through its legislative procedure.²¹⁷ Nevertheless, the original proposal by the Commission suggests regulatory trends that warrant analysis in this article.²¹⁸

Unlike the GDPR, the AI Act is not grounded on the fundamental right to data protection but on the European Union's competence to regulate the single market.²¹⁹ This change in legal grounds ensures that the AI Act applies to contexts that fall outside the scope of data protection law. Various tax-related functions are indeed covered by Annex III of the AI

-
210. AI Act (Commission Proposal) (2021). As of February 2022, the Slovenian and French presidencies of the Council of Ministers have proposed compromise texts that alter certain provisions of the original commission text but do not lead to substantial changes to the discussion in this section. As such, unqualified references to "AI Act" refer to the text as proposed by the commission, whereas references to the Commission document point to the explanatory memorandum.
211. AI HLEG, *A Definition of AI: Main Capabilities and Disciplines* (Independent High-Level Expert Group on Artificial Intelligence, European Commission 2019); AI HLEG, *Ethics Guidelines for Trustworthy AI* (Independent High-Level Expert Group on Artificial Intelligence, European Commission 2019).
212. Art. 5 AI Act.
213. AI Act (Commission Proposal) (2021), p. 4.
214. Art. 6 AI Act. Annexes II and III of the AI Act detail the categories that are deemed to be high-risk applications.
215. M. Veale & F.Z. Borgesius, *Demystifying the Draft EU Artificial Intelligence Act — Analysing the Good, the Bad, and the Unclear Elements of the Proposed Approach*, 22 Computer Law Review International 4 (1 Aug. 2021), available at <https://doi.org/10.9785/crl-2021-220402>.
216. AI Act (Commission Proposal) (2021), p. 1.
217. For more information about the procedure the AI Act must follow if it is to become law in the European Union, see, inter alia, T.S. Cabral, *A Short Guide to the Legislative Procedure in the European Union*, 6 UNIO – EU Law Journal 1 (5 July 2020), available at <https://doi.org/10.21814/unio.6.1.2711>.
218. On the AI Act and its connections with trends in EU regulation of digital affairs, see, inter alia, L. Edwards, *Regulating AI in Europe: Four Problems and Four Solutions* (Ada Lovelace Institute 2022); G.D. Gregorio & P. Dunn, *The European Risk-Based Approaches: Connecting Constitutional Dots in the Digital Age*, 59 Common Market Law Review 2 (2022); Veale & Borgesius, *supra* n. 215.
219. AI Act (Commission Proposal) (2021), p. 7. Under art. 114(2) TFEU, the residual competence for market harmonization does not extend to "fiscal provisions", a restriction that might raise questions about the applicability of the AI Act to fiscal matters. However, the same provision also precludes provisions "relating to the rights and interests of employed persons", which are nevertheless present in Annex III of the AI Act. For the sake of argument, this section assumes a narrow reading of "fiscal provisions" that does not include the use of AI in tax administration; an in-depth examination of this point would, however, exceed the scope of this article.

Act, which provides a list of applications deemed to pose a high risk to fundamental rights. Point 5(a) of this annex states that “AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for public assistance benefits and services, as well as to grant, reduce, revoke, or reclaim such benefits and services” are classified as high-risk systems. While such benefits are not necessarily related to taxation, this formulation covers the use of AI systems for assessing matters such as need-based tax credits.²²⁰

Annex III also establishes that a broad range of applications of AI in law enforcement contexts imposes a high risk to fundamental rights.²²¹ Those applications become relevant for taxation to the extent that AI systems are deployed as part of law enforcement action directed at tax crimes, for example, by profiling potential tax fraudsters.²²² Finally, the AI Act also defines as high-risk those AI systems used to assist judicial authorities in “researching and interpreting facts and the law and in applying the law to a concrete set of facts”,²²³ which means the safeguards of Section III of the AI Act apply to judicial activity in tax-related cases.²²⁴ The definition of “high-risk system” in the AI Act thus covers various forms through which states exercise their power in the tax domain.

Classification as a high-risk system leads to the application of specific regulatory requirements.²²⁵ Quite a few of these provisions are directed at increasing the transparency of high-risk systems, such as the need to automatically register events that occur during a system’s operation²²⁶ or to register a system’s technological properties – such as accuracy – and risks within its technical documentation.²²⁷ The mere existence of this information in an organized form facilitates the assessment of a system by judicial authorities and potentially by administrative bodies. In addition, the AI Act also provides information to other actors: the users of AI systems acquired or licensed from third-party providers²²⁸ and the humans who act as overseers to the decision-making processes involving AI systems.²²⁹

220. S. Buijsman & H. Veluwenkamp, *Spotting When Algorithms Are Wrong*, 2022 Minds & Machines, available at <https://doi.org/10.1007/s11023-022-09591-0>; Hadwick & Lan, *supra* n. 6.

221. Point 6 of Annex III to the AI Act. Recital 38 of the AI Act states that systems “specifically intended to be used for administrative proceedings by tax and customs authorities” should not be considered as being used by law enforcement authorities for such purposes. Given the use of “specifically” in this formulation, this recital does not exclude systems that are regularly used for law enforcement purposes, though it still suggests that occasional data sharing with law enforcement does not render an AI system a high-risk system.

222. Point 6(e) of Annex III AI Act.

223. Point 8(a) of Annex III AI Act.

224. On automation in the courts, see Leenes, Kosta & Kamara, *supra* n. 56; S. Lim, *Judicial Decision-Making and Explainable Artificial Intelligence*, 33 Singapore Academy of Law Journal 280 (2021); T. Sourdin, *Ethical Issues in Judge AI and Judicial Technology Use*, in *Judges, Technology and Artificial Intelligence* (Edward Elgar Publishing 2021), available at <https://doi.org/10.4337/9781788978262.00013>.

225. For comprehensive analyses of the governance system that Chapter 2 AI Act establishes for high-risk AI systems, see M. Ebers, *Standardizing AI - The Case of the European Commission’s Proposal for an Artificial Intelligence Act*, in *The Cambridge Handbook of Artificial Intelligence: Global Perspectives on Law and Ethics* (L.A. DiMatteo, M. Cannarsa & C. Poncibò eds., Cambridge University Press 2022); N. Smuha et al., *A Response to the European Commission’s Proposal for an Artificial Intelligent Act* (LEADS Lasb @ University of Birmingham 2021); Veale & Borgesius, *supra* n. 215.

226. Art. 12 AI Act.

227. Art. 11 AI Act.

228. Art. 13 AI Act.

229. Art. 14 AI Act.

Yet, this broad set of transparency measures does not require explanations of high-risk AI systems or their output. Under article 11 of the AI Act, a system's technical documentation must provide its design specifications, including but not limited to "the general logic of the AI system and of the algorithms".²³⁰ A sufficiently detailed description of this logic might amount to an explanation, as discussed in the context of the GDPR above. Still, such an explanation will be global, as it focuses on the AI system in its entirety and not any specific decisions. Furthermore, any such explanation provides an *ex ante* view of the overall logic applied when the documentation was written (or last revised),²³¹ which might not reflect the logic used for any specific decision. There is no obligation to provide local post hoc explanations of specific decisions made with the involvement of an AI system.

Nevertheless, the adoption of local post hoc explanations can be useful for complying with monitoring requirements. Under article 9(2) of the AI Act, high-risk AI systems must be accompanied by an effective risk management system that allows for the identification and analysis of known and foreseeable risks, as well as for the estimation and evaluation of risks that arise during operation. Likewise, article 14(3)(a) requires the identification and construction of technically feasible measures to ensure effective human oversight before a high-risk system is placed on the market or put into service. Since article 14(4)(c) states that human overseers must be able to correctly interpret AI output, post hoc explanations can be used to meet this requirement. There is no general obligation to adopt such measures, however: as discussed in section 4.1.1. concerning data protection by design, these risk-based approaches do not mandate the adoption of specific measures. Therefore, post hoc explanation measures are only required if there is no other way to address risk in a particular context of application.

Even those transparency requirements that *do* exist under the AI Act are subject to limitations. By design, they only apply to high-risk systems,²³² a classification that excludes applications in the tax domain beyond those described above. None of the rubrics under Annex III of the AI Act cover, for example, the use of AI to calculate the amount of tax owed in a given context or the use of AI by taxpayers, which is likely to account for a substantial share of tax AI applications. In addition, many of the transparency requirements are directed toward the "user" of the system, that is, "any natural or legal person, public authority, agency or other body using an AI system under its authority".²³³ This means taxpayers – individual or otherwise – cannot invoke the transparency requirements laid down in the AI Act to obtain information about the systems used by tax authorities.²³⁴ Explanation under the AI Act is directed towards the users who deploy AI systems for tax applications, not the taxpayer.

230. Point 2(b) of Annex IV AI Act.

231. Under art. 11 AI Act, documentation must be drawn up before the system is placed on the market or put into service and kept up-to-date thereafter.

232. Art. 52 AI Act sets up transparency obligations for systems that interact directly with natural persons, such as chatbots, for emotional recognition systems and biometric categorisation systems and for AI systems that generate so-called "deep fakes". Unlike the bulk of the measures laid down in the AI Act, these specific requirements also apply to systems not deemed to pose a high risk to fundamental rights, but their relevance for tax applications is somewhat narrow. Therefore, this article will not analyse those specific requirements.

233. Art. 3(4) AI Act.

234. Art. 60 establishes a database of high-risk AI systems, to be maintained by the European Commission, which is accessible to the public. Annex VIII AI Act lists all information this database must contain, but this list does not include any form of explanation of AI decisions.

To sum up the previous discussion, the AI Act requires general ex ante explanations of high-risk systems and stimulates local post hoc explanations as part of its high-risk management system. Such measures, however, do not apply to all uses of AI in the tax domain, and they offer little recourse for taxpayers. Consequently, the AI Act does not provide grounds for a general right to an explanation of tax decisions concerning a taxpayer.²³⁵

4.2. *Non-legally binding guidelines*

In this section, the authors briefly present and analyse the various international recommendations that provide general standardization rules for using AI systems. At the national level, various countries have adopted some form of national AI policy,²³⁶ and some countries have also produced guidance documents that aim to encourage the use of XAI technologies. A particularly relevant example of the latter is a recent report from the US National Institute of Standards and Technology (NIST), which establishes explainability as a key requirement for trust in AI systems²³⁷ and defines four principles that should guide XAI. It should provide an explanation by delivering evidence or reasons for the output or processes at hand,²³⁸ and such explanations should be meaningful for their intended recipients,²³⁹ accurately depict the decision-making process and reasons why a system operates as it does or produces a given output²⁴⁰ and identify and declare the limits of the knowledge that guides system operation.²⁴¹ In addition to these general principles, national documents have included reviews of the state of the art in XAI,²⁴² frameworks for explanation²⁴³ and proposals for setting up technical standards for explainability,²⁴⁴ although none of these focuses on the specific challenges of AI in the tax sector.

Internationally, various (inter)governmental organizations and commissions have taken steps toward a regulatory framework for AI, such as the African Commission on Human

235. The inclusion of provisions on a right to explanation is, in fact, a major demand by civil society organizations: EDRi et al., *An EU Artificial Intelligence Act for Fundamental Rights. A Civil Society Statement*. (30 Nov. 2021), available at <https://www.accessnow.org/cms/assets/uploads/2021/11/joint-statement-EU-AIA.pdf> (accessed 3 Dec. 2021).

236. For an overview of approaches to AI throughout the world, see M. Rotenberg, M. Hickok & K. Caunes eds., *Artificial Intelligence and Democratic Values Index* (Center for AI and Digital Policy 2022).

237. P.J. Phillips et al., *Four Principles of Explainable Artificial Intelligence 1*, National Institute of Standards and Technology (29 Sept. 2021), available at <https://doi.org/10.6028/NIST.IR.8312> [hereinafter *NIST Four Principles of XAI* (2021)]. Similarly, see *FEAT Principles Assessment Methodology*, White Paper Document 3, p. 47 (Veritas Phase 2, Monetary Authority of Singapore 2021); *German Standardization Roadmap on Artificial Intelligence 201* (DIN & DKE 2020) [hereinafter *German AI Roadmap* (2020)].

238. *NIST Four Principles of XAI* (2021), *supra* n. 237, at p. 3.

239. *Id.*, at pp. 3-4.

240. *Id.*, at pp. 4-5.

241. *Id.*, at p. 5.

242. See, e.g., *id.*, at pp. 6-29.

243. See, e.g. *Transparency Principles Assessment Methodology*, White Paper Document 3C (Veritas Phase 2, Monetary Authority of Singapore 2021); D. Leslie & M. Briggs, *Explaining Decisions Made with AI: A Workbook (Use Case 1: AI-Assisted Recruitment Tool)* (The Alan Turing Institute 2021).

244. See, e.g. UK National AI Strategy, p. 56; *German AI Roadmap* (2020), *supra* n. 237, at p. 59.

and Peoples' Rights,²⁴⁵ UNESCO,²⁴⁶ the Council of Europe²⁴⁷ and the OECD.²⁴⁸ Given that the recommendations made by the Council of Europe and the OECD have been recognized in judicial proceedings and by the G20, respectively,²⁴⁹ the authors have chosen to analyse those recommendations as part of the international non-legally binding guidelines.

Those sources can be seen as "soft law" capable of influencing law-making bodies and serving as a source of inspiration when interpreting hard law.²⁵⁰ By referring to the delegation theory, scholars explain that "states choose soft law when they are uncertain whether the rules they adopt today will be desirable tomorrow and when it is advantageous to allow a particular group of states to adjust expectations in the event of changed circumstances".²⁵¹ Soft law is thus well-suited to the rapid evolution of AI systems. Over time, soft law standards could morph into hard law regulations once the former proves effective. At the level of interpretation, soft law regarding AI has already proved to be useful, as follows, for example, from the Slovak Constitutional Court's judgment in *eKasa*, in which the Court broadly referred to non-legally binding Recommendation CM/Rec(2020)1 of the Committee of Ministers to Member States on the human rights impacts of algorithmic systems. The following analysis aims to verify whether the selected non-legally binding sources of standards for the implementation and application of AI systems could prove useful in setting a minimum legal standard for XAI in tax law.

4.2.1. Recommendation of the OECD Council on Artificial Intelligence

The Recommendation of the Council on Artificial Intelligence²⁵² (the OECD AI Recommendation) emphasizes that AI actors²⁵³ should commit to transparency and responsible disclosure regarding AI systems.²⁵⁴ The OECD Recommendation requires AI actors to provide meaningful information appropriate to the context in which an AI system is used to (i) foster a general understanding of AI systems; (ii) make stakeholders aware of their interactions with AI systems; (iii) enable those affected by an AI system to understand the out-

245. ACHPR, *Resolution on the Need to Undertake a Study on Human and Peoples' Rights and Artificial Intelligence (AI), Robotics and Other New and Emerging Technologies in Africa* (African Commission on Human and Peoples' Rights 2021).

246. UNESCO, *The Recommendation on the Ethics of Artificial Intelligence* (United Nations Educational, Scientific and Cultural Organization 2020).

247. Council of Europe, Recommendation CM/Rec(2020)1 of the Committee of Ministers to Member States on the Human Rights Impacts of Algorithmic Systems (8 Apr. 2020) [hereinafter Recommendation CM/Rec (2020)1 (2020)].

248. OECD, *Recommendation of the Council on Artificial Intelligence* (22 May 2019), available at <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449> (accessed 3 Mar. 2022) [hereinafter *OECD AI Recommendation* (2019)].

249. See the Slovak Constitutional Court's judgment of 17 December 2021 in the *eKasa* case, discussed in sec. 3.1.; G20, *G20 AI Principles* (9 July 2019).

250. The OECD Model and its Commentaries serve as important non-legally binding documents for the implementation of hard tax treaty laws and their interpretation worldwide: E.A. Baistrocchi, *The International Tax Regime and the BRIC World: Elements for a Theory*, 33 Oxford Journal of Legal Studies 4, p. 744 (2013).

251. A.T. Guzman & T.L. Meyer, *International Soft Law*, 2 Journal of Legal Analysis 1 (2010).

252. *OECD AI Recommendation* (2019), *supra* n. 248.

253. The OECD defines AI actors as organizations and individuals that deploy or operate AI and other entities playing an active role in the AI system lifecycle (*OECD AI Recommendation* (2019), *supra* n. 248, at p. 7).

254. The OECD defines an AI system as a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments (*OECD AI Recommendation* (2019), *supra* n. 248, at p. 7). An AI system capable of generating such output satisfies our broad definition of an AI system (see *supra* n. 1).

come; and (iv) enable those adversely affected by an AI system to challenge its outcome.²⁵⁵ The relevance of the OECD AI Recommendation has been recognized by the G20, which has drawn inspiration from this recommendation to establish the G20 AI Principles.²⁵⁶

All these points appear relevant for the explainability of AI systems insofar as they require deployers and users of AI systems to ensure that everybody affected by the outcome of such systems receives a relevant explanation in the given circumstances. This includes the use of AI systems in all possible settings, including tax matters. Since all deployers of AI systems are AI actors, entities from the private sector should ensure that the AI systems they provide to the tax authorities can be used in a way that is understandable both to taxpayers and the authorities themselves. Thus, whenever the tax authorities rely on AI systems to reach decisions affecting taxpayers, the tax authorities should give each taxpayer an explanation understandable to the taxpayer that clarifies why and how the AI systems led to a particular tax decision. Moreover, a taxpayer should have the right and the opportunity to challenge any negative tax decision based on an AI system. This is in line with constitutional rights and principles and the right to a fair trial under the ECHR, as analysed in sections 4.1. and 4.2.1.

The aforementioned recommendations could trigger tensions in terms of tax and trade secrecy, as discussed in section 2.2. Since they, as opposed to rules regulating tax and trade secrecy, are not legally binding, the OECD AI Recommendation is not likely to be effective until it is implemented as hard law, with the resultant lifting of secrecy for the benefit of explainability.

4.2.2. *Council of Europe Recommendation on the human rights impacts of algorithmic systems*

Council of Europe Recommendation CM/Rec (2020)1²⁵⁷ formulates a set of guidelines regarding the design, development and ongoing deployment of algorithmic systems.²⁵⁸ One of the pivotal points under this Recommendation refers to the explainability of the processes and output of algorithmic systems,²⁵⁹ which is addressed in the context of the transparency of algorithmic systems and the contestability of automated decisions.

These points are relevant to XAI insofar as they require the developers of AI systems to develop AI systems that are technically capable of providing explanations and impose duties on the users of AI systems to ensure that those affected by an automated decision can effectively challenge those decisions. The formulated guidelines do not exclude tax authorities. Council of Europe Recommendation CM/REC (2020)1 explicitly states that the guidelines are designed to advise “States, and public and private sector actors”.²⁶⁰ The points on

255. OECD AI Recommendation (2019), *supra* n. 248, at p. 8.

256. G20, *supra* n. 249.

257. Recommendation CM/Rec (2020)1 (2020).

258. *Id.*, p. 5.

259. Within the context of Recommendation CM/Rec (2020)1, algorithmic systems are understood as “applications that, often using mathematical optimisation techniques, perform one or more tasks such as gathering, combining, cleaning, sorting, classifying and inferring data, as well as selection, prioritisation, the making of recommendations and decision-making” (p. 5). An algorithmic system capable of generating decisions and making recommendations falls within the scope of our definition of AI systems as established in *supra* n. 1.

260. *Id.* The term “actors” has a broad scope, including any public and private sectors “in all their actions regarding the design, development and ongoing deployment of algorithmic systems”.

explainability do not distinguish between algorithmic systems used by States or by private sector actors and even call upon States to establish appropriate levels of transparency “with regard to the public procurement, use, design and basic processing criteria and methods of algorithmic systems implemented by and for them”. Therefore, the guidelines established for XAI can be considered relevant in tax matters, meaning that AI systems used by tax authorities should be capable of technical explainability as well as legal explainability. This is in line with the constitutional principles analysed in section 3.1. As far as contestability is concerned, Council of Europe Recommendation CM/REC (2020)1 requires that the explainability of processes and output be held to high standards where algorithmic systems are used in decision-making processes that carry high human rights risks.²⁶¹ This is relevant for XAI in tax law insofar as tax authorities use AI systems to detect tax fraud, as the explainability of those systems safeguards the right to a fair trial, as analysed in section 3.2.

It is noteworthy that Council of Europe Recommendation CM/Rec (2020)1 emphasizes that legislative frameworks for intellectual property and trade secrecy should not deprive AI systems of their explainability.²⁶² That is indeed very important for the effective explainability of AI systems, but as long as this point of the Recommendation is not implemented as hard law, it will prove difficult to displace trade secrecy for purposes of the explainability of AI systems.²⁶³ Moreover, where AI systems serve tax purposes, the problem of tax secrecy will remain. Council of Europe Recommendation CM/Rec (2020)1 does not address that problem at all.

The relevance of Council of Europe Recommendation CM/REC (2020)1 for XAI in tax law was demonstrated by the Slovak Constitutional Court in its judgment of 17 December 2021 in *eKasa*, which is discussed in section 3.1. By broadly referring to that Recommendation, the Court established a set of rules that directly deal with the legal aspects of XAI.

5. No Minimum Legal Standard for XAI in Tax Law and Fragmentation of the AI Legal Environment

Sections 3. and 4. map out various situations in which taxpayer rights protection requires some explanation. As section 3. shows, current constitutional and human rights frameworks come under strain when inexplainable AI systems are used in tax matters, as the lack of explanation potentially erodes the precision, transparency and predictability expected of tax law and its execution. Indeed, various judicial decisions within the framework of the ECHR provide grounds to maintain that legislators are obliged by case law to ensure that the use of AI does not interfere with principles such as the right to a fair trial, respect for private and family life and non-discrimination.²⁶⁴ While the *Ferrazzini* ruling shows the ECtHR is reluctant to extend the fair trial provisions of the ECHR to tax matters, the *SyRI* judgment by a Dutch court of first instance finds an implicit connection between the violation of the right to respect for private and family life and the violation of the right to a fair trial because of a feature common to all legislation governing the use of AI systems – their lack of explainability.²⁶⁵ In the case at hand, the use of inexplainable tax systems was

261. Recommendation CM/Rec(2020)1, p. 13.

262. *Id.*, at pp. 13 and 15.

263. Possibly only via litigation, as, for instance, the Slovak Constitutional Court’s judgment of 17 December 2021 in the *eKasa* case shows. *See* sec. 3.1.

264. *See* sec. 3.2.

265. *See* sec. 3.2.

deemed to impact the rights of taxpayers to effectively protect themselves from the adverse effects of tax authority decisions, and the adoption of technical explanation measures would therefore appear to be a requirement for ensuring the effectiveness of the ECHR itself. Therefore, even the well-established reticence to apply ECHR provisions to tax law matters is not enough to dispel the relevance of XAI for the lawful use of AI in the tax domain, which suggests that human rights frameworks can play a role in the development of a minimum legal standard for the use of XAI in tax matters.

However, international legal standards have yet to catch up with the constitutional diagnosis presented above. The European Union is generally deemed to be at the forefront of global AI regulation, yet its main legal instruments in that field make no explicit mention of tax matters. As of March 2022, the proposed AI Act applies only to a narrow set of tax-related applications, which include criminal prosecutions but not the everyday tasks of tax administrations, while the GDPR is more broadly applicable, covering all non-law enforcement contexts in which the personal data of natural person taxpayers is processed. While neither legal instrument explicitly requires any form of post hoc explanation, a general right to explanation is to be found in the GDPR, with post hoc explanation models perhaps being the only practical means to meet the transparency requirements imposed by the risk-based approaches in both legal instruments. Consequently, these laws impose only a limited obligation to adopt explanation in tax contexts, which could nevertheless be useful as an inspiration for defining the contours of explanation in other jurisdictions or contexts beyond the scope of these norms.

Some of the non-legally binding sources of recommendations on AI systems specifically point to the need for XAI to protect fundamental human rights. The OECD AI Recommendation and Council of Europe Recommendation CM/Rec (2020)¹²⁶⁶ both require deployers and users of AI systems – in all settings, including tax matters – to ensure that everybody affected by the outcomes produced by such systems receives a relevant explanation and has the right to effectively challenge a negative outcome. Trade secrecy cannot stay in the way of achieving this end, as acknowledged in Council of Europe Recommendation CM/Rec (2020)¹²⁶⁷. Nevertheless, this postulate could end up being respected by some courts²⁶⁷ while ignored by others, especially if they are non-European Council members,²⁶⁸ as long as the legally binding rules do not displace trade secrecy for explainability purposes. This problem also extends to all recommendations found in national instruments that, to varying extents, deal with the need for the explainability of AI systems at the technological level.

Non-legally binding guidelines do not address the issue of tax secrecy, perhaps because none of them deals specifically with AI systems in tax law. The fact remains, however, that the problem of XAI in tax law is either completely overlooked, or the lack of transparency on the part of the tax authorities – and, thus, their efficiency – appears to be more important under the rule of law than the protection of fundamental human rights.

All in all, the analysis in sections 3. and 4. reveals that there is no minimum legal standard of XAI in tax law, and we are probably far away from having one. Currently, we observe huge fragmentation in AI regulations, with the European Union at the forefront of legal

.....

266. See secs. 4.2.1. and 4.2.2.
267. E.g. the Slovak Constitutional Court’s judgment of 17 December 2021 in the *eKasa* case. See *supra* sec. 3.1.
268. E.g. the Wisconsin Supreme Court’s judgment of 13 July 2016 in the *Loomis v. Wisconsin* case: Washington, *supra* n. 101.

initiatives. There is, however, a glaring lack of regulations devoted to the field of tax law, particularly in respect of XAI. Although the threats posed by inexplainable AI systems in tax law are no less real and serious than those stemming from tax avoidance, the initiatives of individual countries and international organizations in that field are close to nil. The aforementioned threats have most likely been vastly overlooked and misunderstood by tax policymakers up until today. This might follow from the fact that XAI in tax matters can conflict with tax secrecy concerns and the overall trend to increase the efficiency of tax administration via the use of new technologies, including AI systems.²⁶⁹ It is unacceptable that the efficiency of tax administration is allowed to displace the effective protection of fundamental taxpayer's rights due to the lack of explainability of AI systems in tax matters.²⁷⁰

6. Concluding Remarks

Although AI is a global phenomenon with powerful implications for almost every aspect of our daily lives, including taxation, little has been done yet in terms of regulating the use of AI in the least harmful way (harmful to all interested parties, including taxpayers, tax authorities and the developers of AI systems). Tax AI is one of the more conspicuous regulatory gaps in that regard. Tax and trade secrecy legislation that adversely affects tax XAI is the only clear-cut type of legislation to have gained purchase worldwide. The problem with XAI is exacerbated by the technical complexity of the world's most efficient AI systems, which rely on vast amounts of data and machine learning algorithms, typically deep neural networks. The path towards XAI in tax law would therefore appear to be cluttered with unsurmountable legal and technical obstacles. Our analysis shows, however, that these obstacles not only should be overcome but also *can* be overcome, albeit with major efforts on the part of the international research community and national legislators.

More effort needs to be invested in creating procedures that allow the use of new and improved methods such as AI or, for that matter, automated solutions of any kind for decision-making. Currently, there is little pressure on tax administration organizations and legislators to formulate standard procedures. Given the recent litigation, this is likely to change. This means that existing academic research into best practices in model training could contribute to creating procedures appropriate for legal decision-making and legal model training. The ultimate goal is to minimize the potential harm to taxpayers, which could be achieved by following XAI guidelines.

These obstacles should be overcome to avoid protracted litigation with AI systems hovering in the background. The authors appreciate that tax administrations worldwide will not be able to resist increasing their efficiency through AI systems. Still, they must not lose sight of the need to respect taxpayers' rights ensuing from constitutional principles and legal provisions intended to safeguard human rights. Instead, the use of AI systems is promoted by the OECD under the motto of "seamless and frictionless tax administration".²⁷¹ The authors wonder how seamless and frictionless tax administration can be if it uses tax AI systems that are explainable to neither taxpayers nor judges.

269. OECD, *Digital Transformation Maturity Model* (OECD Maturity Model Series, OECD 2021); OECD, *supra* n. 8.

270. Cf. D. Hadwick, *Behind the One-Way Mirror: Reviewing the Legality of EU Tax Algorithmic Governance*, 31 EC Tax Review 4, p. 201 (2022).

271. OECD, *supra* n. 8, at p. 42.

This is not to say that XAI will solve every issue introduced by the opacity of AI in tax applications. Indeed, section 2.2. shows that, in many cases, secrecy can just as easily ensue from the law as from technical complexity. However, the existence of non-technical sources of opacity does not make the technical opacity of AI any less real, which means there is an urgent need for a minimum legal standard of tax XAI. Otherwise, the large gaps and fragmentation in regulations in the field of AI and tax law will continue to grow and steadily trigger even more ferocious disputes and scandals, thereby undermining society's trust in digitalized tax administration. Future research will be needed to support the construction of a regulatory framework that addresses these concerns at the global level, to identify how legal requirements for explanation could be met by XAI technologies and to design adequate hard law and soft law instruments for ensuring the global adoption of suitable XAI. This will make it possible to eliminate the currently huge fragmentation in AI regulations and fill the regulatory vacuum in the tax field. To this end, tax lawyers and computer scientists should work together to help reach a consensus on legal and technical solutions that secure XAI's place in tax law. This ambitious goal is not impossible, but it will require shared multilateral solutions in the form of a global minimum legal standard for XAI in tax law. The authors' future research will aim to contribute to the more technical side of this challenge by focusing on how to select and develop the XAI techniques best suited to the needs of taxpayers and judges, i.e. local and global post hoc explainability techniques. The authors' goal remains to show a way to balance legal requirements and the efficiency of the AI systems for the benefit of the tax administration and taxpayers.