TECHNICAL

NOTES & MANUALS

Understanding Artificial Intelligence in Tax and Customs Administration

Joshua Aslett, Ignacio González, David Hadwick, Stuart Hamilton, Michael Hardy, and Azael Pérez

TNM/2024/06

TECHNICAL NOTES AND MANUALS

Understanding Artificial Intelligence in Tax and Customs Administration

Joshua Aslett, Ignacio González, David Hadwick, Stuart Hamilton, Michael Hardy, and Azael Pérez

Authorized for distribution by Ruud de Mooij

This technical note addresses the following questions:

- What is artificial intelligence (AI), and why is it important?
- Why are legal and ethical matters at the forefront of discussions on AI?
- What is an AI "use case" and what practical examples exist?
- What actions can be taken to promote responsible use of AI?
- How should an AI use case be evaluated and introduced?

This note was prepared by Joshua Aslett, Ignacio González, David Hadwick, Stuart Hamilton, Michael Hardy, and Azael Pérez. The authors acknowledge the valuable feedback and contributions provided by Debra Adams, Doris Akol, Thomas Cantens, Margaret Cotton, Dmitri Jegorov, Andrea Lemgruber, Andrew Masters, John McAlister, Andrew Okello, Herve Tourpe, and Shan Zhong.

Cover Design: IMF Creative Solutions

Cataloging-in-Publication Data
IMF Library

Names: Aslett, Joshua, author. | González García, Ignacio, 1955-, author. | Hadwick, David, author. | Hamilton, Stuart Gordon, author. | Hardy, Michael Andrew, author. | Pérez, Azael, author. | International Monetary Fund, publisher.
Title: Understanding artificial intelligence in tax and customs administration / Joshua Aslett, Ignacio Gonzalez, David Hadwick, Stuart Hamilton, Michael Hardy, and Azael Pérez.
Other titles: Understanding AI in tax and customs administration. | Technical notes and manuals | International Monetary Fund. Fiscal Affairs Department.
Description: Washington, DC : International Monetary Fund, 2024. | Nov. 2024. | Includes bibliographical references.
Identifiers: ISBN:
9798400290435 (paper)
9798400290473 (ePub)
9798400290459 (web PDF)
Subjects: LCSH: Tax administration and procedure—Technological innovation. | Artificial intelligence.
Classification: LCC HJ2305.A8 2024

Recommended citation:
Aslett, Joshua, Ignacio González, David Hadwick, Stuart Hamilton, Michael Hardy, and Azael Pérez. 2024. "Understanding Artificial Intelligence in Tax and Customs Administration." IMF Technical Notes and Manuals 2024/06. International Monetary Fund, Washington, DC.

Please send orders to:
International Monetary Fund, Publication Services
PO Box 92780, Washington, DC 20090, USA
Tel: (202) 623–7430 | Fax: (202) 623–7201
publications@IMF.org
eLibrary.IMF.org
bookstore.IMF.org

Contents

ABBREVIATIONS

AI	artificial intelligence
AIWG	Artificial Intelligence Working Group
ER	expert rule
EU	European Union
GenAI	generative AI
GPT	generative pretrained transformer
HITL	Human-in-the-Loop
HOTL	Human-on-the-Loop
IT	information technology
LLM	large language model
MIT	Massachusetts Institute of Technology
ML	machine learning
MV	machine visioning
NLP	natural language processing
OECD	Organisation for Economic Co-operation and Development
SNA	social network analysis
WCO	World Customs Organization

I. Introduction

This technical note provides an overview of current thinking on artificial intelligence (AI) in tax and customs administration. Written primarily for senior officials, the intent of the note is to provide an awareness of AI that helps to inform decision making. To that end, the note opens by addressing both historic and ongoing AI developments. Recognizing the importance of legal and ethical aspects of AI's use, the note then describes principled concerns and recent regulatory initiatives. The remainder of the note speaks to foundational aspects of AI's management. It presents the concept of an AI "use case," provides guidance on how to promote AI's responsible use, and describes a process for introducing AI into an operational setting. The note concludes by highlighting a selection of open questions around AI being actively explored by experts.

Beyond conveying theory, the annexes to this note include tools designed to help implement important AI management practices. While aspects of AI have long been used in tax and customs administration, the advent of accessible generative AI (GenAI) has significantly raised AI's profile. At the time of writing, GenAI has sparked an aggressive race among suppliers of AI technology to rapidly develop, iterate, and evolve their offerings. One likely effect of this is that regardless of prior applications of AI in tax and customs, AI's use in performing everyday tasks will increase substantially. This is likely whether in direct support of core operations or, as one important example, as an enhancement to common desktop tools used daily by staff. As the use of AI is not without complications, it is increasingly important that it be deliberately managed. To aid in that, Annexes 1–3 provide tools for AI policy, strategy, and risk assessment.

Across the many aspects of AI's use, most contemporary discussions revolve around one underlying idea—certainty—and how to benefit from AI while mitigating its risks. A topic long studied in both tax and customs administration, "certainty of action"—strengthened by clear, predictable, fair, and consistent application of laws and regulations—is vital for the promotion of voluntary compliance, sound financial planning, and economic stability more broadly. Where AI is raising uncertainty is in (1) its potential economic impact on tax bases and trade flows, and the follow-on effects to policy, and (2) its practical implications, utility, and management from an administrative perspective. While the discussion on AI's effects on economics and fiscal policy continues, tax and customs administrations can and should be taking deliberate steps to improve the levels of certainty around their own use of AI tools and technology.

As has long been the case in adapting to new technology, the first step toward improving certainty with the use of AI is to better understand it. In recent decades, the emergence of information technology and digitalization has had profound implications. For those in tax and customs administration, understanding and transitioning to digital operations has been far from smooth and remains an ongoing process. Similarly, AI's expected evolution may be both transformative and challenging. As AI permeates daily life, being informed to make decisions that reduce risks and promote certainty starts with one question: *What, exactly, is artificial intelligence?*

II. Understanding Artificial Intelligence

What Is Artificial Intelligence?

Artificial intelligence (AI) is the simulation of human intelligence using computer systems.[1] Considered a subject within the realm of computer science, AI is an interdisciplinary field that draws heavily from data science, statistics, and a range of other subjects.[2] Accessible through many mediums—desktop and mobile applications, interaction with bots (short for robots), tools for analytics, and others—AI can, with important caveats, be applied to any business problem or task that is sufficiently understood and supported by adequate data. In doing so, AI technology can draw from vast repositories of information, facilitate the exploration and analysis of complex questions, present convincing conclusions, and mimic human-like interactions.

In tax and customs administration, the use of AI spans decades of investments in digital operations. Generalized in Figure 1, these investments have tended to be progressive, building from the automation of core processes in the 1970s and 1980s.[3] During this period, AI emerged in tax and customs with the use of "expert systems" to enhance risk-based targeting of enforcement activities. Subsequently in the late 1990s, machine visioning was introduced in customs.[4] Over the next 30 years, due primarily to the introduction of electronic services, the volume of data captured grew exponentially, leading to wide adoption of tools for advanced analytics, data warehouses, big data platforms, and machine learning (ML) techniques.[5] Since the 2010s, ML has been a category of AI heavily invested in for compliance research and risk analysis.

Figure 1. Digital Operations and AI: 50 Years of Investment

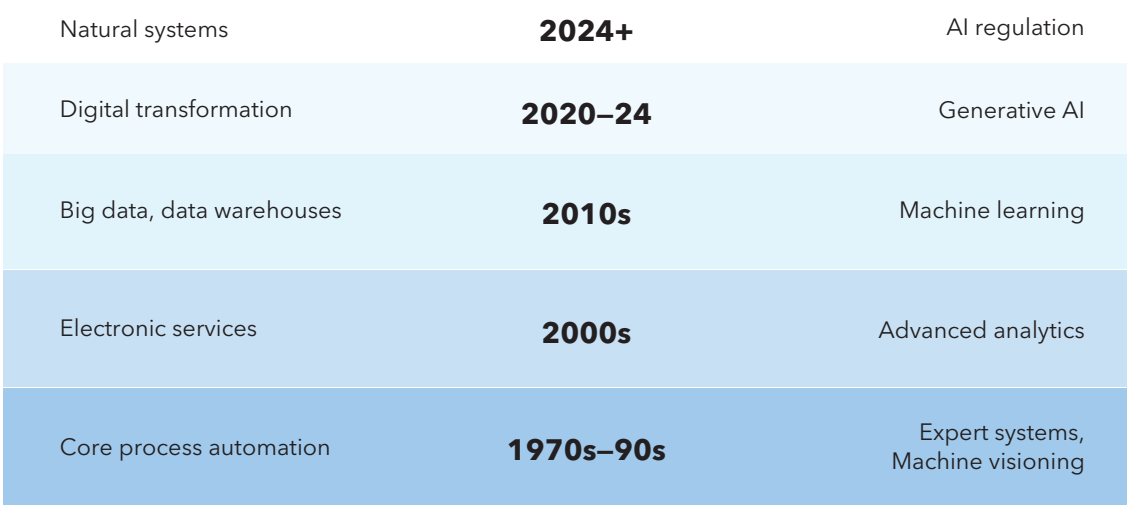

Source: Authors.

[1] The definition of artificial intelligence (AI) is an evolving subject of deliberation and debate. Throughout this note, the use of plain language has been prioritized over technical specificity, as with the definition of AI presented here.

[2] Other subjects include mathematics, neuroscience, psychology, linguistics, and economics.

[3] See González 2018.

[4] Machine visioning is used for purposes that include automated cargo inspection and license plate recognition.

[5] Expert systems build from a legacy of using expert rules and knowledge sourced "top-down" from human subject matter experts. By comparison, machine learning uses algorithms and quantitative methods to learn from data "bottom-up," without explicit programming (or expert rules). In contemporary expert systems, both bottom-up and top-down techniques are integrated and applied.

The AI technology now drawing public attention, GenAI, reflects a major evolution that senior officials must be aware of. GenAI is capable of producing content (text, data, images, videos) comparable with outputs created by humans. Its advancement and integration into daily life are taking place in a period already defined by digital transformation, the reengineering of processes for a digital environment, and concepts of taxation embedded in "natural systems," where tax "just happens."[6] The timing of these changes along with GenAI's arrival—and its accessibility through desktop applications and mobile apps—will elevate AI into reform planning discussions for many administrations.[7] In these discussions, regulatory, legal, and ethical concerns around AI's use will be a priority and the likely next chapter in its adoption.

Because AI is expected to continue to rapidly evolve, understanding what it can and cannot do is vital to making informed decisions. Box 1 explores some of AI's current limits, addressing common myths and the reality of its practical applications.

BOX 1. AI Myths Versus Reality

Myth: AI is mostly hype that can be ignored.

Reality: AI is already powerful and rapidly improving. For example, when questions from professional and academic examinations are supplied to current GenAI tools, the answers that the tools provide score very highly, in many subjects outperforming human test takers. These and similar capabilities are being introduced into commercial desktop applications and mobile apps which, in tandem with other channels for accessing AI, are creating an "AI everywhere / AI always on" scenario likely to heavily influence the work of public servants, including tax and customs officers.

Myth: AI's benefits are overstated and not applicable to the public sector.

Reality: AI offers significant benefits for the public sector, including enhanced efficiency, accuracy, and the ability to uncover insights from large data sets that would be impractical for humans to analyze manually. For tax and customs functions, this could mean more effective fraud detection, improved compliance rates, and more personalized services for both taxpayers and foreign trade operators. As with the adoption of smartphones, GenAI is likely to become ubiquitous in office systems as a tool to enhance productivity.[8]

Myth: AI will lead to large job losses and very soon.

Reality: In tax and customs administrations, AI can streamline operations, reduce errors, and allow staff to focus on higher-value work, such as investigative and strategic decision making. The increase in productivity from the use of generative AI may also lead to an ability to redeploy staff from administrative, back-office tasks to more human contact roles. The scale of this change may be like that experienced by tax and customs administrations in the past when moving from administrative assessment to self-assessment paradigms, and may require a similar level of planning to navigate the choices available, account for costs, and realize potential benefits.

6 A reference to Tax Administration 3.0. See OECD (2020) for details.

7 See WCO (2023) for potential implications specific to customs administration.

8 For example, Copilot, Microsoft's implementation of aspects of ChatGPT 4.0, has been integrated into the Microsoft Office suite of products, such as Microsoft Teams, Excel, Word, and PowerPoint.

Myth: AI can fully replicate human intelligence and decision-making processes.
Reality: AI mimics certain aspects of human intelligence, such as pattern recognition and problem-solving, but currently does not possess self-consciousness or emotional intelligence. AI systems operate within a set of parameters defined by their programming and the amount and quality of the data used for their training, which is often a limitation. This means that AI currently lacks the human capacity for understanding context in a broad, intuitive way, a distinction that is crucial in scenarios where nuanced decision making or ethical considerations are involved.

Myth: AI can surpass human intelligence in all aspects.
Reality: AI excels in processing large data sets and performing some specific tasks faster and more accurately than humans. However, it lacks the moral reasoning and innate ability for creative thinking that humans have. This distinction is crucial as it highlights AI's role as a tool ideal for augmenting human capabilities. For example, AI can already diagnose diseases from medical imaging with high accuracy but relies on human medical professionals for the final judgment and patient care, demonstrating its potential utility in supportive rather than substitutive roles.

Myth: AI tools and technology are self-aware.
Reality: AI systems function within the scope of their programming and the data they have been trained on. They do not possess consciousness or self-awareness, nor do they have an independent goal orientation. AI's decisions are based on algorithms and data patterns, lacking genuine understanding or intent. This is evident in natural language processing applications, where AI can generate human-like text but cannot currently comprehend the meaning of text in the same way that humans do.[9]

Myth: AI is inherently unbiased and objective.
Reality: AI algorithms can perpetuate and even amplify biases present in their training data. This is particularly relevant for tax and customs administrations, where biased AI and machine learning tools could lead to unfair targeting or overlooking certain behaviors (including new or previously unknown forms of fraud not reflected in historic data sets). It underscores the importance of carefully curating training data and continuously monitoring AI systems for biased outcomes, ensuring that automated processes support formal procedures that are fair and equitable.

Source: Authors.

Recent AI Developments

On November 30, 2022, ChatGPT was released, rapidly becoming the technology with the fastest uptake and use in history.[10] In the same broad family of AI-powered tools, such as Apple Siri and Amazon Alexa, OpenAI's ChatGPT is a chatbot capable of interacting in natural language text. Based on input prompts, it can analyze and answer complex questions, as well as generate tailored content (for example, email messages, summaries of documents, or jokes and poetry).[11] The level of technical sophistication in

[9] The distinction between AI and human comprehension may be debatable. See Webson and Pavlick (2022).

[10] See Hu (2023). Two months after launch, ChatGPT reached an estimated 100 million active users.

[11] Siri, Alexa, and ChatGPT all rely on natural language processing, a major field of AI development. A crucial distinction exists in that the technical origins of Siri and Alexa derive from rules-based processing and algorithms, whereas ChatGPT was developed from the start to rely on a type of deep learning AI model.

ChatGPT's first release, and its ability to reason across subjects ("knowledge domains"), represented a major evolution in AI that sparked a race for innovation. In the short time since ChatGPT's arrival, AI capabilities have increased at a pace exceeding historic IT and computing developments, warranting special attention.[12]

Why Is the Arrival of ChatGPT and Similar AI Tools Important?

Advances in newer AI tools have reached a stage where they can now match or outperform median human results in a growing number of knowledge domains.[13] The technology behind leading AI products has been trained on most of what has been written by humans that is digitally available, including much of the content of the internet—far more than any individual could read in a lifetime. With its ability to understand natural language questions and to reason across these vast repositories of knowledge, AI's utility can be tested in human terms. One way to do that is using standardized examinations, providing exam questions to AI tools, and evaluating responses. To illustrate, OpenAI has reported that in a battery of tests, ChatGPT 4.0 scored in the 90th percentile on the Uniform Bar Exam; in the 88th percentile on the Law School Admission Test; and in the 80th, 99th, and 54th percentiles on the Graduate Record Exam Quantitative, Verbal, and Writing tests.[14]

The performance of AI is rapidly improving in many domains, including in key subjects of importance to tax and customs administration. In OpenAI's tests, a substantial improvement was achieved between releases of ChatGPT 3.5 and ChatGPT 4.0, where, for instance, performance on the Uniform Bar Exam increased from the 10th percentile to the 90th. Like the knowledge required to pass a law exam, AI tools will be increasingly capable of drawing from economics, finance, accounting, taxation, trade, and the actual business of tax and customs administration itself. The ability of AI to reason across these domains, inform decisions, and produce helpful content—in the form of responses to questions, advice, guidance, drafts of emails, and memoranda—is likely to place AI at the center of daily office work, a dimension of use for AI very different than in support of research and core operations, but nonetheless important.

Why Is This AI Technology Emerging Right Now?

Increased Data Availability. The ongoing digitalization of transactions and interactions in daily life has resulted in an exponential increase in the amount of data available for ML purposes. AI can process and analyze this "big data" at scale, offering insights that can power tools like chatbots that, in turn, can inform policy, improve compliance, and help streamline operations. In doing so, AI can identify trends and patterns not visible to humans. AI's reliance on data, big or small, is paramount, underscoring the need for robust data management and privacy measures that emphasize legal and ethical considerations in AI's development and deployment.

Advancements in ML Algorithms. Recent improvements in ML, particularly with algorithms using "Transformer" architectures, have led to the development of Large Language Models (LLMs). ChatGPT relies on an LLM called GPT, built using the "Generative Pre-trained Transformer" (GPT) architecture.[15] These technologies have dramatically improved AI's capabilities in language modeling, data analysis, pattern recognition, and predictive modeling. For tax and customs administrations, they contribute to an enhanced

[12] Referring to Moore's law, an observation first made by Gordon Moore in 1965 that the number of transistors on computer chips tended to double approximately every 18–24 months, leading to a significant increase in computing power over time. By comparison, the size of Large Language Models, measured by their number of parameters, has been observed as doubling roughly every 6–12 months.

[13] The median human result in a knowledge domain is the "score" or capability achieved by humans performing at the 50th percentile (half scoring less and half scoring more).

[14] See OpenAI 2023.

[15] Generative pretrained transformer is a type of "deep learning" neural network.

ability to identify tax evasion, smuggling, and fraud patterns much more efficiently (at lower cost) and effectively (with higher accuracy) than legacy approaches to the use of analytics, such as additive risk scoring.[16]

Computational Power. The growth in computational power at lower costs has made it feasible to run complex AI models on enormous data sets, including most of the content on the internet, to build the LLMs underlying the current suite of GenAI tools. Although expensive to initially build and train (costing tens to hundreds of millions of dollars), LLMs and other "foundation models" can sometimes be fine-tuned at a relatively low cost to improve their abilities in particular knowledge domains. This potential is likely to allow tax and customs administrations to eventually develop new, tailored AI systems and tools that benefit from direct integration and use of GenAI technology, bringing with them a wealth of important, new capabilities.[17] One example includes the use of chatbots for taxpayer services, where the chatbots can respond in multiple languages (including languages not normally supported by a tax or customs administration).

Recognizing AI's Systemic Adoption

In the coming years, AI will influence all of the stakeholders involved in taxes and trade. To best appreciate the scale of these changes, a short exploration of some of the scenarios that are already developing can be helpful. The following examples are based on observations made by the authors in different regions of the world. They have been purposefully embellished to highlight human aspects of AI's ongoing and systemic integration into the workplace.

Policymakers. A Member of Parliament is placed on a tax and customs committee. Before a committee meeting, the member instructs their staff to review a new legislative proposal. In the review process, an assistant uploads a copy of the proposal to a GenAI chatbot, inquiring how specific provisions in the legislation may or may not be comparable with those in other similar jurisdictions. The response from the chatbot is detailed, convincing, and accepted at face value. Later, the assistant reads an article about AI "hallucinations" and worries that the chatbot's response may have been incorrect and that unreliable information may have been introduced into discussions.[18]

Taxpayers. In a trade association meeting, a presentation is made on techniques that can be used to reduce taxes. Concerned that the advice could be wrong, a taxpayer at the meeting uses popular search engines to research the validity of the claims made. Unsatisfied with the results, the taxpayer turns to a GenAI mobile app that provides tax advice. On describing the scenario in detail, the app presents very specific reasons why the techniques described may not be valid. Satisfied with the answer and unwilling to pay for professional advice, the taxpayer decides to ignore the presentation, not wishing to risk being faced with hefty penalties.

Importers and Exporters. As part of a reform program, a customs administration introduces a virtual assistant to provide instant customer support. Configured as a chatbot, the assistant can help importers and exporters determine the correct classifications for goods. This capability is developed by training the assistant on the relevant classification codes and descriptions, along with the administration's prior rulings, both public and private, without disclosing the specific traders involved. Although tested extensively, the virtual assistant's reliability is known to be less than 100 percent, and, in some instances, it has provided traders with poor information.

Auditors. A small team of auditors is assigned to audit a taxpayer. As part of preparatory work, the team reviews the taxpayer's risk profile. These profiles are maintained in a computer system that uses expert rules and ML algorithms to highlight areas of interest and suggest adjustment values. The intent of the system is to help prioritize and focus examinations. The audit is performed and, in accordance with computer-generated

[16] See Aslett and others 2024.

[17] See MIT (2023). Also consider the use of RAGs, Retrieval Augmented Generation, for tailored AI solutions.

[18] An AI hallucination refers to a scenario where AI outputs based on patterns in data do not reflect actual reality.

guidance included in the taxpayer's risk profile, discrepancies in reporting are identified. The audit team, believing that it sufficiently understands the underlying issues, considers an adjustment amount suggested by the computer system. The adjustment seems large but reasonable and, based on positive experiences with the computer system, the amount is proposed to the taxpayer. In consultation with counsel, the taxpayer immediately agrees—leaving the auditors wondering if they should have taken a closer look.

Fraudsters. Recognizing the powerful capabilities of GenAI tools, a network of criminals learns how to "trick" a commercial AI chatbot into providing advice on how to best organize a value-added tax fraud program. Based on its access to deep repositories of knowledge and the content of the internet, the chatbot describes, in detail, specific jurisdictions, industries, business structures, and legal arrangements having low risk of detection. Beyond just planning, the criminal network learns how to use the technology to generate fictitious registrations and invoices. Inquiring with the chatbot as to what typical flags might trigger a warning or intervention by the authorities, the invoices are generated accordingly to avoid suspicion.

III. Accounting for Legal and Ethical Challenges

An expectation of AI's pervasive integration into daily life has provoked concerns over the negative consequences of its use. These concerns have both legal and ethical implications. As one example, in some jurisdictions, tax administrations may not lawfully process biometric data except under specific conditions.[19] Any AI system that does so would contravene the law. Even if the authority exists for a particular use of AI, aspects of specific applications may be unethical. These issues are complicated by the nature of AI itself, which often produces conclusions with a degree of perceived randomness that, even when documented, can be difficult, and in many instances impossible to easily explain.[20] Accordingly, while procedures for the use of AI may be defined, AI itself can provide seemingly unpredictable conclusions, having implications ranging from perceptions of bias to actual liability in the event of AI-influenced decisions leading to harm.

An awareness of key issues shaping discussions on legal and ethical perspectives on AI is essential to inform planning. This section introduces five principled concerns and a summary of recent regulatory initiatives that seek to address them.

Principled Concerns

Concern #1: Biases and Unjustified Use of Data

The term bias refers to subjective prejudices, whether of a human or an AI system or tool, that unjustifiably favor or disfavor individuals and groups. The mitigation of potential bias requires proactively proving the impartiality of decision making through objective examination of evidence. In doing so, every element influencing a decision must be scrutinized to ensure that it reflects deliberate intent and objective reality. In the context of AI, biases can be present in:

- **Training Data.** Training data is used to teach the algorithms relied on by AI systems and tools how to formulate and draw conclusions. This data may contain inherent biases, as a by-product of historical policies, social norms, institutional and operational practices.
- **Algorithmic Design.** Design choices and technical development processes can embed unintended biases in algorithms as a result of human errors or incorrect assumptions.
- **Outcomes.** AI systems may, even with unbiased data and meticulous design, generate biased outputs because of unexpected interactions between key data points (variables) or unforeseen contextual differences.
- **Feedback.** Biased feedback loops can reinforce existing biases over time, leading to future decisions that further disadvantage specific groups.

Concern #2: Simplexity

Simplexity refers to the trade-off between simple and complex aspects of systems where, particularly with AI, value and risk need to be carefully considered. As illustrated in Figure 2, overly complex AI may result in black boxes, limiting, in human terms, the ability to explain the conclusions that AI draws. However, oversimplification could lead to the use of inadequate or inappropriate data, introducing bias.[21] When adopting AI, a balance must be struck between use of simple, accessible AI systems and tools that

[19] For example, administrations subject to the European Union's General Data Protection Regulation.

[20] A reference to probabilistic variations, especially in generative AI, and algorithms that intentionally incorporate randomness or stochastic processes, all of which can make outputs appear unpredictable.

[21] Simpler AI technology and systems make use of rules-based algorithms and regressions, whereas more complex approaches tend to involve machine learning (including generative AI platforms, services, and tools).

add value while limiting underlying complexities. This requires consideration of the sophistication of users, organizational capacity, and data.

Figure 2. Simplexity: Balancing Simplicity and Complexity

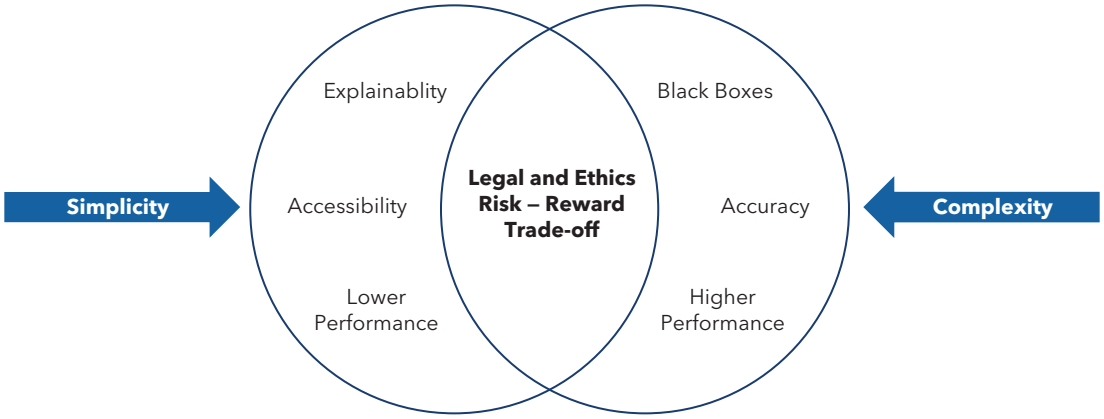

Source: Authors.

Concern #3: Transparency

Transparency for AI requires clear and concise information about AI systems and tools to be available in a way that is intelligible to all stakeholders. Transparency is a cornerstone principle to determine the legality and ethical use of AI. It can be explored through the following three components, each addressing a specific aspect of transparency. Answers to the questions may be useful to identify areas of improvement in a particular administration. More generally, Figure 3 presents a set of administrative actions for consideration to help regulate AI use.

- **Openness.** The supply of information: *What information on AI is openly shared?*
- **Intelligibility.** The receiver's ability to understand: *How is the information received?*
- **Explainability.** Information on AI use: *Why was an AI-influenced decision taken?*

Concern #4: Privacy and Data Governance

In the context of AI, privacy refers to taxpayers' right to retain some control over what, how, and when their personal information is collected and processed. Taxpayer data plays a critical role in the

Figure 3. Actions to Promote Transparency with AI

Action	Publicize the use of AI	Mandate notifications	Ensure explainability
Rationale	Before deploying an AI system, the public should be informed of the specific means of enforcement leveraged by the tax or customs administration(s).	Taxpayers and traders should be informed when they are interacting with AI (for example, a chatbot) or subject to AI-influenced decision making.	Taxpayers subject to redress decisions must be presented with an explanation that discloses the underlying logic of the relevant AI-influenced decisions.

Source: Authors.

development of accurate, high-performance, and human-centered AI systems, creating opportunities for innovation but also raising potential risks of data breach and intrusive surveillance. Figure 4 presents basic actions that promote data privacy protection.

Figure 4. Actions That Promote Data Privacy

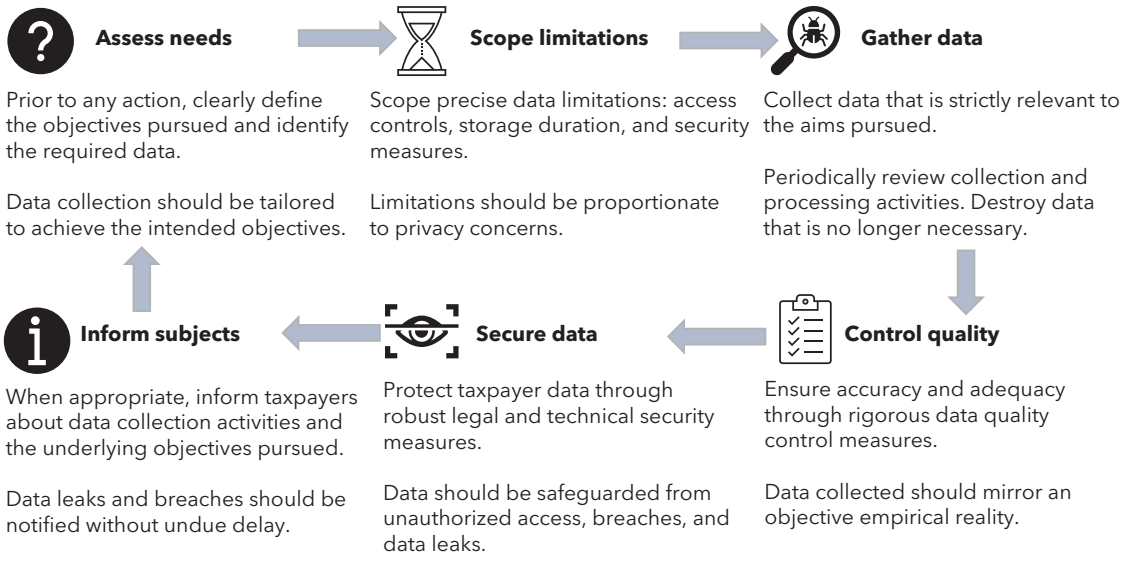

Source: Authors.

Concern #5: Accountability

Accountability requires the active enforcement of responsibility to mitigate unintended or adverse outcomes and guarantee the objectivity of decisions. Accountability plays a pivotal role in ensuring the ethical use of AI. The establishment of an accountability framework can help to identify, trace, and remedy any potential harm, fostering public trust and confidence in the use of AI. Figure 5 illustrates how such a framework should be used in practice.

Figure 5. Actions That Promote Accountability with AI

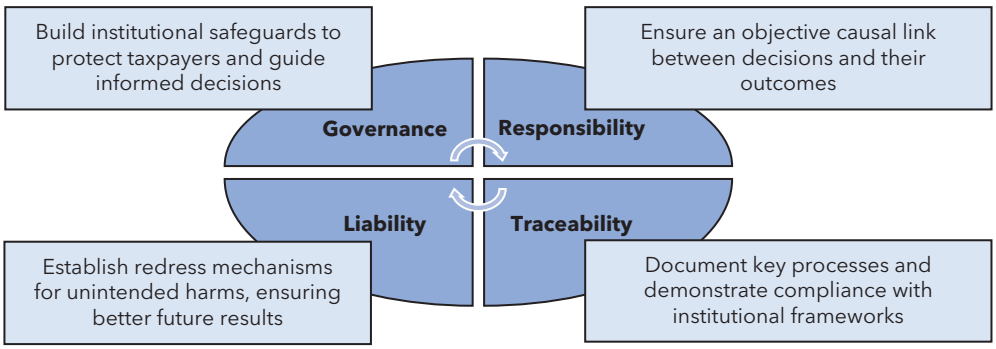

Source: Authors.

Recent Regulatory Initiatives

Numerous regulatory initiatives have emerged in recent years, at both international and national levels, to address AI concerns. A selection of the major initiatives, to date, is presented in Table 1. These

initiatives involve various stakeholders, including national legislative bodies and intergovernmental organizations. While the initiatives vary in their approaches, all seek to introduce more certainty with the use of AI, with some accounting for social implications. These initiatives are expected to be complemented by guidance published by industry and standards organizations, along with options for certification.[22]

Table 1. Recent AI Regulatory Initiatives

Jurisdiction/ Agency	Regulation	Scope	Approach
European Union (EU)	EU AI Act draft proposal (January 21, 2024)	AI systems, including GenAI	**Hard law regulation** Ex ante risk assessment with specific rules for selected use cases deemed "high risk"
Council of Europe	Draft Framework Convention on AI (December 18, 2023)	AI systems, including GenAI	**International Treaty** Common ethical principles and guidelines on enforcement mechanisms
Colombia	Proyecto de Ley Estatutaria 059 (November 14, 2023)	AI systems	**Hard law regulation** Ex ante risk assessment with specific rules for selected use cases deemed "high risk"
OECD	AI Principles (November 8, 2023)	AI systems, including GenAI	**Intergovernmental guidelines** Common ethical principles on the use of AI systems
United States	Executive Order on the Safe, Secure and Trustworthy Artificial Intelligence (October 30, 2023)	AI systems and "dual use" applications, including GenAI	**National guidelines** Decentralized guidelines based on common ethical principles
China	Interim measures on GenAI (August 15, 2023)	Public usage of GenAI applications	**Hard law regulation** Ex post liability based on ethical principles and rules
South Korea	Artificial Intelligence Responsibility and Regulation, Bill 2123709 (August 8, 2023)	AI systems, including GenAI	**Hard law regulation** Ex ante risk assessment with specific rules for selected use cases deemed "high risk"
Kyrgyzstan	Draft Digital Code of the Kyrgyz Republic (August 5, 2023)	Digital technologies, including AI systems (Chapter 23)	**Hard law regulation** Ex ante risk assessment with specific rules for selected use cases deemed "high risk"
Panama	Anteproyecto de Ley n°014 que regula la Inteligencia Artificial en la Republica (July 6, 2023)	AI systems, including GenAI	**Hard law regulation** Common ethical principles on the use of AI systems
Argentina	Marco legal para la regulación del esarrollo y uso de la Inteligencia Artificial (June 1, 2023)	AI systems, including GenAI	**Hard law regulation** Ex ante risk assessment with common ethical principles and guidelines
Mexico	Ley para la obótican ética de la inteligencia artificial y la obótica (May 23, 2023)	AI systems	**Hard law regulation** Common ethical principles on the use of AI systems
Chile	Boletin 15869-19 (April 24, 2023)	AI systems	**Hard law regulation** Ex ante risk assessment with specific rules for selected use cases deemed "high risk"
The Philippines	House Bill No. 7396 (March 1, 2023)	AI systems	**Hard law regulation** Common ethical principles on the use of AI systems
Israel	Policy on Artificial Intelligence, Ministry of Innovation, Science and Technology (October 30, 2022)	AI systems	**National guidelines** Decentralized "soft law" guidelines based on common ethical principles

Source: Authors; and Cantekin 2023.

Note: Legal norms generally conform to a hierarchy within which international treaties and hard law regulations (laws enacted by parliaments or legislative bodies) take precedence over executive orders and guidelines. AI = artificial intelligence; GenAI = generative AI; OECD = Organisation for Economic Co-operation and Development.

[22] For example, see ISO/IEC 42001 and its concept of an Artificial Intelligence Management System.

IV. An Overview and Selection of AI Use Cases

As generalized in Figure 1, digital operations and AI have a well-established history in tax and customs administration. The extent of AI's adoption today can be explored in the results of surveys and research. As one notable example, the Organisation for Economic Co-operation and Development administers a global survey in support of its Inventory of Tax Technology Initiatives, covering aspects of AI.[23] Although only surveying tax administrations, its findings provide insights likely to approximate a minimum level of AI use in customs administration. In the most recent survey, 55 percent of respondents (44 of 80) self-reported some use of AI, with higher levels of use reported in advanced economies. Separately, ongoing research from the University of Antwerp's Centre of Excellence DigiTax has identified the use of AI in 88 percent of EU member states' tax administrations (24 of 27), out of which most AI implementations are supporting enforcement functions.[24]

Specific operational applications of AI are best understood as scenarios generally referred to as "business use cases" or, more simply, "use cases." This section introduces the concept of a use case and provides examples for both tax and customs administration.

Understanding a Use Case

The expression "use case" originated in the field of software development as a mechanism to help organize and document requirements. Providing utility by encapsulating functionality for a particular scenario, use cases achieved popularity in their ability to break complex systems into pieces often more manageable (and verifiable) for both IT and business audiences. Today, the concept has achieved widespread adoption beyond software development—including in AI, where it is a foundational mechanism for managing the use of AI services, systems, and tools.

Tables 2 and 3 present a selection of illustrative tax and customs administration AI use cases. Categorized as either established or emerging, the use cases are intended to help inform administrations at early stages of planning to adopt or scale up the use of AI in an operational setting. The attributes in the tables have been included to help explore the dimensions of an AI use case in a general sense. They are based on evolving dialogue on AI with indicative values presented based on observations of actual AI implementations (for established use cases) or expectations of eventual use (for emerging use cases). For each use case presented in the tables, the respective legal and business requirements, AI technologies employed, and parameters for actual use will vary considerably based on specific operating contexts. As such, the information presented should be considered illustrative to build an awareness of AI rather than a literal assessment of any given use case's relative qualities.

[23] See OECD 2023.

[24] See Hadwick 2024.

Selected Use Cases: Tax Administration

Table 2. Illustrative AI Use Cases: Tax Administration

	AI Use Case	Indicative AI Technology	Potential Value[1]	Indicative Maturity	Indicative AI Explainability	HITL / HOTL Suggested?[2]	Indicative Risk Classification[3]
1	**ESTABLISHED USE CASES**						
1.1	Risk Management: Case Selection–Audit	ERs, ML	High	High	Medium	✓	High
1.2	Risk Management: Case Selection–Collections	ERs, ML	High	High	Medium	✓	High
1.3	Risk Management: Refunds Risk Assessment	ERs, ML	High	High	Medium	✓	Medium
1.4	Risk Management: Profiling–Web Scraping	NLP, ML	Low	Low	Low	✓	High
1.5	Risk Management: Profiling–Social Media Ingest	NLP, ML	Low	Low	Low		High
1.6	Risk Management: Profiling–Taxpayer Risk Profiling	ERs, ML	Medium	Medium	Medium		Low
1.7	Risk Management: Exploratory Analysis & Data Mining	NLP, SNA, ML	High	High	Medium		Low
1.8	Investigations: Social Network Analysis	NLP, SNA, ML	Low	Medium	Medium		Low
1.9	Investigations: Unstructured Content Analysis	NLP, SNA, ML	Low	Low	Medium		Low
1.10	Enforcement: Predicting Default Assessment Values	ERs, ML	Medium	Medium	Medium	✓	Medium
1.11	Taxpayer Services: Interactive Virtual Assistance	ML, GenAI	Medium	Medium	Low		Low
1.12	Taxpayer Services: Voice Recognition	NLP, DL	Low	Medium	High		Low
1.13	Taxpayer Services: E-Filing Suggested Values	ERs, ML	Low	Medium	High		Low
1.14	Taxpayer Services: Nudging Programs	ML, GenAI	Low	Low	Medium		Low
2	**EMERGING USE CASES**						
2.1	All Staff: GenAI Support for Desktop Productivity	GenAI, NLP	Low	Low	Low	✓	Medium
2.2	Taxpayer Services: GenAI Content Creation	GenAI, NLP	Low	Low	Low	✓	Low
2.3	Risk Management: GenAI Support for Risk Analysis	GenAI, NLP	Medium	Low	Medium	✓	Medium
2.4	Training: GenAI Support for Staff Education Tools	GenAI, NLP	Medium	Low	Medium		Low

Source: Authors.

Note: DL = deep learning; ERs = expert rules; GenAI = generative AI; ML = machine learning; NLP = natural language processing; SNA = social network analysis

[1]Potential Value is an estimate of a use case's indicative ability to positively influence compliance behavior or operational effectiveness.

[2]HITL is "Human-in-the-Loop," HOTL is "Human-on-the-Loop." The concepts describe the degree of human involvement in an AI-supported process. See Section V for details.

[3]Indicative Risk Classification is an estimate of the relative potential for a use case to create ethical, reputational, or other harm when compared with other use cases.

Selected Use Cases: Customs Administration

Table 3. Illustrative AI Use Cases: Customs Administration

	AI Use Case	Indicative AI Technology	Potential Value[1]	Indicative Maturity	Indicative AI Explainability	HITL / HOTL Suggested?[2]	Indicative Risk Classification[3]
1	**ESTABLISHED USE CASES**						
1.1	Risk Management: Pre-Arrival Cargo Targeting	ERs, ML	High	Medium	Medium	✓	High
1.2	Risk Management: Selectivity for Channeling	ERs, ML	High	High	Medium	✓	High
1.3	Risk Management: Post-Clearance Case Selection	ERs, ML	High	High	Medium	✓	High
1.4	Risk Management: Automated Cargo Inspection	MV, ML	Medium	Medium	Medium	✓	Low
1.5	Risk Management: Profiling–Trader Risk Profiling	ERs, ML	Medium	Medium	Medium		Low
1.6	Risk Management: Cross-Border Data Matching	ERs, ML, NLP	High	Medium	Medium	✓	Low
1.7	Risk Management: Exploratory Analysis & Data Mining	NLP, SNA, ML	High	High	Medium		Low
1.8	Border Management: Facial Recognition	MV, ML	Medium	High	High	✓	High
1.9	Border Management: Passenger Risk Profiling	ERs, ML, MV	Medium	High	High		High
1.10	Border Management: Vehicle Tag Recognition	MV, ML	Medium	High	High	✓	Medium
1.11	Investigations: Social Network Analysis	NLP, SNA, ML	Low	Medium	Medium		Low
1.12	Investigations: Unstructured Content Analysis	NLP, SNA, ML	Low	Low	Medium		Low
1.13	Customs Services: Interactive Virtual Assistance	ML, GenAI	Medium	Low	Low		Low
1.14	Customs Services: Automated Document Processing	ML, NLP	Medium	Medium	Medium		Low
2	**EMERGING USE CASES**						
2.1	All Staff: GenAI Support for Desktop Productivity	GenAI, NLP	Low	Low	Low	✓	Medium
2.2	Customs Education Services: GenAI Content Creation	GenAI, NLP	Low	Low	Low	✓	Low
2.3	Risk Management: GenAI Support for Risk Analysis	GenAI, NLP	Medium	Low	Medium	✓	Medium
2.4	Training: GenAI Support for Staff Education Tools	GenAI, NLP	Medium	Low	Medium		Low

Source: Authors.

Note: ERs = expert rules; GenAI = generative AI; ML = machine learning; MV = machine visioning; NLP = natural language processing; SNA = social network analysis.

[1]Potential Value is an estimate of a use case's typical ability to positively influence compliance behavior or operational effectiveness.

[2]HITL is "Human-in-the-Loop," HOTL is "Human-on-the-Loop." The concepts describe the degree of human involvement in an AI-supported process. See Section V for details.

[3]Indicative Risk Classification is an estimate of the relative potential for a use case to create ethical, reputational, or other harm when compared with other use cases.

V. How to Promote Responsible Use of AI

While regulatory concepts for AI continue to evolve, senior leaders can take action to begin implementing prudent management practices. Intended to help promote certainty across a portfolio of AI use cases, the following 10 actions should be understood as suggestions based on observations of tax and customs administrations in different regions of the world. They are expected to be applicable to most administrations, regardless of their scale and level of technical sophistication. Like other technology-oriented topics, such as data management and information security, the governance and management of AI is an institutional responsibility that extends far beyond the remit of an IT department. Accordingly, the first action that should be considered is introducing an administrative policy on AI's adoption and use.

1. Implement a Formal Policy on AI

Develop and implement a policy document that opens by explaining what AI is and why regulation is needed. Having established context for the policy, prioritize statements regarding (1) senior leadership's expectations and intent, (2) ethical principles and values governing the use of AI, (3) compliance with external regulations and standards, (4) requirements for transparency and explainability, (5) use of external AI services, (6) governance and oversight mechanisms, (7) procedures for managing AI risks, (8) training and awareness, and (9) monitoring and evaluation. For reference, Annex 1 provides an example of a formal AI policy for tailoring and adoption.

2. Sensitize Staff to AI and Individual Duty of Care

Regardless of whether a formal AI policy is in place, organize a deliberate project designed to sensitize staff to AI and their responsibilities as public servants. Staff, particularly senior leadership, must be encouraged to proactively understand the benefits, limitations, and potential consequences of the AI systems and tools that support regular work. They should understand that their responsibilities may require that they can explain how the relevant AI makes decisions, what potential biases could exist in the AI, the implications of relying on AI output, and where human involvement is essential. With this knowledge, they should carry out their duties emphasizing the rights of taxpayers and the ethical use of AI tools in compliance with laws and regulations.

3. Build from Strong Fundamentals with an AI Strategy

Invest the time needed to prepare a strategy for AI that is appropriate and consistent with broader reform and modernization objectives. Aspects of AI are likely to be transformational, presenting both opportunities and risks. However, the core fundamentals of tax and customs administration will persist well into the future. Because of this dynamic, a dedicated strategy for AI should be created in many administrations. But although this strategy may speak to advanced aspects of IT, work on AI should not be at the expense of strengthening core operations in areas that are weak. Annex 2 provides a concept to help navigate these issues.

4. Establish a Formal Inventory of AI Use Cases

Compile a comprehensive inventory of AI that may currently be used or planned, referencing the concept presented in Tables 2 and 3. The use of AI may be obvious and explicit (for example, as part of a narrow, standalone risk management system) or embedded in a toolkit relied on for many different purposes (for example, as part of a suite of advanced analytics tools used for exploratory analysis and data

mining). In either scenario, all deliberate and other nontrivial uses of AI should be documented as use cases and cataloged with attributes. In doing so, care should be taken to account for future plans and to record the categories of AI technology relied on (for example, expert rules, ML, natural language processing, and GenAI).

5. Subject Use Cases to Legal and Ethical Review

Scrutinize each use case in the inventory to ensure compliance with legal requirements and sound ethical practices. Annex 3 provides a methodology to support a risk assessment of each use case. The methodology accounts for each of the principled concerns described in Section III (biases and unjustified use of data, simplexity, transparency, privacy and data governance, and accountability). Once complete, and using the concept provided in Tables 2 and 3, specify an overall, relative risk classification for each use case (high, medium, low). As part of any enterprise risk-management processes, ensure that high-risk use cases in the inventory are regularly elevated for review and consideration of mitigation strategies.

6. Where Appropriate, Keep a "Human-in-the-Loop"

Mandate a requirement for human involvement in scenarios where AI-influenced decisions have the potential to cause significant harm. The concepts of "Human-in-the-Loop" (HITL) and "Human-on-the-Loop" (HOTL) have emerged as important mechanisms in AI regulation. Their premise is to mitigate risks while combining the benefits of AI analytical capabilities with human judgment and, in doing so, leveraging the advantages of both to improve overall effectiveness. In HITL, humans may be directly involved in all final decisions. With HOTL, humans play more of an oversight role, fine-tuning AI and stepping in only on higher-level, critical decisions. In tax and customs administration, HITL and HOTL concepts are important in areas where a temptation may exist to fully "automate and forget" processes (for example, selectivity in customs).

7. Risk-Assess New Use Cases Prior to Introduction

Apply the same risk review methodology to all new, future AI use cases that are explored for introduction. Although AI may facilitate new automation and levels of effectiveness, its use comes with costs beyond acquisition and implementation of the respective technology. Before deciding to implement a use case, the full consequences should be evaluated. These include costs relating to acquisition, ongoing operations, and mitigation if any AI used leads to unintended or unforeseen consequences, some of which might impact future compliance and revenues (for example, resulting in audit results overturned on appeal).

8. Regularly Publish the Use Case Inventory

Publish a filtered copy of the AI use case inventory to advance transparency and promote certainty. Exclude from the inventory AI use cases that (1) reasonably relate to the protection of national security interests (particularly in the realm of customs administration border security) or (2) have little potential impact on core business operations or the affairs of taxpayers and traders (for example, excluding trivial, often embedded AI tools used by staff such as "spellcheck" in desktop productivity software). Once filtered, the inventory should be published and a channel provided for soliciting feedback.

9. Prominently Disclose the Use of AI in Regular Operations

Introduce the changes required so that all interactions with AI are made clear and obvious, regardless of the audience, internal or external. When employing AI systems or tools, users should be notified that AI is being used. For example, when interacting with a virtual assistant, the assistant should inform the user that it is a bot and not a human. Similarly, all outputs that include AI-generated content—emails, drafts

or summaries of documents, data, or images—should be clearly labeled. This may be achieved, for example, by affixing a statement such as "elements of this content were generated by the use of AI systems or tools."

10. Evaluate Use Cases for Performance and Intent

Having established control over an inventory of AI use cases, regularly reevaluate the actual impact of each. Beyond questions relating to the value of new digital investments, AI is unique in that while the intent behind a use case (its business purpose and justification) may be clearly defined, AI's actual use can result in unanticipated and undesirable outcomes. Given these consequences, in practice and beyond metrics, AI may not deliver the basic functionality originally intended. Once an AI use case moves into an operational state, it is crucial that the actual decisions made or influenced by the AI be regularly evaluated.

VI. How to Introduce an AI Use Case

The introduction of an AI use case is best facilitated through concepts that already exist for managing digital services. Figure 6 presents a structured approach based on an IT service's life cycle. It relies on the use of "control gates" through which any significant change must be evaluated and approved (through governance structures, often in the form of committees). The concept illustrated integrates change control across business (project management), technical development (systems and services), and IT operations, using mechanisms from industry standard methodologies. It has been tailored for AI and is suitable primarily for mid- to large-sized administrations.

Managing Change across the Life Cycle of an AI Use Case

Figure 6. Integrating AI into a Portfolio of IT Services

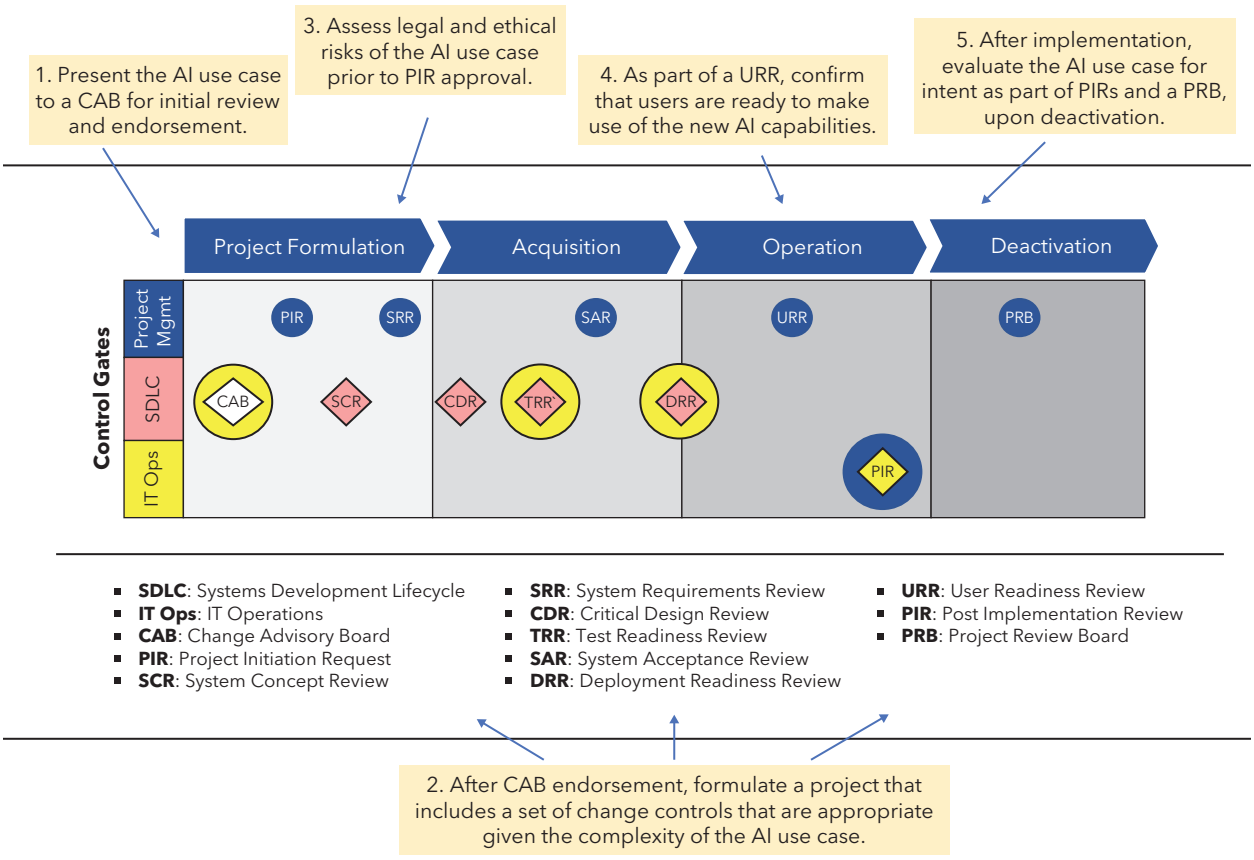

Source: Aslett 2024. The diagram has been tailored for AI from material included in the Virtual Training to Advance Revenue Administration (VITARA) Information Technology and Data Management online course.

Note: In practice, this concept is applied to any change that qualifies as a "project" by meeting or exceeding defined criteria.

VII. Questions Worth Asking

As the field of AI is undergoing a period of rapid change, discussions on many aspects of its use remain unsettled. Among the many open questions being debated, the following may be especially relevant to senior officials involved with tax and customs administrations.

Who should have primary accountability for AI at an institutional level?

A director for modernization or innovation may be an appropriate figurehead for leading and providing oversight of AI initiatives. Because AI draws from multiple domains—business, technology, data—a leadership role having authority for modernization may be ideal. Other valid options could include a Chief Information Officer, a Chief Data Officer, or a cross-functional working group, depending on the scale and complexity of a particular administration.

Where do risks and liability exist in AI service supply chains?

In some scenarios, AI services are provisioned through layers of external providers, and it is not clear what upstream suppliers are being relied on or what risks exist. If, for example, biased data or content is supplied by one firm to another firm that trains an LLM used by GenAI services, and then that service influences an unjustified enforcement decision—what are the repercussions? Taking the time to identify and think through these scenarios may help inform sound contracting and service arrangements that allow safer exploration and introduction of AI.

How much or what level of AI explainability should be targeted?

The ideal amount of explainability expected from AI is likely to be proportionate to an administration's tolerance for risk and resources for mitigation. In instances where there is little capacity to evaluate and manage AI, less AI should probably be used—and, when used, care should be taken to avoid high-risk scenarios where significant revenue may be at stake.

To what extent should an AI strategy account for external AI adoption and use?

While prioritizing internal institutional matters, some degree of planning around the use of AI by external stakeholders is prudent. Consistent with principles that promote responsible use of AI, external stakeholders should also be encouraged (or, when possible, directed) to prominently disclose the use of AI in official interactions with tax and customs officers.

How advanced is AI likely to become in the next 5, 10, or 15 years?

In short, very advanced—and to the extent that individuals and agencies not using AI as part of regular work will be at a significant disadvantage. This reality will provoke disruption that, if well-managed, will allow staff to refocus their time on more complex, value-adding work activities than would otherwise be possible.

How should an approach to AI vary depending on the level(s) of institutional maturity?

Although established thinking around digital development remains valid, the arrival of GenAI warrants special attention. Historically, "leapfrogging" through new digital investments has been viewed as a valid strategic choice, with the caveat that care should be taken to ensure that core fundamentals are put in place. Because GenAI will filter into tax and customs administrations through external channels, a policy for AI use is important regardless of overall digital maturity.

OVERVIEW OF ANNEXES

The annexes to this note provide tools that are designed to be tailored for use. They are not intended to be copied and pasted into an operational setting. Tailoring and customizing the details is an essential requirement. To support this process, in many administrations, a cross-functional team with representation from across key AI domains—business, technology, and data—should be formed, possibly as part of an AI committee. From within this group, each tool should be understood as follows.

Example of an AI Policy

Annex 1 provides the text of an illustrative AI policy document. Streamlining AI policy into other documents could be sensible in some instances (for example, as part of an IT policy). Regardless of approach, AI policy should be endorsed by leadership and formally disseminated, supported by outreach to sensitize staff. The example provided is based on a review of public references and internal IMF consultations. Its development and practical adoption could benefit from the input of external agencies. In advance of adoption of an AI policy, some administrations may find it appropriate to issue short notes or guidance on key topics, which could include, as one example, the appropriate use of GenAI chatbots.

References for Developing an AI Strategy

Annex 2 provides two diagrams to help explore and think through AI strategy. Although AI offers new capabilities, the development of a formal strategy remains a business task consistent with well-established good practices. Where AI requires special care is in understanding the new capabilities it offers, its potential to advance transformation, and the layers of dependencies that it relies on. The first diagram in Annex 2 provides a structured concept for thinking through aspects of a strategy, in general—with commentary on issues specific to AI. The second diagram provides an example of a road map that addresses AI dependencies, use cases, and sequencing.

AI Risk Assessment Methodology

Annex 3 provides a concept for risk-assessing AI use cases. Building from a range of literature on AI and the concepts in Section III, the methodology is applied in practice by evaluating an AI use case against 40 specific questions.[25] The questions are organized into three areas covering AI (1) input, (2) throughput, and (3) output. Within these, two types of questions exist—simple yes/no (binary) and open-ended. The most important questions are highlighted and, in risk assessing a use case, should be prioritized. As the contextual applications of AI can vary dramatically, the methodology does not provide logic to output an overall risk classification label. Instead, results of the assessment are intended for subjective interpretation on a case-by-case basis.

[25] Including Cantekin (2023), and Gesley and others (2023).

ANNEX 1. Example of an AI Policy

Artificial Intelligence Use Policy

I. Introduction

The < Tax or Customs Administration > has determined to make strategic use of artificial intelligence (AI) to advance its mandate to improve voluntary compliance < and facilitate trade >. In doing so, the < Tax or Customs Administration > is committed to AI's responsible use, ensuring at all times full compliance with legal requirements and generally accepted ethical norms. The decision to use AI is reflected in the < Tax or Customs Administration > strategic plan and digital transformation strategy.

In support, this policy outlines the expectations to be adhered to by all internal staff, contractors, service providers, suppliers, and other third parties that carry out work on behalf of the < Tax or Customs Administration >.

II. Definitions

Artificial Intelligence. The simulation of human intelligence using computer systems to make predictions, inform decisions, and generate content, often relying on machine learning algorithms, expert rules, machine visioning, and other related technologies.

Artificial Intelligence Working Group. The pinnacle body in the < Tax or Customs Administration > tasked with oversight and compliance with this policy, along with the evaluation of current and future AI use cases. Referred to hereinafter as the "AIWG."

Artificial Intelligence System. Any dedicated computer system (or service) that makes use of AI technology to generate outputs, make predictions, recommendations, or decisions that influence staff, taxpayers, service providers, or other third parties that carry out work on behalf of < Tax or Customs Administration >, excluding "Embedded AI Tools," as discussed in the following section.

Artificial Intelligence Use Case. A specific business scenario within which the use of AI is employed to carry out defined business objectives.

Embedded AI Tools. Tools that make use of AI technology that have already been approved for use within defined constraints and do not require additional review by the AIWG. An example of an embedded AI tool would include the "spellcheck" function in popular desktop productivity software.

Personal Information. Information that may identify, relate to, or is capable of being associated with a physical person, whether directly or indirectly.

Taxpayer Information. Information that may identify, relate to, or is capable of being associated with a physical person or legal entity, whether directly or indirectly, that may be or become a registered taxpayer with the < Tax or Customs Administration >.

III. Governance and Oversight

An AIWG will be established and operate as the pinnacle governing body for AI in the < Tax or Customs Administration >. The AIWG will convene quarterly to provide direction and oversight of all AI-related topics. This includes implementation of AI strategy, and review and approval of current and future AI use

cases. The AIWG will be chaired by and accountable to the < Tax or Customs Administration >'s Director of Modernization, who will approve the AIWG's terms of reference.

IV. General Standards for AI Use

With the limited exception of Embedded AI Tools available in approved software, and within any additionally defined constraints, all uses of AI must be approved by the AIWG to ensure that it is:

- **Lawful.** All uses of AI must comply with applicable laws and regulations, which may be present in contractual obligations in the form of limitations or other restrictions.
- **Ethical.** All uses of AI must comply to the extent practical with the ethical guidelines defined subsequently in this policy.
- **Justifiable.** All operational uses of AI must be for a valid and justified business purpose. This includes experimentation in a development context, which, if it occurs, must have clearly defined objectives. Open-ended, uncontrolled experimentation is not permitted.

V. Specific Requirements for AI Use

- **Preapproval Before Introduction.** All uses of AI must be evaluated and approved in advance by the AIWG. Exceptions will be reviewed on a case-by-case basis. Approvals authorizing use are recognized once an AI use case is registered in the < Tax or Customs Administration > use case inventory (in either development or production status) along with a formal risk rating.
- **Adherence to Approved AI Use Case Controls.** Restrictions, procedures, and limitations must be documented as part of an approved AI use case and complied with at all times. The primary responsibility for ensuring adherence lies with the ranking official of the office designated as the business owner for the use case (hereinafter referred to as the "Owner").
- **Adherence to Approved Business Intent.** An AI use case, once approved, may not deviate in operation significantly from its originally scoped purpose. Any substantive change must be raised for consideration by the AIWG prior to introduction.
- **Adherence to External Requirements.** All use cases must ensure ongoing compliance with external legislation, regulations, and policies. These may include, but are not limited to, data privacy and the protection of personal and taxpayer information. The primary responsibility for ensuring compliance rests with the Owner, to be supported by guidance from the AIWG.
- **Adherence to Domain Expertise Requirements.** The approval, implementation, and operation of an AI use case must be accompanied by the requisite domain expertise and human resources (in the form of staff having the subject matter knowledge to administer and use the AI, ensuring appropriate use of Human-in-the-Loop or Human-on-the-Loop concepts).
- **Adhere to Staff Sensitization and Training Requirements.** All staff must be sensitized to AI and this policy. The staff that will interact with particular AI systems, services, or tools must receive additional appropriate training on that AI in advance of use (confirmed through a formal user readiness review, where applicable).
- **Adherence to Information Security Policy.** Prior to approval, all use cases must be formally evaluated for compliance with the < Tax or Customs Administration >'s Information Security Policy, using a procedure specified by the Office of the Chief Information Security Officer. Once approved, the Owner must ensure continued and ongoing compliance.
- **Disclosure of AI-Generated Content.** Any content produced by AI (including generative AI tools) must be labeled accordingly (for example, by affixing the statement "elements of this content were generated or influenced by the use of AI systems or tools").

- **Notification of AI Interactions.** Unless granted a waiver by the AIWG, all uses of AI must be prominently disclosed to users. This disclosure may take different forms but should make clear and obvious the presence and use of AI (for example, notifying a user that they are interacting with a chatbot rather than a human).

VI. Ethical Guidelines

The < Tax or Customs Administration > recognizes that in some instances particular uses of AI may be legal but unethical. Accordingly, the < Tax or Customs Administration > has determined that its use of AI will be, at all times, consistent with generally accepted ethical norms by deliberately evaluating and risk assessing AI use cases in the AIWG's approval process for:

- **Potential Biases.** Use cases must be evaluated to ensure the impartiality of decision making and that AI decisions reflect deliberate intent and objective reality. To do so, the evaluation must, to the extent practical, consider (1) training data, (2) algorithmic design, (3) outcomes, and (4) feedback.
- **"Simplexity."** The complexity of AI adopted or developed for a use case must be deliberately limited to the extent practical (limiting "black boxes") while ensuring that the AI provides the minimum capabilities required to satisfy the intended business objectives.
- **Transparency.** Use cases must be evaluated for (1) openness, (2) intelligibility, and (3) explainability. In instances where a lack of transparency may exist, mitigating measures must be developed, approved, and implemented prior to the use case entering an operational state.
- **Privacy and Data Governance.** Use cases must be evaluated to confirm that an appropriate set of controls are (or will be) in place to guarantee the privacy of personal and taxpayer information, in a manner consistent with existing data governance policies.
- **Accountability.** Use cases must be evaluated to determine whether adequate mechanisms are in place to ensure accountability of AI outcomes. To do so, the evaluation must consider (1) governance arrangements (safeguards), (2) liability (redress mechanisms), (3) responsibility, and (4) traceability (documented processes and institutional frameworks).

VII. Prohibited Uses of AI

Certain uses of AI are prohibited in the absence of explicit authorization from the AIWG, including:

- **Supplying Internal Data or Information to External AI Services.** You may not provide *any* of the < Tax or Customs Administration >'s data or information to AI systems, services, or tools that are operated outside of the < Tax or Customs Administration >. This prohibition includes "chatbots" accessible over public networks.
- **Using Embedded AI Tools to Circumvent Controls.** You may not use embedded AI tools in a manner that circumvents controls that should normally apply to an AI use case that has been formally reviewed and approved by the AIWG. When in doubt, it is your responsibility to contact the AIWG for guidance.

VIII. Special Requirements for High-Risk AI

Any approved AI use case designated as "high-risk" is subject to the following:

- **Requirement for High-Quality Data.** Using a standard to be defined by the AIWG, the data sets used for training high-risk AI systems, services, or tools must be objectively designated as high quality prior to use in any operational environment.
- **Requirement for Consistent Performance.** Because of the risk of harm, a high-risk AI use case, once developed, must be thoroughly tested and function within well-defined parameters agreed by the AIWG before introduction into an operational setting.

- **Requirement for Regular Evaluation.** With the well-defined parameters describing a "regular" state of operation, the AI use case must be evaluated on at least a < time period > basis for performance and alignment with approved business objectives.

IX. Reporting Noncompliance

The < Tax or Customs Administration >'s internal staff, contractors, service providers, suppliers, and other covered third parties aware of any conduct that may violate this policy must report it. Reports may be filed by internal staff through regular management channels or, as appropriate, directly to a representative of the AIWG. All reports will be promptly reviewed and, if warranted, investigated.

ANNEX 2. References for Developing AI Strategy

Annex Figure 2.1. Exploring Aspects of AI Strategy

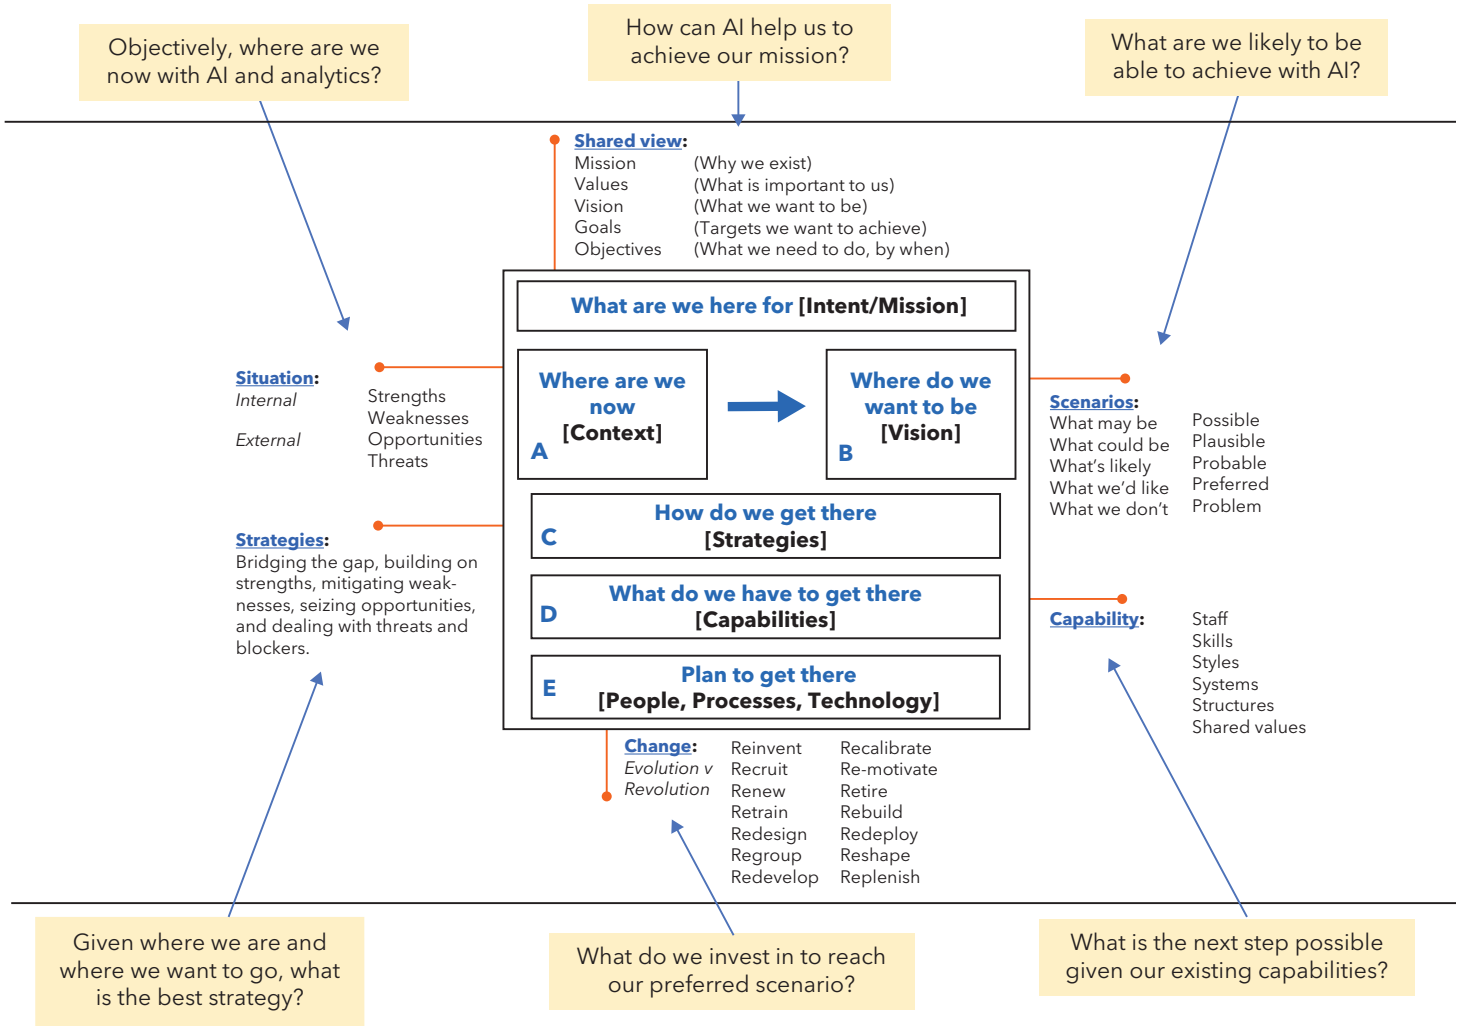

Source: Authors.

Annex Figure 2.2. Illustrative AI Road Map: Dependencies, Use Cases, and Sequencing

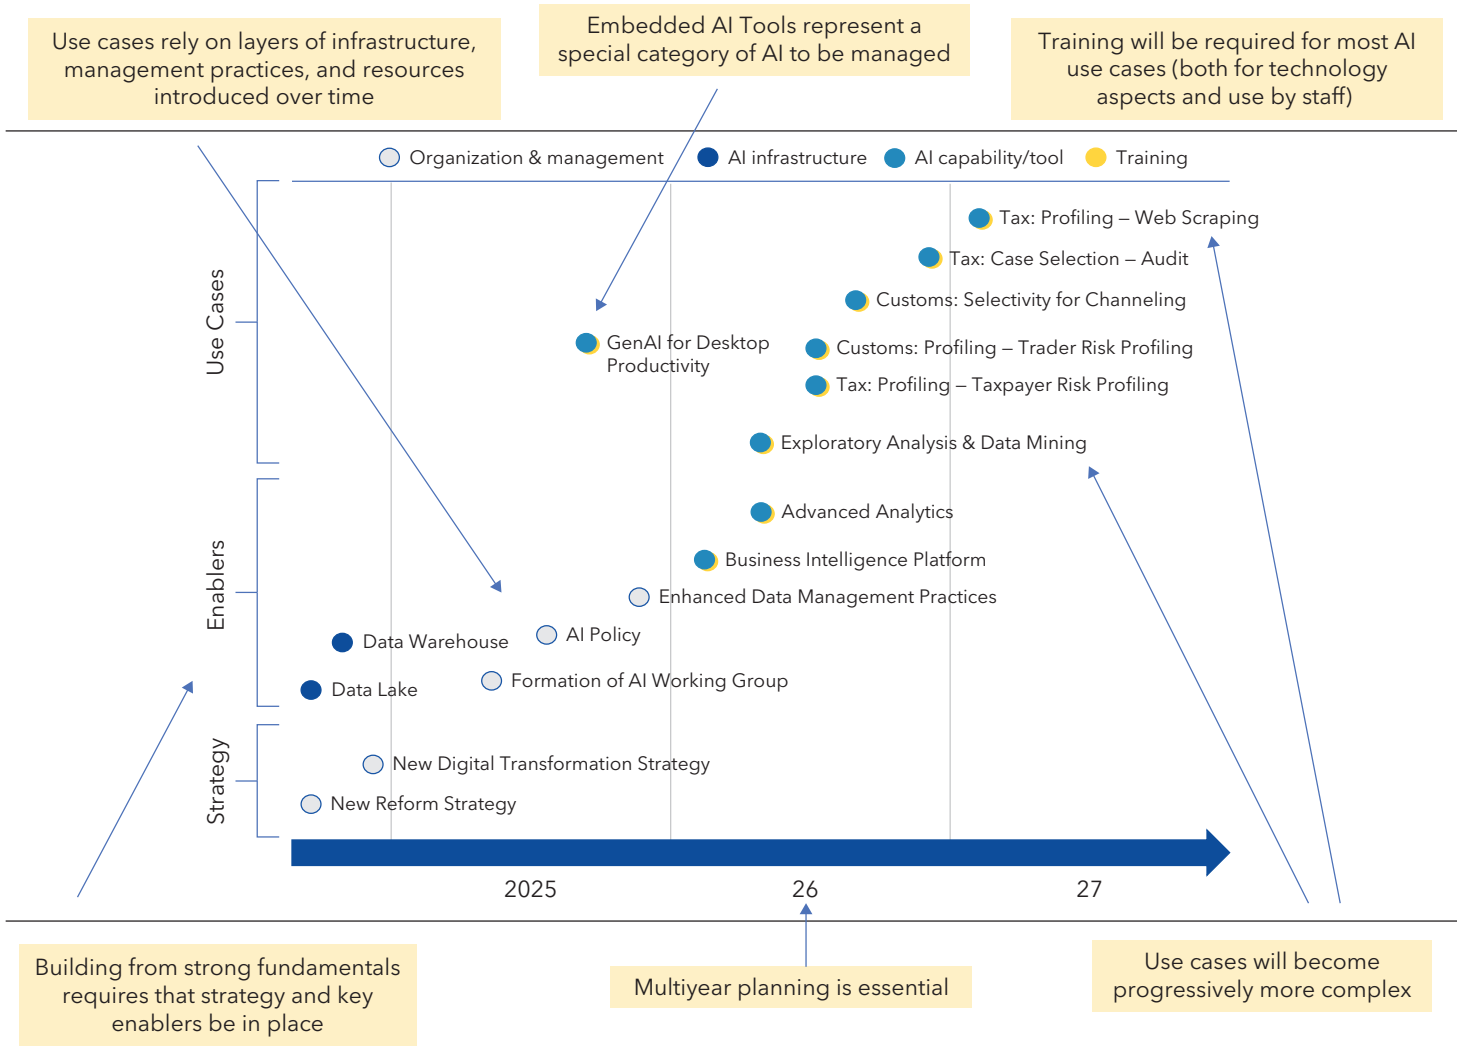

Source: Authors.

ANNEX 3. AI Risk Assessment Methodology

Annex Figure 3.1. AI Risk Assessment Methodology

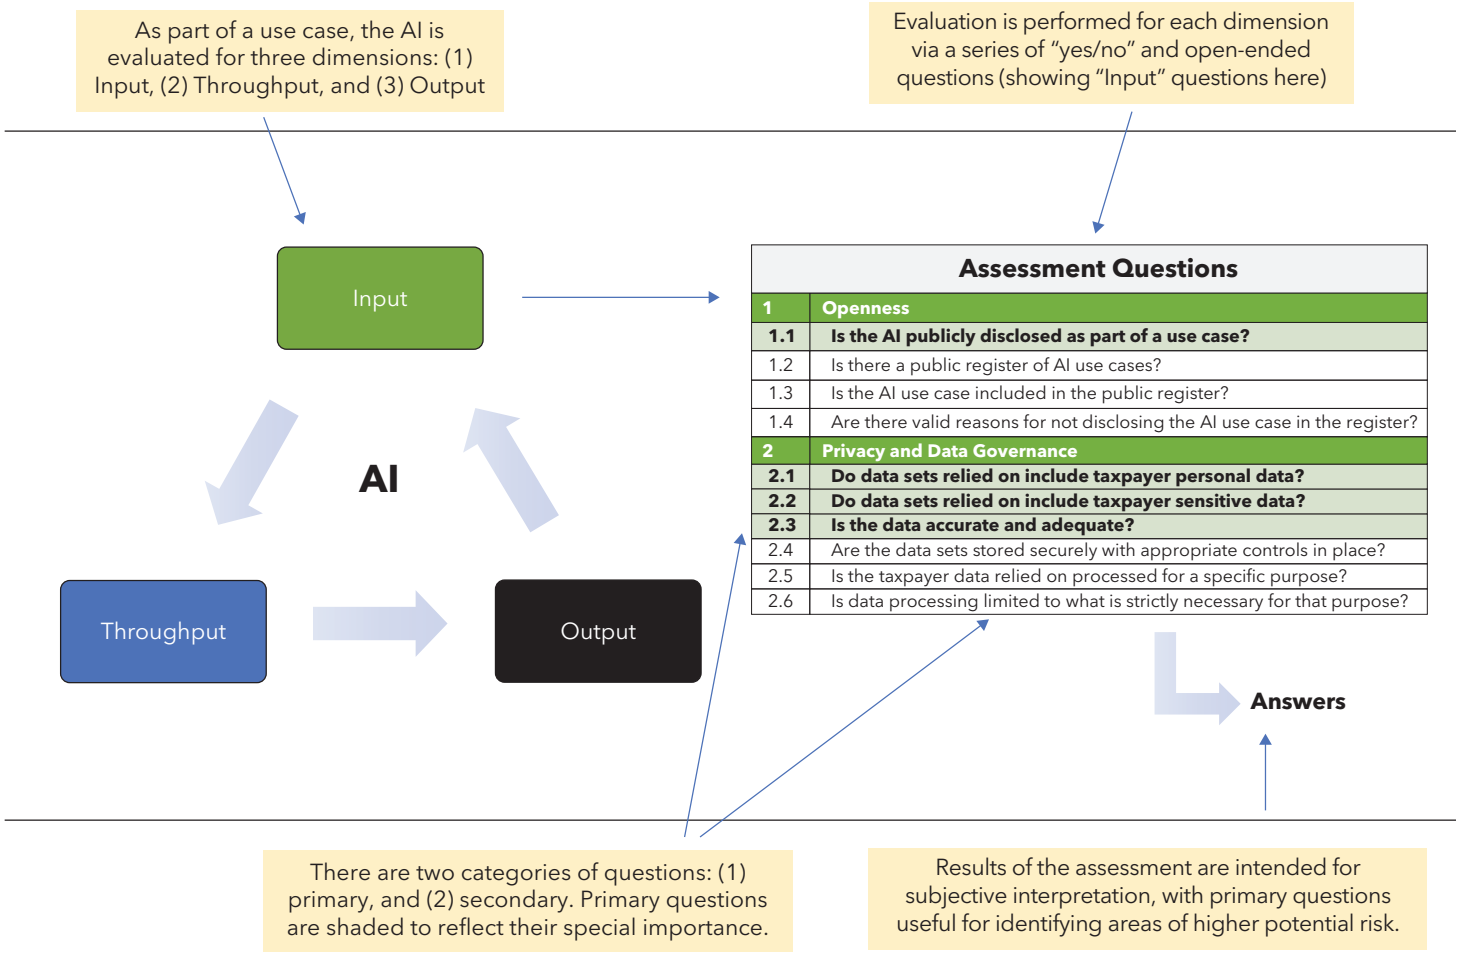

Source: Authors.

Annex Table 3.1. Input Dimension: Risk Assessment Questions

	Assessment Questions	Type of Question	Response Indicating More Risk	Response Indicating Less Risk
1	**OPENNESS**			
1.1	**Is the AI publicly disclosed as part of a use case?**	Binary (Yes / No)	**N**	**Y**
1.2	Is there a public register of AI use cases?	Binary (Yes / No)	N	Y
1.3	Is the AI use case included in the public register?	Binary (Yes / No)	N	Y
1.4	Are there valid reasons for not disclosing the AI use case in the register?	Binary (Yes / No)	N	Y
2	**PRIVACY AND DATA GOVERNANCE**			
2.1	**Do data sets relied on include taxpayer personal data?**[1]	Binary (Yes / No)	**Y**	**N**
2.2	**Do data sets relied on include taxpayer sensitive data?**	Binary (Yes / No)	**Y**	**N**
2.3	**Is the data accurate and adequate?**	Binary (Yes / No)	**N**	**Y**
2.4	Are the data sets stored securely with appropriate controls in place?	Binary (Yes / No)	N	Y
2.5	Is the taxpayer data relied on processed for a specific purpose?	Binary (Yes / No)	N	Y
2.6	Is data processing limited to what is strictly necessary for that purpose?	Binary (Yes / No)	N	Y
2.7	What mechanisms uphold data accuracy and adequacy?	Qualitative Evaluation	Subjective	Subjective
2.8	What measures safeguard data security?	Qualitative Evaluation	Subjective	Subjective

Source: Authors.

[1]Note: Data in the context of "Input" refers to both training data and user input.

Note: The shaded rows indicate the most important questions, which should be prioritized.

Annex Table 3.2. Throughput Dimension: Risk Assessment Questions

	Assessment Questions	Type of Question	Response Indicating More Risk	Response Indicating Less Risk
3	**ALGORITHMIC DESIGN**			
3.1	**Is there an objective justification for the use of each variable?**	Binary (Yes / No)	**N**	**Y**
3.2	**Is the AI's output accurate?**	Binary (Yes / No)	**N**	**Y**
3.3	Is the output coming from a black box or unexplainable?	Binary (Yes / No)	Y	N
3.4	Is the output intelligible to officials internally?	Binary (Yes / No)	N	Y
3.5	What is the degree of complexity in the AI?	Qualitative Evaluation	Subjective	Subjective
3.6	What measures mitigate the risks of biases?	Qualitative Evaluation	Subjective	Subjective
3.7	What measures safeguard the accuracy of the AI's output?	Qualitative Evaluation	Subjective	Subjective
4	**ACCOUNTABILITY**			
4.1	**Is there an objective link between each decision and outcomes?**	Binary (Yes / No)	**N**	**Y**
4.2	**Are there any mechanisms to redress unintended outcomes?**	Binary (Yes / No)	**N**	**Y**
4.3	Are all the key processes well-documented?	Binary (Yes / No)	N	Y
4.4	Are relevant staff sensitized to the AI, its security, and risks?	Binary (Yes / No)	N	Y
4.5	What measures guarantee the objectivity of decisions?	Qualitative Evaluation	Subjective	Subjective
4.6	What measures guarantee accountable outcomes?	Qualitative Evaluation	Subjective	Subjective

Source: Authors.

Annex Table 3.3. Output Dimension: Risk Assessment Questions

	Assessment Questions	Type of Question	Response Indicating More Risk	Response Indicating Less Risk
5	**FAIRNESS**			
5.1	**Is there a Human-in-the-Loop (HITL)?**	Binary (Yes / No)	**N**	**Y**
5.2	**Is there a Human-on-the-Loop (HOTL)?**	Binary (Yes / No)	**N**	**Y**
5.3	Are there mechanisms to evaluate algorithmic performance?	Binary (Yes / No)	N	Y
5.4	What is the degree of discretion in human and AI decision making?	Qualitative Evaluation	Subjective	Subjective
5.5	Are there mechanisms to adjust the AI to handle poor outputs?	Qualitative Evaluation	Subjective	Subjective
6	**EXPLAINABILITY**			
6.1	**Are algorithmic decisions explained?**	Binary (Yes / No)	**N**	**Y**
6.2	Are explanations intelligible to end users externally?	Binary (Yes / No)	N	Y
6.3	Are there measures to provide an explanation, even for black boxes?	Binary (Yes / No)	N	Y
6.4	What measures enable explained and intelligible decisions?	Qualitative Evaluation	Subjective	Subjective
7	**IMPACT ASSESSMENT**			
7.1	**Can the output adversely impact taxpayers?**	Binary (Yes / No)	**Y**	**N**
7.2	**Is any fundamental right impacted?**	Binary (Yes / No)	**Y**	**N**
7.3	Can the output directly or indirectly generate legal consequences?	Binary (Yes / No)	Y	N
7.4	What is the nature and impact of each output?	Qualitative Evaluation	Subjective	Subjective
7.5	How will the outputs impact taxpayers' rights?	Qualitative Evaluation	Subjective	Subjective
7.6	What measures mitigate interferences with fundamental rights?	Qualitative Evaluation	Subjective	Subjective

Source: Authors.

REFERENCES

Aslett, Joshua. 2024. "IMFx: VITARA—Information Technology and Data Management." edX. https://www.imf.org/en/Capacity-Development/Training/ICDTC/Courses/VITARA-ITD

Aslett, Joshua, Gustavo González, Stuart Hamilton, and Miguel Pecho. 2024. "Tax Administration: Essential Analytics for Compliance Risk Management." IMF Technical Notes and Manuals 2024/01, International Monetary Fund, Washington, DC.

Cantekin, Kayahan, and U.S. Global Legal Research Directorate Law Library Of Congress. 2023. "Regulation of Artificial Intelligence Around the World." https://www.loc.gov/item/2023555920/

Gesley, Jenny, Kelly Buchanan, Michael Chalupovitsch, Tariq Ahmad, Clare Feikert-Ahalt, and U.S. Global Legal Research Directorate Law Library Of Congress. 2023 "Safety and Security of Artificial Intelligence Systems." https://www.loc.gov/item/2023555908/

González, Ignacio. 2018. "Analytics and Big Data the New Frontier." CIAT Review No. 44. https://www.ciat.org/Biblioteca/Revista/Revista_44/ingles/2018_TR44_ebook_ing.pdf

Hadwick, David. 2024. "AI TAX ADMIN EU—Open-Access Repository of the Use of AI by Tax Administrations in the EU." University of Antwerpen. https://www.uantwerpen.be/en/projects/aitax/

Hu, Krystal. 2023. "ChatGPT Sets Record for Fastest-Growing User Base." Reuters. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

MIT Technology Review Insights (MIT). 2023. "The Great Acceleration: CIO Perspectives on Generative AI." MIT Technology Review, Cambridge, MA.

OpenAI. 2023. "GPT-4 Technical Report." https://cdn.openai.com/papers/gpt-4.pdf

Organisation for Economic Co-operation and Development (OECD). 2020. *Tax Administration 3.0: The Digital Transformation of Tax Administration.* Paris: OECD Publishing. https://doi.org/10.1787/ca274cc5-en

Organisation for Economic Co-operation and Development (OECD). 2023. "Inventory of Tax Technology Initiatives." OECD Forum on Tax Administration. https://www.oecd.org/tax/forum-on-tax-administration/tax-technology-tools-and-digital-solutions/

Webson, Albert, and Ellie Pavlick. 2022. "Do Prompt-Based Models Really Understand the Meaning of Their Prompts." Department of Computer Science, Brown University. https://arxiv.org/abs/2109.01247

World Customs Organization (WCO). 2023. "Research & Policy Note on Generative Artificial Intelligence for Customs." https://www.wcoomd.org/-/media/wco/public/global/pdf/topics/research/papers/researchpolicynote_genai_en_06122023.pdf?la=en

GLOSSARY

Term	Definition
Analytics	The practice of using mathematics and statistics to analyze data, identify patterns, develop insights, and support decision making.
Advanced Analytics	A subset of analytics that uses more complex techniques to predict future outcomes, including machine learning, predictive modeling, data mining, and statistical analysis.
Algorithm	Logic, methods, or sets of specific steps for solving a problem that may or may not be codified for use in a computer programming language.
Artificial Intelligence (AI)	The simulation of human intelligence, carried out using computer systems and software, relying heavily on statistics and data science.
Bias	Subjective prejudices, whether of a human or an AI system or tool, that unjustifiably favor or disfavor individuals and groups.
Black Box	A system or model whose internal workings are not transparent or easily understood, making it difficult to interpret how it processes inputs to produce outputs.
Big Data Platform	Specialized software and hardware designed to store, process, and analyze large and complex data sets.
Bot	Short for "robot," an automated software program designed to perform tasks or interact with users, often through text or voice interfaces, mimicking human behavior.
Chatbot	A software program designed to simulate conversation through text or voice interactions.
Control Gate	A decision point or checkpoint that must be passed before a change can proceed to the next phase of the change management process, often through governance structures in the form of committees.
Data	Facts, figures, measurements, or observations collected and used as a basis for reasoning, analysis, discussion, or calculation.
Data Governance	The processes and mechanisms by which data life cycle strategies are administered and enforced, accounting for data quality, security, controls, and compliance.
Data Mining	The use of statistics and data science to analyze large data sets and identify anomalies, correlations, patterns, and trends.
Data Science	An interdisciplinary field that combines statistics, computer science, and domain knowledge to analyze data and develop intelligence products that support decision making and solving of complex problems.
Data Set	A set of logically related data often organized in one or more structured tables and used for analysis, processing, and research.
Data Warehouse	A central repository of data used for analytics and reporting, which sources its data from other databases and is sometimes designed to serve as a single source of truth for an organization.
Digital Services	The provision of IT resources to facilitate or perform tasks often, but not always, making use of business systems.
Digital Transformation	The process of digitizing work, optimizing, and then transforming business processes to improve efficiency, enhance customer experiences, and create new business models.
Expert Rule (ER)	Logic in the form of a rule derived from the specialized knowledge and experience of experts in a particular field, used to inform decision making, problem-solving, and the creation of automated systems.
Expert System	A computer program that uses artificial intelligence to simulate the decision-making ability of a human expert in a specific field, building from a legacy of using expert rules and knowledge sourced "top-down" from subject matter experts.
Explainability	The degree to which the logic applied by AI to reach a conclusion can be intelligibly understood.
Exploratory Analysis	A shorthand description for Exploratory Data Analysis, which is the process of using analytics to identify, understand, and summarize the main characteristics of data sets, often using visual methods to reveal patterns, trends, and anomalies.
Foundation Model	A general-purpose, often large AI model that serves as a base for fine-tuning or adapting to specific tasks and applications.

Term	Definition
Hallucination	A scenario where AI outputs based on patterns in data do not reflect actual reality.
Generative AI (GenAI)	A category of AI technology that is capable of producing content (text, data, images, videos) comparable to outputs created by humans.
Individual Duty of Care	The obligation of public servants to act responsibly and ethically in the performance of their duties, ensuring that their actions do not cause harm to individuals or the public.
Large Language Model (LLM)	A category of AI foundation model trained on vast amounts of text data, designed to understand and generate human-like language.
Machine Learning (ML)	A major branch of artificial intelligence, involving the use of data and algorithms to enable computers to learn without explicit programming.
Machine Visioning (MV)	The use of computer systems to capture, interpret, and analyze visual information, enabling them to understand and take decisions based on images and videos.
Natural Language Processing (NLP)	A category of AI focused on enabling computers to understand, interpret, and generate human language.
Simplexity	The trade-off between simple and complex aspects of systems where, particularly with AI, value and risk need to be carefully considered.
Social Network Analysis (SNA)	The analysis of social relationships and structures through examination of networks formed by individuals or organizations, often relying on a category of specialized technology and tools.
Statistics	A body of mathematical science that seeks to describe and interpret data using methods that can provide measures of probability, uncertainty, and variability.
Training Data	Data used to train algorithms by enabling them to identify and learn patterns, which help them make informed predictions or decisions.
Use Case	A specific business scenario within which an IT system, service, or other tool is employed to carry out defined business objectives.
Use Case Inventory	A formal catalog of AI use cases, along with both summary and use case attributes, allowing for active and deliberate management.
Virtual Assistant	More robust than a Chatbot, a digital service, software application, or system designed to perform tasks, provide information, and assist users through text or voice interactions.

PUBLICATIONS

Understanding Artificial Intelligence in Tax and Customs Administration

TNM/2024/06

ISBN 9798400290435