December 2025 Edition 4

# In AI We Trust?

## Inside the New Era of Auditable AI in Tax

TTMAG

Curated by e-Bright

# Contents

# TRUST IN AI

In the past year, conversations about artificial intelligence (AI) in tax have moved from speculation to consequence. Tools that sat at the edges of experimentation now influence how data is organized, designed, and how authorities examine filings in increasingly granular and automated ways. The question we now raise is whether these systems can be trusted, and under what conditions.

This shift reveals a more complex form of trust than the profession initially anticipated. Trust in AI is not limited to the technical qualities of a model. Elements such as accuracy and quality matter, but they address only half the problem. The other half sits inside organizations themselves. Professionals remain cautious, sometimes openly skeptical, about relying on technology they do not yet fully understand. Much of that skepticism is rational. When the stakes involve compliance exposure and reputational risk, hesitation is a form of responsibility, not resistance.

Thus, this edition makes this tension clear, showing that trustworthy AI in tax is built through two parallel processes. One strengthens the architecture through data that is complete and well-structured; systems that allow traceability; controls that prevent models from drifting into unreliable territory; workflows that ensure human review where interpretation is required. The other strengthens the professionals who use it. Understanding the limits of automation; teams who learn how to question outputs instead of deferring to them; organizations that create space for experimentation rather than expecting instant certainty; leadership that treats adoption as a cultural shift rather than a technical rollout.

Both dimensions must advance together. A system can be technically sound and still fail if the people using it do not trust it. Likewise, enthusiasm alone cannot compensate for weak data, unclear processes, or inadequate governance. The year ahead will require organizations to confront both realities at once. The landscape entering 2026 will reward those who can integrate these two forms of trust. The future of tax will not be shaped by AI alone. It will be shaped by the quality of the systems beneath it, and the confidence of the people who rely on them.

Zoe Frattini, Editor-in-Chief

INTERVIEW WITH

# Prof. Dr. Robert Risse

WU-Director of Tax Law Technology Center (TLTC)

# Governance First, Technology Second: Rethinking AI's Role in the Tax Function

AI has accelerated the conversation about the future of the tax profession, but beneath the excitement lies a deeper reality: AI's potential is limited not by technological capability, but by whether organizations understand their own processes, their data structures, and the governance frameworks that shape decision-making. Few voices articulate these foundations as clearly as Dr. Robert Risse, Director of Tax Law Technology Center at WU (Vienna University of Economics and Business).

With a background rooted in accounting, data-driven decision processes, and technology-enabled tax execution, Risse has spent years analyzing how tax functions evolve when confronted with automation, algorithmic tools, and, most recently, large language models (LLMs). He approaches AI neither with idealism nor alarmism, but with a precise, structural perspective: tax professionals cannot meaningfully leverage AI unless they first understand how information is created, how processes are designed, and how rules can be translated into machine-executable logic.

In this interview, he reflects on the relationship between symbolic and sub-symbolic AI, the true nature of hallucination, the opportunities and limitations of standardization, the changing competency profile of tax experts, and why encoding tax law into machine-readable formats may be the most transformative development of the coming decade.

**Before we dive into AI, can you briefly describe what your new book is about and what drove you to write it?**

The book began out of five years of lectures, and bachelor theses focused on the digital side of taxation. I run a center called 'WU Tax Law Technology (TLTC)', and over time it became clear that while everyone in the profession loves talking about technology, very few truly understand the underlying principles. What motivated me was precisely this gap. We needed a resource that returned to the fundamentals, one that explains what the key processes are, how they connect, and which types of technology are suitable for which tasks. When ChatGPT 3.5 emerged in late 2022, the hype around LLMs escalated enormously. It reminded me of the earlier blockchain wave, a surge of enthusiasm with little execution.

My book argues that before moving to the next trend, we need to understand what we are actually capable of executing. We cannot allow ourselves to be carried from hype to hype without grounding.

The book also introduces practical examples so readers can see how developments translate into real execution. Importantly, it revisits the old IBM concept from the 1960s: people-process-technology, to which I now add data, because LLMs force us to rethink what sits at the foundation. Technology is no longer the limitation. Technology can now do almost anything; the limitation has become our understanding of how to structure information and how to execute. This is why the book ends by reflecting on symbolic and sub-symbolic AI. Symbolic AI refers to rule-based systems, and sub-symbolic to data-driven approaches like LLMs. Experts in physics and computer science now argue that LLMs will reach certain limits, not because they aren't impressive, but because the next frontier lies in combining rules and data more intelligently. That's both exciting and challenging, because it tells us we still have much to learn.

**In one of your posts, you said AI and tax involve not just technical nuances but temporal ones as well. What does that mean in practice?**

For tax professionals, the first step is always to understand how to convert tax rules, what I call 'taxable sentences', into process logic. Process management is the newest frontier in business management. If you can't structure your processes, you can't automate anything; and if you can't automate, AI becomes meaningless.

The temporal nuance is that technology should only enter the picture after you have structured your processes. First, you understand what you need to execute; you define the process, and only then do you ask which technology is appropriate.

This is where symbolic versus sub-symbolic AI as a next AI development becomes relevant. Symbolic AI is rule-based, and sub-symbolic AI is data-driven. In taxation, both are needed. In addition, the profession often jumps straight to LLMs without considering simpler tools like robotic process automation or chatbots, which might solve 70% of needs. Not everything requires an LLM.

LLMs are transformative precisely because they remove barriers that used to require heavy data cleaning. Years ago, you had to clean your data before RPA could touch it, but today, an LLM can interpret messy data directly. The challenge now is learning how to combine these tools, how to connect RPA, APIs, chatbots, and LLMs into a coherent workflow.

NotebookLM, Google's new AI environment,

is a good example. It encapsulates sources on the left, PDFs, websites, internal documents, and restricts the LLM to those sources that can be described as a RAG (Retrieval-Augmented Generation). This dramatically increases accuracy. You can tell the LLM: 'Summarize this' or 'Draft a memo based on the summary'. On the right, you now have LLM agents that turn those summaries into blogs, presentations, mind maps, tasks that tax managers spend massive amounts of time on. This is 70% of the profession's administrative workload, and it can be automated today.

When people worry about hallucination, I remind them that hallucination is merely the result of insufficient input. If you feed an LLM only two PDFs, it can only work with the information in those PDFs. If the PDFs contain gaps, the model reflects those gaps. Humans hallucinate too if they receive only half the information they need. Hallucination is therefore an input problem, not a technology problem.

**Our theme this year is whether we can trust AI in tax. What do you see as the biggest barriers to trust?**

A major OECD project I participated in, the 'Trustworthiness of AI' working group, explored precisely this. Tax administrations and universities debated whether authorities should be allowed to use all the data taxpayers send them, and whether algorithms can legitimately draw inferences from that data.

The obstacles to establishing trust can be traced back to a fundamental principle articulated by Bill Gates in a straightforward manner: "garbage in, garbage out." While AI systems are not equipped to evaluate the ethical nature of an input, they can determine its factual accuracy. The onus for this responsibility ultimately rests with humans.

For the foreseeable future, the 'human-in-the-loop' model is essential. AI can execute simple tasks autonomously where the error rate is negligible, but in complex contexts such as transfer pricing documentation, a human must review the AI's output before anything is submitted to authorities. You must examine whether the result is reasonable. Trust is therefore not a matter of believing in the algorithm, it is a matter of verifying the output based on solid tax expertise.

When people say algorithms are black boxes, I think that misunderstands the issue. We also cannot 'see' how the human brain makes its decisions. What matters is not whether we understand every technical computation, but whether the results are reviewable and testable. AI does not need to be transparent in a mechanical sense, it needs to be auditable. That is where trust comes from.

**There is ongoing debate about whether AI in tax needs global standards. What is your view?**

We cannot standardize technology, because it evolves too fast. Any attempt to standardize AI models or architectures will be outdated long before it is written. But what we can standardize is the tax law.

The OECD's 'Rules as Code' initiative is one of the most promising developments in decades. The idea is simple: take the textual law and convert unambiguous parts into machine-readable code. You then have two versions of the law, the human-readable version and the machine-readable version, and the machine-readable version becomes the basis for automated execution. For simple provisions, this is entirely feasible.

This is the realistic path toward standardization. It is not about standardizing

technology; it's about standardizing inputs. If the law is encoded, software engineers can build consistent systems on top of it, regardless of whether they rely on symbolic AI, LLMs, or future architectures we haven't imagined yet.

We must avoid using standards from countries whose governance we may not want to depend on, China already has blockchain standards, for instance. But the biggest mistake would be believing we can standardize AI tools themselves. We cannot, instead we should focus on standardizing the law where possible.

**You said your book also distinguishes between hype and real-world use cases. What do you see as the most promising applications of AI in tax today?**

The most meaningful use cases are those that insert AI into specific process steps. Legal tech has already paved the way here. LLMs excel at interpreting law and summarizing guidance. If you ask ChatGPT (version 4o onwards) what documentation is required for a given inter-company service, it will give you a highly accurate summary of OECD guidelines and country-specific expectations, something that previously required hours of manual work.

Use cases evolve because process steps evolve. If the process step involves extracting information, validating data, drafting text, or searching for structured and unstructured data, then LLMs can already enhance that step dramatically. With the emergence of agents, speed and accuracy are even greater.

The crux of the matter is that tax professionals must possess a comprehensive understanding of the process structure to ensure the effective integration of AI. While proficiency in Python programming is not a prerequisite, a fundamental understanding of Python's capabilities is essential. This includes the ability to comprehend the integration of statistical models with LLMs within the Microsoft Power Platform. Additionally, it is crucial to discern when an LLM provides added value and when simpler tools are sufficient. This task, while not novel, necessitates a deepening of expertise.

**What competencies will the next generation of tax experts need?**

Tax professionals must broaden their horizons. They need to understand business processes end-to-end. In procurement, logistics, customs, supply chain management, tax sits inside all of these. Years ago, no one in tax bothered understanding the purchasing process. Today you must, because critical tax-relevant data is created there.

The same applies to customs processes, VAT processes, equal terms in procurement, rental contracts for property, and tax due diligence workflows. You must understand not only the tax rules, but the business process generating the underlying data.

This shift applies to companies of all sizes. Even tax advisers must evolve. Instead of selling only services, they will increasingly sell products, LLM-based agents that automate parts of the tax workflow for clients. These agents will not be coded by mid-sized firms, but they will be deployed by them, and advisers must understand how they function.
Knowledge must expand, not deep technical coding knowledge, but process knowledge, data knowledge, and the ability to integrate LLMs into workflows. Tax knowledge alone is no longer enough.

**Finally, if we revisit this conversation in 2030, what do you think will have changed?**

By 2030, I believe much of the tax law that is currently textual will be converted into code.

Once the law exists in machine-readable form, the application of technology becomes significantly easier and more reliable. A chatbot supported by LLMs could receive that code, interpret exactly which data it needs to retrieve from a company's systems, and apply the encoded rules consistently.

LLMs and agents will not disappear. Even if experts argue they may reach technical limits, it does not matter for tax, as it doesn't require optimizing production processes or complex engineering tasks. It is administrative, rule-driven, and based on structured logic. LLMs are perfectly suited for this environment.

VAT is a good example. The law is the same for small, medium, and large companies. And yet the execution varies wildly, sometimes disastrously. Coding could harmonize this by providing consistent, mechanized interpretation. It will simplify processes, strengthen compliance, and allow LLMs to apply the encoded law immediately.

This is the next stage. Coding the law will fundamentally reshape the relationship between AI and tax. And by 2030, it will be normal.

# We can't solve problems by using the same kind of thinking we used when we created them

Albert Einstein

INTERVIEW WITH

# Brigitte Baumgartner Garcia

Head of Product Development Compliance at TPBaumgartner

# Beyond Automation: Building Trustworthy AI in Tax Compliance

'AI in tax' has been synonymous with automation, data extraction, and workflow efficiency. Yet, as these tools begin to mature, a more fundamental conversation has taken shape, one centered on trust.

How can organizations ensure that the AI systems governing financial and legal decisions remain accurate, ethical, and explainable? How can professionals maintain control and accountability as machine reasoning grows more sophisticated?

To explore these questions, we spoke with Brigitte Baumgartner Garcia, who has spent her career at the crossroads of law, technology, and product innovation. In this interview, she shares her view on what makes AI truly trustworthy in tax, why companies struggle with adoption, and how human judgment and governance by design must evolve alongside technology.

**You've worked at the intersection of law, technology, and innovation. How do you define "trustworthy AI" in the context of tax compliance?**

To define trustworthy AI, you first need to understand the foundation of large language models. Their success, the reason they can now perform so well across domains, comes from advances in neural network architectures and what's known as retrieval-augmented generation (RAG).

traditional RAG has two components. The first retrieves relevant information from a vector database, where text is represented as mathematical vectors. The second, the augmented generation, uses the retrieved data as context to craft an answer. Here lies a fundamental flaw: these models have been trained to always give you an answer, even if it's wrong.

They're designed to satisfy you, not necessarily to tell the truth. In highly regulated domains like tax, you don't need an answer, you need a correct answer. That's where trust begins.

To move closer to that, we're now seeing the emergence of graph-RAG, which uses knowledge graphs in addition to vector databases. Companies that want trustworthy AI must train their models on controlled, internal data – legislation, case law, and their own corporate records – and structure that data in a way that allows logical connections.

Combining vector-RAG and graph-RAG creates a hybrid-RAG system. Hybrid-RAG reduces hallucinations dramatically because the model draws from verified, structured relationships instead of a loose pool of internet text. For me, a trustworthy AI model is one built with a hybrid system, trained on controlled data sources rather than the open web. In tax compliance, that's the difference between a model that guesses and one that reasons.

**Many professionals still find AI adoption daunting. What makes tax such a complex environment for implementation?**

Because too often companies begin with the technology instead of the problem. They jump straight into automation without first asking, "What do we actually want to solve?". The first step in any adoption process should be a discovery phase, understanding what problem you want to solve and what outcome you want to achieve.

Instead, many companies start by buying tools before defining objectives. In tax, this is dangerous because you're operating in a high-risk environment. You can't hand over critical decisions to an AI system before knowing precisely what tasks are suitable for delegation. And you should be automating when it makes sense and what makes sense. That will always depend on the process and people.

AI implementation should follow a low-risk, high-impact logic: delegate repetitive, time-consuming tasks that don't require complex judgment, while keeping humans firmly in the loop for compliance decisions. Knowing precisely what tasks are suitable for delegation. You should be automating when it makes sense and what makes sense. That will always depend on the process and people.

AI implementation should follow a low-risk, high-impact logic: delegate repetitive, time-consuming tasks that don't require complex judgment, while keeping humans firmly in the loop for compliance decisions.

Another challenge is architecture. Once you define the desired outcome, the AI infrastructure must be built around it, not the other way around. If you skip that strategic step, every subsequent decision, from data structuring to model selection, becomes a guess. Thus, the biggest obstacle is not technical capability; it's the lack of business clarity and process design.

**Governance is becoming the keyword in every AI discussion. What does "governance by design" actually look like?**

Governance by design means planning for traceability, auditability, compliance, and adaptability from day one.

A well-governed system is traceable. You can see how every answer was generated. That's another advantage of the hybrid-RAG approach: it doesn't operate on a black box. It's AI built on knowledge, not just data. You can follow the reasoning path, understand why the model reached a specific conclusion, and correct it when it's wrong.

Then comes auditability. Every step (data input, processing, output) must be documented and retrievable. And compliance must be built in, not added later. The model should reflect legal requirements and

regulatory updates in real time.

Finally, it demands adaptability. Legislation changes constantly. If your AI architecture can't adapt quickly, if fine-tuning takes weeks, then it's not governance-ready. A trustworthy model must evolve at the same speed as regulation.

**Tax has become increasingly cross-functional. How do you see collaboration between tax, legal, IT, and business teams shaping AI development?**

Tax used to be a self-contained world. You dealt with the tax team, they filed returns, and that was it. It has now become deeply interconnected. Cross-functional collaboration is not optional; it's the only way to make AI both innovative and compliant.

However, it's true that these interactions can slow things down, not because cross-functionality itself is the problem, but because large organizations aren't agile enough, they have rigid processes and long decision chains. Europe is in love with process, and while that ensures safety, it often kills speed.

So, if we want AI to thrive, companies must connect departments but also rethink structure by streamlining approvals, encouraging experimentation, and embracing agile principles.

**What will it mean to be an "AI-ready" tax professional next year?**

Even today, an AI-ready tax professional can't just understand tax. They must also understand how products are built, the language of technology and how to communicate effectively with developers. Tax experts are the subject-matter managers, but to shape the tools they'll use, they must explain what problems need solving and that makes them innovators not only subject matter experts. That doesn't mean everyone has to become a prompt engineer or a coder. It means learning to collaborate, to define use cases, evaluate results, and understand the limits of the technology.

Accessibility is also improving. With new tools, Agentic tools launched just recently, you can build your own AI workflow with almost zero coding. All you need is curiosity and the persistence to keep learning and testing until it works. That's what the future of "AI readiness" will look like, less about engineering and more about applied understanding.

**As AI becomes more capable, how do we avoid over-reliance and the erosion of human expertise?**

If we design systems where humans retain visibility, agency, and accountability, expertise won't disappear; it will be amplified. The fear of losing control is understandable, but we're the ones building these systems. We decide their purpose.

The real question is what kind of AI we are building. Is our goal to replace human judgment, or is it to enhance it? We should be developing systems that challenge us, that make us think better. Imagine an AI tool that doesn't just give you the answer, but asks: "Have you considered this case law? Have you noticed this inconsistency?" That's the kind of AI that would make tax professionals smarter, not redundant.

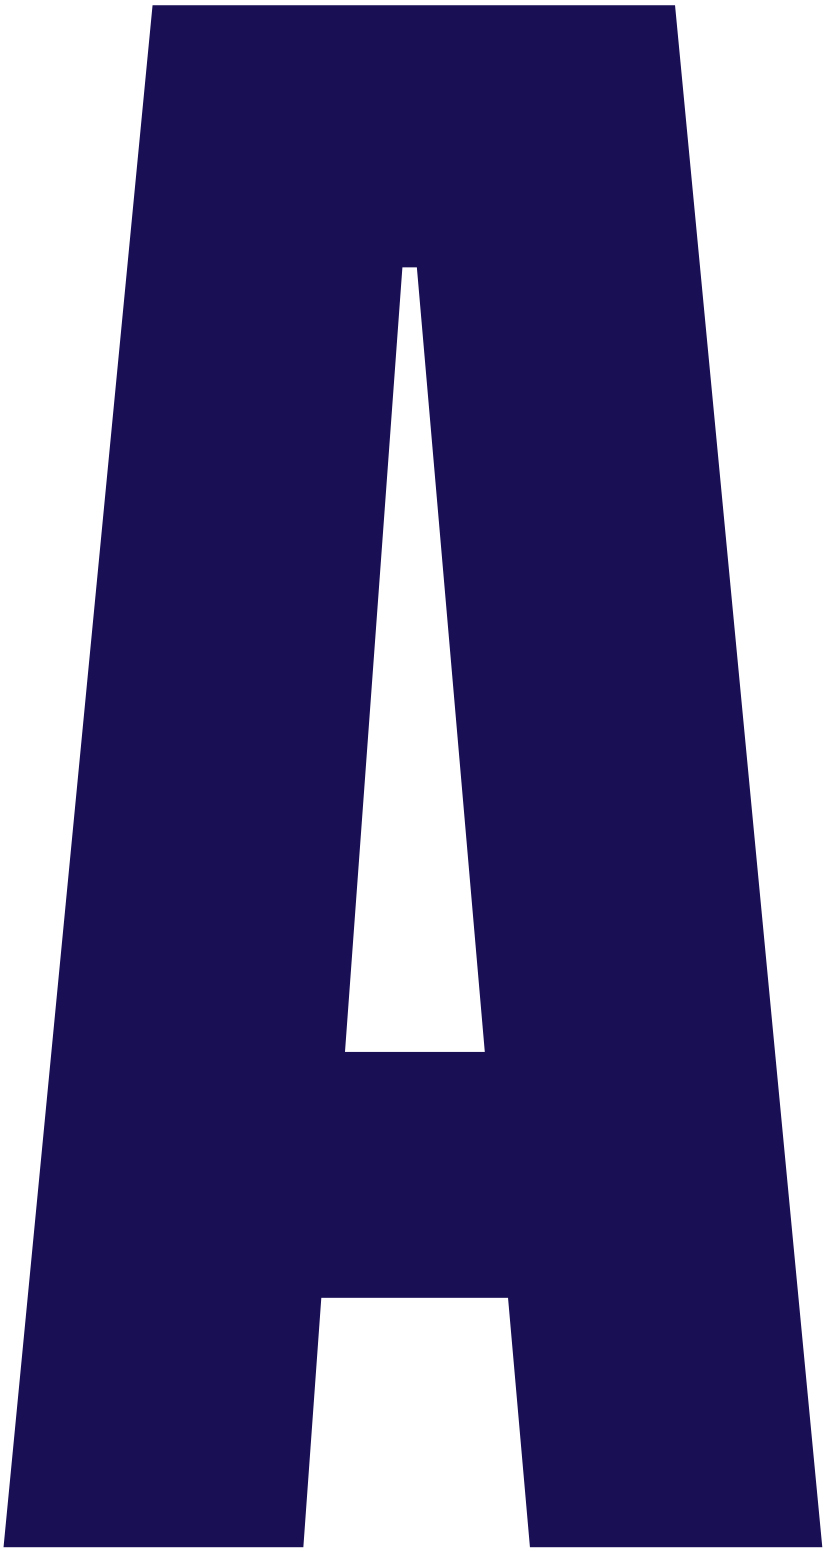

AI

is a mirror, reflecting not only our intellect, but our values and fears

**Ravi Narayanan**
VP of Insights and Analytics at Nisum

INTERVIEW WITH

## Keval Hutheesing

CEO at Cygnet.One

# From Proof of Concept to Proof of Trust: Scaling AI in the Tax Function

For years, tax departments have experimented with artificial intelligence, running pilots, building prototypes, and exploring automation through generative models. As the field moves from experimentation to execution, a critical question surfaces: can we trust AI in tax?

In this interview, Keval Hutheesing, a technology leader with deep experience in scaling enterprise AI platforms across jurisdictions, shares what it truly means to build trust in AI, what stands in the way of production-scale deployment, and how explainability, governance, and human oversight will define the next era of digital transformation in tax.

**Let's begin with the central question of this issue: what does "trust in AI" mean for tax operations?**

We usually frame it across three pillars: security, accuracy, and compliance. First, security is non-negotiable. AI systems handle financially sensitive and client-specific data, often critical to business operations. Unless organizations can guarantee data security throughout the processing lifecycle, users won't risk scaling these systems with real client data. Strong testing, encryption, and secure data-handling pipelines are foundational to trust. Second, there's accuracy. We all know the concerns around generative AI, such as hallucinations, inconsistency, and context loss. The goal isn't perfection; it's reliability. Accuracy will never be 100%, but achieving a stable, auditable level of correctness is key to user confidence.

Finally, compliance. Many early AI tools, like ChatGPT, were trained on outdated datasets, leading to responses that didn't reflect current regulations. In tax, that's unacceptable. Models must be kept current, with regular updates and governance checks to ensure they align with evolving compliance standards. Trust, in this context, isn't just

about whether the AI "works." It's about whether it's secure, reliable, and regulation ready.

**Moving from pilot projects to production-scale AI exposes new challenges. What's the biggest gap you see in that transition?**

The biggest gap is demonstrating real value. According to a recent MIT study, 95% of enterprises have yet to see measurable P&L impact from their AI initiatives. That figure is telling.

Several issues explain it. One is accuracy, if people don't trust the output, they won't adopt it. Another is predictability, as AI systems that behave inconsistently or produce variable results undermine confidence and scalability. However, the most fundamental problem is data. Without a comprehensive data strategy, you can't scale AI responsibly. Data is the foundation on which all models are built. Many companies underestimate the quality, completeness, and contextual structure of the data feeding their systems.

Another lesson is that general purpose tools like ChatGPT rarely succeed at enterprise scale. The companies that are making progress are those combining domain expertise and technical specialization, bringing together professionals who understand tax logic with engineers who build solutions tailored to that domain. Specialized AI systems, rather than generic ones, are what generate true operational value.

In tax, audit trails are central to accountability. What principles do you follow to ensure AI outputs remain explainable and traceable?

We design AI auditability the same way we design software audit trails. There are three essential layers.

The first is identity verification. Every interaction must be attributable to a specific, authenticated user. Whether someone is querying an internal AI assistant or deploying an automated decision engine, all actions should occur after login. Similarly, any code change or modification in a system must be linked to a verified identity. The second layer is comprehensive logging. Every query, every system call, every output should be recorded. Who initiated it, what response was generated, and what process was triggered. This makes the AI system as traceable as an IT infrastructure audit. The third is multi-agent traceability. We often use multi-agent systems where one agent interprets the query and routes it to another specialized agent. For example, one might handle compliance guidance while another executes calculations. By logging each agent's action and decision path, we can identify precisely where an error occurred, whether in routing, reasoning, or execution.

Finally, we audit data flow itself. Each layer of the system must have defined controls to verify what data goes in, what transformations occur, and what comes out. Only through this level of granularity can AI outputs be transparent enough to withstand regulatory scrutiny.

**Generative AI also introduces the problem of hallucinations, outputs that are factually wrong or fabricated. How do you mitigate that risk?**

We spend a significant amount of time addressing this, because hallucination in a compliance heavy environment can be catastrophic.

One method is tight scoping. The narrower and more structured your prompt, the lower the risk of the model drifting into creative or inaccurate territory. When the AI knows

exactly what to answer, in what format, and within which parameters, reliability increases dramatically.

Another highly effective approach RAG. Instead of "making up" answers, the AI retrieves information from a curated internal repository – such as a company's tax knowledge base or document archive – and serves it back with context. Since the source material is known and fixed, hallucinations are minimal.

We've even implemented citation-based generation. Each answer includes references to its source documents (laws, rulings, or policies), so users can verify where the output came from. For instance, if the AI provides advice on a specific tax treatment, it cites the legal basis and links directly to the governing regulation or case law. That traceability reinforces confidence in the system.

More recently, we've begun testing automated reasoning checks, like those offered by AWS. They allow us to define logic parameters and run validation tests against every AI output. If the response violates those parameters, it's automatically flagged.

At the end of the day, human-in-the-loop oversight remains indispensable. AI should accelerate analysis, not replace judgment. Just as calculators didn't eliminate accountants, AI won't eliminate tax professionals, it extends their reach and efficiency while preserving their responsibility.

**Prompt governance is becoming as important as model governance. How can organizations make sure their prompts remain consistent, secure, and auditable?**

It starts with ownership. Someone within the organization, such as a Chief AI Officer, AI Governance Lead, or Prompt Governance Officer, must be accountable for defining and maintaining prompt standards.

Next comes the prompt lifecycle. Prompts should be created, tested, reviewed by independent experts, and only then approved for production use. Especially in regulatory environments, every prompt must follow a documented life cycle with checks at each stage. Then there are guardrails. These ensure that sensitive or personally identifiable information isn't entered into prompts or shared with external systems. Clear boundaries need to be defined for what data can and cannot be used, particularly in finance and tax, where confidentiality is paramount.

Finally, automation can help enforce those guardrails. In our own systems, we've embedded automated prompt-checking features that detect both data sensitivity and misuse. If a user attempts to submit an inappropriate or irrelevant prompt; for example, using enterprise AI to generate unrelated content, the system rejects it and alerts the user. This layered governance approach that combines human oversight, lifecycle control, and automated enforcement, builds a sustainable, scalable framework for responsible AI usage.

**A lot of professionals fear being replaced by AI. You seem to view it differently, as collaboration. What does effective human-AI collaboration look like in tax?**

It's absolutely collaboration. AI shouldn't replace human effort; it should multiply it. We see this clearly in four practical areas.

First, knowledge search. Instead of spending hours combing through databases or PDFs, professionals can now query vast internal knowledge repositories instantly. Second, document processing, particularly for invoices and forms, where AI can extract and classify

data from hundreds of formats faster than humans ever could. Third, AI co-pilots or assistants. These are tools that sit alongside the professional, offering real-time support by answering questions, drafting explanations, or suggesting compliance actions based on the task at hand. We've seen this pattern in software development as well: human-AI collaboration consistently outperforms either working alone. Fourth, customer or internal support. Many first-line queries in tax software are repetitive, "Where can I find this button?" or "How do I upload a file?" AI can handle those, freeing specialists to focus on the complex, judgment-based questions that really matter.

So, the value isn't in replacement; it's in redistribution. AI handles repetition. Humans handle reasoning. Together, they elevate the quality of work.

**For that collaboration to work, do professionals need to reskill, or will entirely new roles emerge?**

Both. Reskilling is essential, but so is redefining roles. Internally, we've launched an AI Champions initiative. Each department designates someone to become its AI lead. These individuals receive specialized training on the latest AI tools, then return to their teams to train others. It creates a self-reinforcing learning culture. Externally, partnerships also matter. Organizations should bring in specialized trainers and domain experts who can contextualize AI for tax use cases. It's not enough to learn the technology; professionals must understand how to apply it responsibly within regulatory frameworks.

We also believe in measurement. Every upskilling initiative should include metrics, such as time saved per task or improvement in accuracy, to quantify ROI. For instance, in our development teams, we compare how long tasks take with and without AI, measuring productivity gains and tracking progress over time.

Reskilling, new roles, and measurable learning outcomes together create a workforce that doesn't just adapt to AI but accelerates with it.

**Finally, where are you seeing real operational deployment of AI in tax today?**

We're beginning to see meaningful production use cases, particularly in four areas. First, knowledge-based search. Many tax firms have deployed RAG-based chatbots on their internal document repositories, giving professionals immediate access to contextual information. When outputs are explainable and traceable, adoption increases significantly.

Second, document and invoice processing. Generative AI now handles unstructured inputs that older systems struggled with. There's no longer a need to manually map fields for every new format.

Third, support query resolution as automated agents now manage standard user queries, reducing response times and improving internal service levels.

Lastly, tax code classification, particularly in consumer goods sectors, where vast product catalogs require correct tax mapping. AI can cross-reference product descriptions with regulatory databases far faster and more accurately than manual teams.

These examples show that we're finally moving from hype to deployment, from proof of concept to proof of trust. The next step is scaling it responsibly, with explainability, governance, and human oversight built into every layer.

INTERVIEW WITH

# Anita Richter

Head of Indirect Taxes at Springer Nature

# Trust Starts With Data: Building the Foundations for AI in Tax

As artificial intelligence becomes embedded in global tax operations, trust has emerged as the defining topic. The focus has moved beyond prototypes and proofs of concept, i.e. it now revolves around data integrity, cross-functional collaboration, and how to prepare organizations for AI that can be audited, explained, and relied upon.

To explore this intersection between technology, compliance, and people, we spoke with Anita Richter, Head of Indirect Taxes at Springer Nature. Her insights offer a candid view of what trust in AI really means for the tax function and why the human and organizational dimensions remain the hardest to engineer.

**In previous conversations you argued that we can't trust AI until we trust our own data and systems. What did you mean by that?**

The biggest obstacle is fragmentation, which creates inflexibility. Whenever you have highly customized or fragmented systems, flexibility disappears. You can't respond quickly to new regulatory or business requirements if your data sits in isolated silos in a multiple systems landscape. Flexibility is key, not just for tax, but for the entire organization. Even if your systems appear to work well today, you must ask: are these agile enough for tomorrow? That's the question we should all be asking ourselves. This is why many companies, including Springer Nature, are moving to Cloud-based environments. Cloud migration is often the first step towards building the infrastructure that allows AI to function properly, because once data is centralized, one can start to think about what AI can actually do.

It can add value in many ways: predictive analytics, pattern recognition, or even basic data cleaning. However, none of that matters if your systems aren't properly connected and your teams don't sufficiently collaborate. From the tax department's perspective, AI will only work if we move closer to IT. The two departments must understand each other's language and be willing to change perspectives. Without that bridge, we'll never see AI's true potential in tax.

**So, trust isn't just about trusting the tool, it's also about trusting people and processes?**

Exactly. We often talk about trust in AI as if it's purely technical, "does the model give the right answer?" But there's another dimension: are the people using it equipped to do so responsibly? AI is still just a machine. It will only perform as well as the person behind the screen. That means professionals need to learn how to interact with it; how to ask the right questions, how to formulate good prompts, how to evaluate whether an answer makes sense.

Curiosity plays a huge role here. You can't simply deploy a system and expect people to trust it. They need to experiment, to learn through trial and error. Building trust is as much about confidence in one's own use of AI as it is about confidence in the technology itself.

**How do you make teams more open to using AI responsibly and effectively?**

Leadership must lead by example. Everyone is talking about AI right now, and that creates pressure. People feel they must use it because others do. But adoption without purpose leads to frustration.

I recently came across an interesting survey stating that businesses claim to see little or no benefit from AI. I don't believe that will hold true for the future. We're simply still at the beginning of an exciting journey. Like any new technology, you need to experiment, - "fail fast, fail forward" - before you can see value. In our tax department, we started small; using AI for everyday support: summarizing documents, drafting emails, performing quick research. It's low risk, but it helps people to get comfortable. Once the team experiences real efficiency gains, curiosity naturally grows.

The second step is identifying use cases that go beyond the basics. I recently attended a conference in Frankfurt where a startup company presented a tool that uses Agentic AI for deep data analysis, completely transparent code, fully traceability. That kind of solution is highly interesting because it showcases how AI can evolve from simple assistance to powerful analytics.

But before any of that, you need to understand your data, where it lives, how it's structured, and how to access it. Otherwise, you can't even begin a conversation with IT about what AI could accomplish within your department.

**That brings us to data readiness. How do you assess whether your organization's tax data is truly AI-ready?**

That's a complex question. In tax, "data" covers much more than vendors and customers. It's transactional, master, and reference data, all of which must be accurate and synchronized.

At Springer Nature, we're moving to SAP Cloud, building on our existing S/4HANA foundation. What's interesting is that SAP itself is evolving towards data governance by design. It now offers modules that automatically feed into master data governance structures, real-time address verification, VAT ID checks, and integration with external validation services.

This is where the concept of a "golden record" becomes crucial. You need one core system that holds the definitive, always up-to-date version of every data point, customer, vendor, transaction, and all other systems draw from it. When your golden record is correct, the rest follows.

Maintaining that quality is another challenge. You can either dedicate internal resources to continuously monitor and update data, which

is difficult, given limited human capacity, or you can rely on external services that do it automatically. Increasingly, AI itself can assist with anomaly detection and consistency checks.

So, the answer to "how do you make data AI-ready?" is twofold: design robust data governance processes and let AI help you maintain them.

**What are the biggest data-quality blind spots tax teams face?**

It's not that tax teams underestimate data, we understand its importance perfectly well. The real issue is dependency. Many critical data points are owned by other departments than the tax department. It is not helpful if projects, that affect tax data, are initiated without tax being involved at an early stage.

As a result, by the time the tax team joins, decisions have been made that impact reporting, compliance, or structure. That's the real blind spot: the assumption that tax data is an isolated concern. Data quality is a shared responsibility, but within most companies it isn't treated that way yet.

**Who should ultimately own the data? The CFO? The IT function?**

Ownership is a fundamental question. In my view, data should sit under the CFO. The CFO's role is inherently strategic, to interpret financial data and turn it into business insight. You can't do that without owning the underlying data and ensuring its quality. The CFO should therefore also oversee data governance teams, defining clear processes, responsibilities, and accountability structures.

Once those are in place, you can tackle legacy ERP systems and start layering AI responsibly. But that sequence matters. You can't modernize with AI on top of weak governance. Governance must come first.

**What does a realistic maturity roadmap look like, from manual controls to trusted AI oversight?**

It's a huge project, and the first step is understanding that AI is not the goal. It's a tool to reach a goal. You begin by defining what success looks like, what is it you want to achieve. The second, and much harder, step is assessing where you stand. What systems do you have? How do these interact? What dependencies exist? Without an honest picture of the status quo, no roadmap will work.

Once that's clear, you slice the transformation into manageable pieces. If you try to fix everything at once, the scale of the task is overwhelming. Progress in finance and tax transformation happens step by step, but every step must move you closer to the final goal: clean data, trusted automation, and explainable outputs.

**Who should be responsible for explainability and AI transparency within a tax organization?**

Explainability can't be outsourced. Humans must stay in control of what AI produces. Every department should have AI representatives, people who understand both the business context and the technical side. They don't need to code, but they must be able to interpret how AI works, validate outputs, and communicate risks.

At Springer Nature, for example, we have strict internal AI policies. We cannot just plug sensitive business or customer data into public tools like ChatGPT or Perplexity. Every tool we want to use must go through a formal approval process because of data-privacy

and confidentiality obligations.

That bureaucracy can feel frustrating at times, but it's necessary. The balance between innovation and compliance is delicate, and transparency is what keeps it safe.

**How do you prevent large-language models from replicating inefficient or flawed processes?**

This is fundamentally a cultural issue. When teams are overwhelmed, they fall back on familiar patterns. They reproduce old processes because deadlines leave no room for reimagining them. But when the same people handle the same projects repeatedly, inefficiencies become invisible, and those inefficiencies end up embedded into AI systems as well. The antidote is perspective. Fresh people, new thinking, teams need space to question assumptions and redesign processes rather than replicate them. Transformation cannot be a one-off program with a beginning and an end, it must be a permanent element of business strategy.

Decisions have been made that impact reporting, compliance, or structure. That's the real blind spot: the assumption that tax data is an isolated concern. Data quality is a shared responsibility, but within most companies it isn't treated that way yet.

**Who should ultimately own the data? The CFO? The IT function?**

Ownership is a fundamental question. In my view, data should sit under the CFO. The CFO's role is inherently strategic, to interpret financial data and turn it into business insight. You can't do that without owning the underlying data and ensuring its quality.

The CFO should therefore also oversee data governance teams, defining clear processes, responsibilities, and accountability structures. Once those are in place, you can tackle legacy ERP systems and start layering AI responsibly. But that sequence matters. You can't modernize with AI on top of weak governance. Governance must come first.

**What does a realistic maturity roadmap look like, from manual controls to trusted AI oversight?**

It's a huge project, and the first step is understanding that AI is not the goal. It's a tool to reach a goal. You begin by defining what success looks like, what is it you want to achieve. The second, and much harder, step is assessing where you stand. What systems do you have? How do these interact? What dependencies exist? Without an honest picture of the status quo, no roadmap will work.

Once that's clear, you slice the transformation into manageable pieces. If you try to fix everything at once, the scale of the task is overwhelming. Progress in finance and tax transformation happens step by step, but every step must move you closer to the final goal: clean data, trusted automation, and explainable outputs.

If you design a process badly, your AI will simply operate badly at scale. That is why human creativity is more important now than ever.

Documentation has always been central to tax. How does it support trust in AI outputs?

Documentation is the backbone of trust. Without it, AI outputs cannot be audited. For tax professionals, it's not just about internal understanding; it's about external assurance. Auditors and tax authorities must be able to trace what happened, how a conclusion was reached, and what checks were in place.

Proper process mapping and documentation

transform AI from a "black box" into an auditable system. They make both regulators and professionals confident that the technology is being used responsibly.

**As tax authorities move toward real-time reporting and e-invoicing, how does that affect the trust equation for AI?**

I actually find the global development from monthly aggregated data reporting to real-time transactional data reporting exciting. For indirect tax, it has the potential to significantly reduce fraud and to increase tax honesty and equality.

At the same time, it is also a massive challenge. Every country has its own rules, timelines, and formats. The question then becomes whether your systems, and your organization are flexible enough to keep up.

AI will play a significant role here, not only for businesses but also for tax authorities who will use it to analyze enormous quantities of data, compare filings across different tax domains, and identify anomalies within seconds. This is both promising and intimidating.

Countries like China are already extremely advanced. They cross-compare banking data, invoice data, and reporting data in real time. Other countries still prefer paper. However, this will change. And yes, data privacy concerns must be taken seriously, but privacy can also slow down progress. Finding the balance will be a delicate debate. We will need to anticipate the kind of questions tax authorities' AI systems will define. That is a completely new dimension of compliance.

**Will we ever reach a point where AI decisions in tax are trusted without human review, and should we?**

I doubt that's a question with a single answer.

At a recent conference, in a room of highly qualified German tax advisers, the question was asked whether they thought they would still be needed in ten years from now. Everyone said yes, of course, "AI will never replace us."

To be honest, I'm not so sure. If tax advisers continue to focus solely on compliance and bookkeeping, I don't see much of a role left for them. Those areas will become largely automated.

AI's capabilities are evolving faster than most people realize. The hallucination problem will likely disappear as models improve and training data becomes richer. The bigger question will be what kind of information we feed into the machine – quality in, quality out.

In the end, AI won't eliminate tax professionals, but it will force them to redefine their role. Expertise, ethics, and judgment will still matter, but the way we express these will change completely.

From ‘Reshuffle: Who Wins when AI Restacks the Knowledge Economy’

AI is becoming an institutional infrastructure of sorts that shapes how organizations manage knowledge and allocate attention, and - with that - how they make decisions.

**Sangeet Paul Choudary**

INTERVIEW WITH

## Vishnu Bagri

Partner at Singhvi Dev & Unni LLP, Chartered Accountants

# Reshaping the Tax Workplace: How AI Redefines Roles, Responsibility, and Reasoning

As organizations race to integrate artificial intelligence into their financial and tax operations, a deeper question emerges: how will this technology reshape not just workflows, but the very identity of the tax professional?

For Vishnu Bagri, a senior tax and technology leader who has advised global enterprises on tax functions and worked on technology led solutions, the future of work in tax must be understood along two dimensions; the individual knowledge worker and the organization that employs them. In this conversation, he explains how AI is driving both greater efficiency and a profound restacking of knowledge, authority, and accountability.

**When we speak about the workplace of tomorrow for tax professionals, what transformation do you foresee?**

We must approach this on two planes. One is the individual, the tax professional whose expertise and habits define how work gets done. The other is the organization, which must decide how it governs, invests in, and embeds AI responsibly.

On the individual side, I often think of a tax professional's role as resting on three pillars: the ability to read, to retain, and to analyze. AI will lighten the memory burden, it can instantly retrieve historical information, prior rulings, and precedents. However, it will make reading and analysis harder, not easier. Professionals will now need to interrogate the system to understand why a model arrived at a conclusion and whether that reasoning is sound in light of first principles. So, while AI introduces efficiencies and productivity gains, it also creates new forms of risk and new intellectual demands. We will need sharper analytical minds, people who can question the machine, not just use it.

From the organizational side, the use of AI is not simply a technological decision; it is an ethical and regulatory one. It demands significant investment to ensure data

confidentiality, regulatory compliance, and governance standards. Tax functions, by nature, handle highly sensitive and business-critical information, and the trust of both regulators and clients depends on how securely that data is managed. However, once those foundations are in place, AI can bring a level of consistency and integration across global tax functions that has long been elusive. Today, filings and returns are often managed by dispersed teams with different interpretations of local data and business context. AI, if implemented with strong architecture and data pipelines, can harmonize this complexity, enabling multinationals to manage tax processes globally with uniformity and control.

**Is it realistic to imagine a standardized AI workspace in tax, or will it always remain fragmented?**

I would say that intelligence, by definition, resists standardization. Just as human intelligence varies across individuals, artificial intelligence will too. What can, and must be standardized, is the AI-enabled workspace, not the intelligence itself. If I draw a parallel from mathematics, think of a lowest common multiple. That's what every tax professional will share: a baseline of governance, data access, and control. Beyond that, individual reasoning will differ.

In a practical sense, standardization should apply to data inputs and governance frameworks. The data that enters the AI ecosystem must be consistent, validated, and traceable. Governance guidelines, defining who can source information, how it enters the system, and how the output is reviewed, should also be standardized.

I imagine a future where each professional may have a digital twin or AI agent that mirrors their role, similar to how each of us has an identity on social media. Within an enterprise, your "tax avatar" will act based on your instructions, authority, and ethics, executing within clearly defined rules. The result would be a workplace that is automated, auditable, and scalable. One where workflows are transparent, and every action leaves a digital trail.

AI's reasoning itself should not be standardized; that would defeat its purpose. Its value lies in diverse reasoning patterns and contextual interpretation. The standardization belongs to the ecosystem with data integrity, audit trails, and compliance processes.

**You've mentioned compliance several times. Why do you see it as the main area for standardization?**

Compliance operates through structured, repeatable logic. That makes it ideal for automation, and therefore, for standardization. The inputs, approvals, and signoffs involved in filings can all follow a defined workflow.

We've seen a similar journey before with robotic process automation (RPA). RPA identified tasks with clear "if-then" logic, coded them, and achieved huge gains in efficiency. AI now amplifies that, it can analyze which processes to automate, generate the code, and even self-optimize those workflows over time. So, the evolution from RPA to AI is less about replacement and more about deepening automation with reasoning.

That's why I see compliance as the first domain where we'll see genuine standardization: in how data flows, how approvals are captured, and how every step is logged. The other critical area is governance: audit traceability, ethical use of

data, and information security. If these aren't standardized, the risks, from hallucination to misuse, become unmanageable.

**Do you expect tax authorities or international bodies like the OECD to push for AI auditing or standard guidelines?**

It's conceivable, but mandating AI use across jurisdictions would be enormously complex. For true regulatory standardization, each country would need its own controlled environment, its own language models and training data, and would then have to integrate those into corporate systems.

I don't see that happening soon, certainly not in the next few years. What I do foresee are governance guidelines that require organizations to demonstrate auditability; showing how an AI system reached a conclusion, what source data it used and what analytical logic it followed. Regulators will likely focus on traceability rather than prescribing the AI tools themselves. That, to me, is a more pragmatic and necessary step, which ensures transparency without stifling innovation.

**How do you see prompt governance and professional certification fitting into this future?**

Prompt governance is emerging as a new layer of control, but we must strike a balance. If governance becomes over-regulation, we'll lose the very creativity and critical thinking that makes AI valuable. In practice, prompt governance should focus on preventing misuse. For example, ensuring that people don't use enterprise AI systems to share highly confidential data, expose proprietary prompts, or train models with irrelevant or risky inputs. It's less about controlling what people think and more about controlling what data they feed into the system.

Now, when we talk about certification, I distinguish sharply between certified workflows and certified professionals. Certified workflows make sense in compliance, but AI is pushing us away from process-driven environments toward interaction-driven ones.

With AI, much of what we once called a process will become automated, the system will execute the "finite set of human variables" that used to define a workflow. What will remain is interaction: humans engaging with outputs, reasoning, and accountability.

That's where I see the emergence of accountability nodes, meaning individuals responsible for reviewing, signing off, and taking ownership of outputs before they move further in the system. So, I'm less convinced by the idea of certified workflows and far more supportive of certified professionals. These are people who understand the AI ecosystem, its audit traceability and its limits. Certification here creates responsibility, not just process compliance.

**How might AI connect tax professionals with data systems in a more auditable way?**

Think of a tax team with two key roles: the preparer and the reviewer.

The preparer uses an AI agent to gather data from multiple systems (ERP, financial statements, transaction records) and compile it into a compliance form. Every action the agent takes is logged: which data

sources it accessed, which analysis framework it applied, and whether it drew on predefined prompts or external data.

The reviewer then receives the same output, with a complete trace of the preparer's digital steps. Using another AI agent, the reviewer can cross-check this with prior filings, detect red flags, or reconcile with audited financial statements. Each agent's scope is tightly defined; they operate within permissioned data boundaries.

As these systems mature, you could imagine agent-to-agent collaboration, where preparer and reviewer agents interact automatically. Even then, every output must remain traceable to a responsible human, the person who validates, "Yes, this is the result my system has produced, and I stand by it." That's the essence of accountable automation.

**Finally, what guiding principles should tax leaders follow to ensure AI builds instead of eroding trust?**

It begins with purpose. Tax leaders must understand why they're introducing AI, what problem it solves and how it impacts both productivity and people. In compliance, the goal is efficiency: automating repetitive tasks. Still, every automation decision should be paired with a reskilling decision. Leaders must ensure that professionals learn data architecture and AI reasoning, to understand how systems analyze, classify, and infer. As AI scales, one professional will manage volumes of work that once required entire teams. That changes not just workflows, but organizational structures.

Beyond compliance lies business planning, where AI can empower real-time decision making. Tax has always been integral to pricing, procurement, and expansion decisions. With AI, tax professionals can move from retrospective analysis to live strategy support, advising business leaders as transactions happen.

Then there's the regulatory dimension. Tax authorities themselves are adopting AI to identify anomalies and red flags. That creates a new dynamic: how does a taxpayer anticipate the moves of a regulator equipped with its own AI? To engage meaningfully, professionals must understand both the law's first principles and the architecture of data.

In the end, trust in AI will depend on two capabilities: a deep grounding in legal reasoning and a sophisticated grasp of how information flows. The professionals who master both will define the next era of the tax profession.

Small and medium-sized enterprises (SMEs) are entering a new phase of tax compliance. Digital reporting requirements, rising expectations around documentation, and the growing influence of AI mean that processes once considered “big company challenges” now apply to businesses of every size. As this shift accelerates, SMEs need tools that simplify complexity without sacrificing accuracy or control.

This is the space where Lukas Zörner, Co-Founder and CEO of Integral, has built his focus. He works on creating intelligent and structured tax workflows that guide users through each step, prevent avoidable errors, and strengthen the relationship between advisers and their clients. His approach is rooted in human-centered design, where technology reinforces judgement by removing friction and clarifying decisions to make compliance manageable.

In this interview, he discusses the realities SMEs face, the importance of structured digital workflows, and how thoughtfully designed systems can bring confidence and clarity to businesses that often feel overwhelmed by compliance demands.

INTERVIEW WITH

## Lukas Zörner

Co-Founder & CEO at Integral

# From Proof of Concept to Proof of Trust: Scaling AI in the Tax Function

**Let’s start with your latest developments: What are the newest projects or features your team has been working on in the past year?**

Over the past year, we’ve sharpened Integral as an intelligent, human-centered tax platform for SMEs with two priorities: first, deeper automation as “silent infrastructure,” and second, seamless collaboration between businesses and experts. In practice, that means faster, context-aware workflows across payroll, tax, and unified data flows instead of document ping-pong. A key milestone was the integration of Cleverlohn following our M&A transaction, combining leading payroll expertise with our platform approach. The result: less friction, more reliability, and advisors who can spend more time on truly strategic work.

**How has your platform evolved in response to the growing adoption of AI across finance and tax departments?**

We shifted from feature-level AI to platform-wide infrastructure. AI now underpins data ingestion, and decision support across payroll, compliance, and reporting. The biggest win is reduced noise for advisors.

**Have you noticed a shift in the types of problems clients are trying to solve compared to a last year?**

Yes, clients are shifting toward real-time visibility to avoid month-end surprises and enable mid-period interventions. They expect inquiry resolution that reduces back-and-forth and delivers clear, context-aware answers, with escalation only when human judgment is needed. At the same time, the focus has moved from basic automation to assurance. Auditability, explainability, and traceable decisions matter more than ever. Clients want to understand why a rule applied, who approved an exception, and how outcomes link to source data. This builds confidence and lets them act quickly without sacrificing trust.

**What partnerships, integrations, or market expansions are you most excited about right now?**

We are deepening integrations with core finance stacks across ERP, payroll, and banking to eliminate data hops and duplicate entry and to enable cleaner, real-time workflows. In parallel, we are exploring distribution partnerships with fintech and accounting platforms to bring integrated tax workflows to more SMEs. Together, these steps expand reach, reduce friction, and make end to end compliance feel seamless.

**Based on the theme of this issue, what does trustworthy AI mean to you in a tax-specific context?**

Trustworthy AI in tax starts with humility. It should assist, not replace, human judgment where interpretation and intent matter. It also means security and compliance are built in from the start. Most of all, trust is earned through clarity and accountability. Outcomes must be explainable, decisions traceable, and exceptions reviewable. Expert oversight closes the loop, ensuring the system supports professionals rather than overshadowing them. When these elements work together, confidence follows and better decisions become faster and safer.

**What guardrails or validation mechanisms have you implemented to mitigate hallucinations?**

Most importantly, a human currently still stays in the loop for cases requiring interpretation to ensure expert review before anything affects compliance or reporting.

**How do you envision the future AI-augmented tax team? Will humans still drive interpretation, or will systems take the lead with human review?**

The future tax team will blend human judgment with AI support. Systems take on ingestion, reconciliation, triage, and first pass classifications, while advisors focus on edge cases, strategy, and narrative. Their skill set expands toward data interpretation and prompt craftsmanship. With real time-shared workflows and continuous model feedback, humans remain in the lead and AI amplifies their impact.

**In your experience, what's the biggest misconception tax professionals have about AI tools?**

A common misconception is that AI tries to replace expert judgment, when its real impact is speeding up routine work and strengthening human decisions. Another is that more automation always improves outcomes, even though quality depends on guardrails, clarity, and policy alignment. Many also assume adoption is a one-time setup, yet continuous calibration is essential. And finally, AI is often seen as a large-enterprise advantage, even though SMEs benefit the most from reduced friction and faster answers.

**How do you think organizations can foster a culture of AI literacy and accountability across tax, finance, and IT functions?**

From my experience, it starts with agreeing on a shared set of principles and talking openly about what trustworthy AI means for the organization. When teams across tax, finance, and IT align on values like transparency, auditability, and human oversight and see them reflected in actual policy, it creates confidence. People feel included instead of replaced, and adoption becomes a shared journey rather than a top-down mandate.

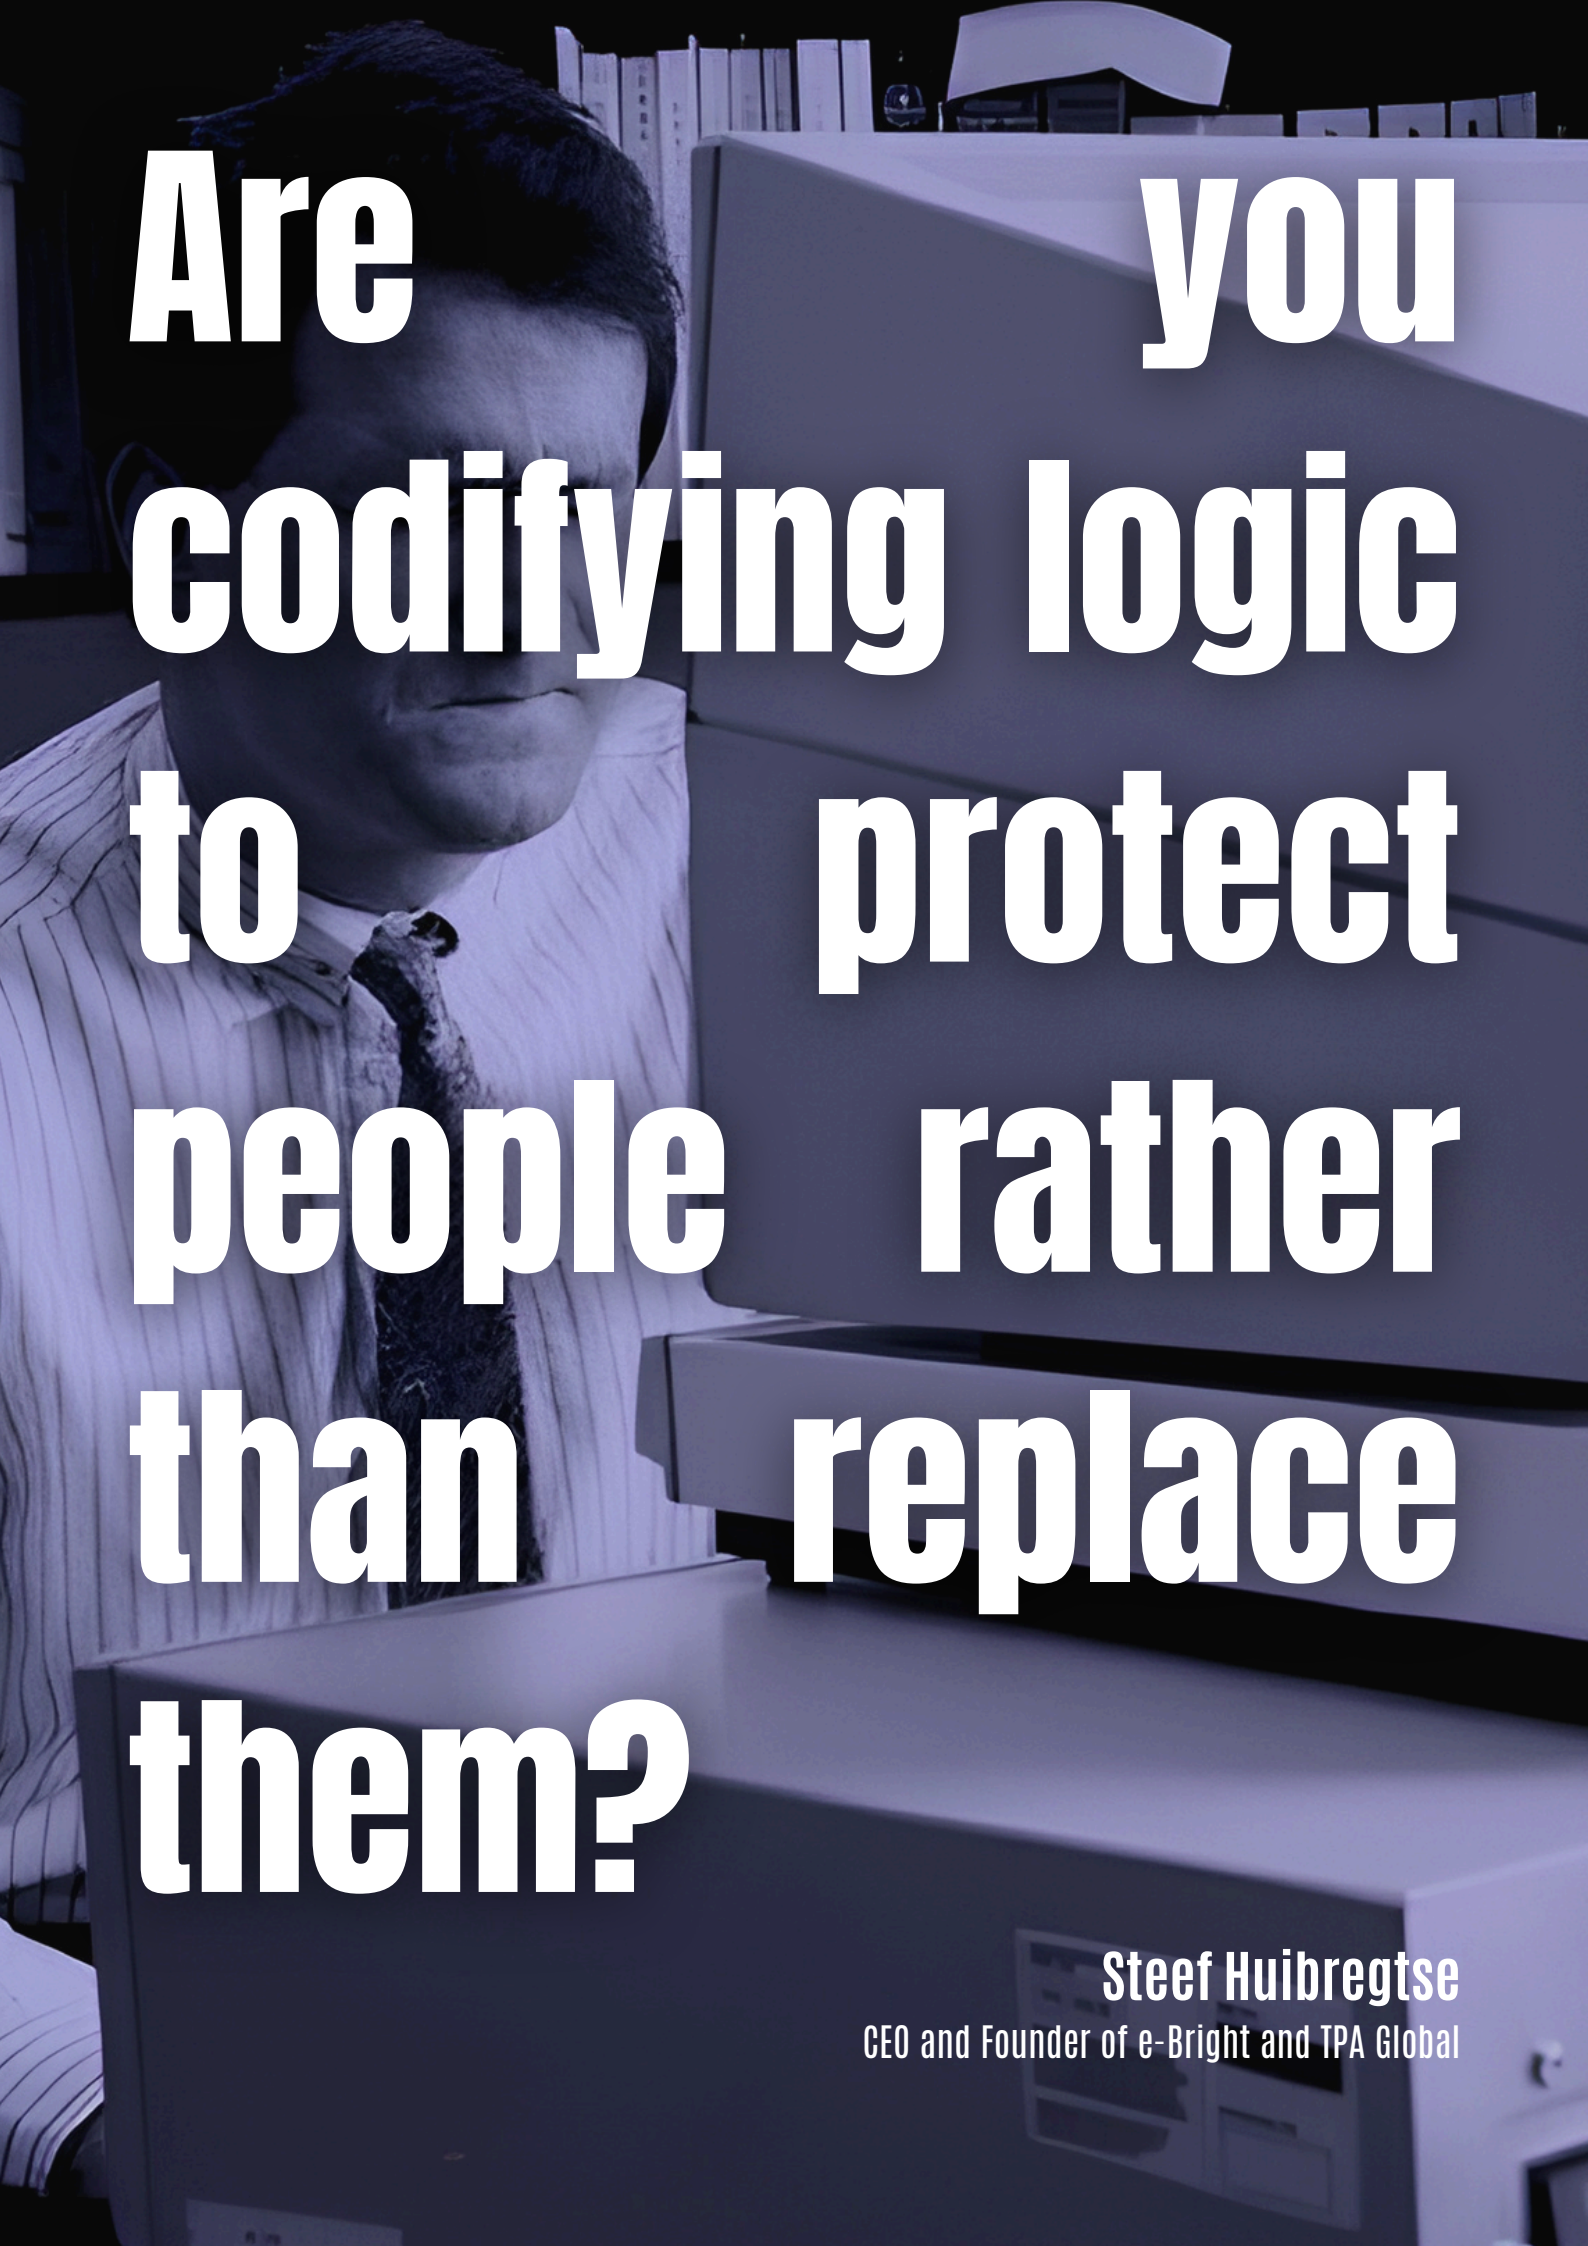

# Are you codifying logic to protect people rather than replace them?

**Steef Huibregtse**
CEO and Founder of e-Bright and TPA Global

# TTMAG Edition 4 Dec 2025

**e-Bright**
Rivium Quadrant 90
Capelle aan den IJssel 2909 LC
welcome@e-bright.com
www.e-bright.com

Published by e-Bright