

Foresight: The CPA Podcast

Season 6 Episode 2: Identifying and controlling the risks of Generative AI for CPAs

Neil Morrison: Welcome to Foresight: The CPA Podcast. I'm Neil Morrison. On episode five of last season, we asked *Will generative AI replace the CPA?* It was a great conversation, but it really just skimmed the surface of this topic. Bill Gates says 2023 gave us a glimpse of how AI will shape the future. The rise of generative AI is unlike any other artificial intelligence tool that CPAs have built into their practices, and Cathy Cobey says there's one word that describes the impact this new capability will have.

Cathy Cobey: Transformative.

Neil Morrison: Cathy Cobey is the EY global responsible AI co-lead. She oversees a global team that works on the ethical and control implications of artificial intelligence and autonomous systems.

Cathy Cobey: So I picked transformative because I don't think that there's been a lot of transformation in the way that we've done accounting for a couple of decades. I think the last time was when we really did the move to ERP systems like SAP and Oracle when we really moved from paper records to more automated records and general ledger systems. I think AI is going to be a similar type of transformation and it's going to instead really change the way that we do work on a day-to-day basis. It's going to allow us to really increase the use of automation in more cognitive areas, allowing us to do a lot more analysis, a lot more correlation.

And accountants are really going to have a very strategic impact in helping their businesses really get to information and insights that allow them to really help the business run better and to be able to react so much faster to things that may be changing or opportunities that may manifesting that might've taken even several days can now be reacted to hopefully much more quicker. So that's why I mentioned by transformative in that I think that this is going to be another big shift in the opportunities in how accountants do their day-to-day work.

Neil Morrison: Now, if that sounds like a rather optimistic read of how generative AI will impact the accounting profession, keep in mind that Cathy is an expert in managing the risks associated with artificial intelligence and she's been doing it for a while. So Cathy, what were the risks with the AI that was in place before generative AI came along?

Cathy Cobey: Well, that AI, which now is called either classical or traditional AI, was very data-science based. It really did end up usually being built within a data-science team, and it had a very singular function. It was ingesting a single type of data type. And so the risks really were, did we understand enough to even know where it was going wrong, how it might go wrong, how would we determine if it was going wrong? We didn't have a lot of real-time monitoring available over these technologies. And so it was a risky situation in that you may not be able to detect until a client called you and you don't want to be in that situation where a client is telling you that your technology has made an error that you weren't aware of. And so that's what has affected some of the slower adoption of these technologies.

Neil Morrison: I think when we spoke with you in the past, another risk that you mentioned was the risk of bias. It was so dependent on the data being unbiased, that if that data wasn't unbiased, it just exacerbated things and just amplified the bias.

Cathy Cobey: Well, a lot of the traditional classical AI was designed to be an optimization engine, which is that it tried to get to the most optimal result, and it determined what that optimum result was based on analyzing historical decisions and the indicators of why you might've made that decision. So in recruiting, for example, maybe you hired a lot more male candidates because historically males held more jobs in certain groups, and so it felt that, oh, a male is a better candidate for this position. And so it would erroneously make a judgment decision on what indicators are important based on historical decisions, not recognizing that our culture has changed or that was not an appropriate gender or age or ethnicity is not an appropriate determinant to make a recruiting decision. But be honest in healthcare might be. That might be a really informative in developing a treatment plan to know someone's ethnicity or their gender.

So it's not like bias is bad in all situations. And so that's why it's really difficult to deal with it. Because you have to really go use case by use case to say, what does bias or unfairness mean in this context? How might that have manifested in the data records in the past and how can I train the AI to ignore it or to work around it or to not consider it? And so there's been a lot of research to try and think about how to do that, but what is required is to start by recognizing it, right? Once you recognize it, you can build some solutions around it. But I think in some of the early AIs, there maybe wasn't as great an appreciation of how important that is.

Neil Morrison: Now, how are the risks from generative AI different from those previous, the classical AI that you're talking about?

Cathy Cobey: Well, in a couple of ways, I think that probably one of the biggest risks is that generative AI for the most part is available through a user interface. So where a lot of the traditional AI was in a scientific form in its outcomes that was being delivered to a data scientist who knew the data science behind it, and they could evaluate the outcomes and the risks with that knowledge. A lot of generative AI right now is available through just a prompt. You can just either-

Neil Morrison: Just a chat interface.

Cathy Cobey: ... Speak. It's a chat interface, right? And so the problem is that... Well, the good thing is that it's allowed it to be available to so many more people to start utilizing it, but that broader group does not really understand the underlying technology, and so they don't know what to look for. So some of the things that a lot of people don't recognize is that they'll ask a question of something like ChatGPT, not recognizing that it was trained on data that only goes up to 2022, for example.

So if you're looking for examples in 2023, it won't provide them. And so if you don't recognize that the training data set has that boundary condition and you don't build that into your outcomes, you may think you've got a very complete analysis to a question you ask when you don't. Or sometimes it doesn't recognize the fact that accuracy is more important than completion. And so if you ask for three examples and it only has two, it'll give you two actual examples and then make up the third. Whereas a human would say, "Look, it's better to tell them I can only give you two than to give you a third, that's not factual."

But AI doesn't always recognize that... It doesn't have that general context to recognize how to make that kind of judgment call.

Neil Morrison: Right. I wonder if we can just quickly zero in on the risks for CPAs when it comes to generative AI. Is there something in particular that's at risk for CPAs when they're using generative AI?

Cathy Cobey: Well, I think the biggest issue for CPAs is that we have a very high threshold for the accuracy of the information that we utilize, and we also use as inputs into other decisions that we make. So whether it's preparing financial statements or it's helping broader operational management in making a business decision or prioritization, we don't always appreciate just how high our bar of the accuracy and completeness of that information is. And so artificial intelligence is by its nature, a probability-based system in that it uses incomplete information to make judgements and its precision and confidence level can change very dramatically depending on the completeness and the accuracy and the relevance of the data that it gets from one decision to the next.

And so I think CPAs need to really start to appreciate how they can validate the reliance that they should place on the outcomes coming out of AI and how they can build some of that uncertainty into our processes where right now there is a very high expectation that when we provide information, it's got very high reliance and relevance and reliability, and yet AI doesn't always meet that threshold, but we don't want to then turn around and not use it. It's going to be a very valuable technology. So how do we balance those two sometimes competing priorities?

Neil Morrison: I guess the way to do that is with developing good governance or controls around the use of generative AI. Can you think of what are some examples of good governance around this technology?

Cathy Cobey: Well, I think probably one of the biggest areas of governance is to really ensure that there's, first of all, good transparency. So transparency when AI is being used, transparency in how that AI system is developing its decisions, what its boundary conditions are, what the precision and confidence levels. The more information that a user has, the more that they can understand what they might need to supplement it for, right? So that transparency has been part of the AI conversation, but I think it certainly has really been regenerated through generative AI.

Neil Morrison: And when you talk about transparency, is there a requirement for transparency with clients about this, making sure that they know at every stage where AI or generative AI has been used? Is that necessary?

Cathy Cobey: I think that a lot of clients are now starting to build that into their feedback or contractual expectations. And so I think you just look in broad society, I think there's been a message sent to a lot of the technology providers to say, "I want to know, right? I want to know if I'm dealing with an AI based decision or if it is using some more traditional technologies or a human." And so I think there's just been over and over again as we've had different use cases come out that there's a pretty high level of transparency. Once that transparency is there, be honest right now, we're still a very trusting society.

So as much as we have conversations about needing to mitigate the risks and manage these risks, a lot of us are still using ChatGPT. We're still using a lot of the facial recognition to open our phones and to

manage the security to a lot of the applications on our phones. And we're using transcription services more and more on productivity applications and meeting sites. And so we may not appreciate just how much we're actually already using a lot of these technologies. And so the more that there's transparency, I think it'll also increase a lot more of the trust where we'll recognize that there's actually a lot more that we've already built some experience and trust around than we're probably recognize.

Neil Morrison: But that also makes us feel like the transparency component. There's a point where that becomes almost impossible to maintain because as you said, there's so many little bits and pieces of AI that are just going into our normal day-to-day practice. You wouldn't want to give a client just a laundry list of all the different components of AI that you used. Over time, you can imagine that that starts to dilute the usefulness of it.

Cathy Cobey: It's true, but I do think that also is where some of the risk assessment processes come in place. And so a lot of times we've all been coming to these AI systems with one size fits all. It's all risky, where be honest, there's actually a wide spectrum of risk as you look from one AI application to the next. And so what we do also need to do is get a better job of understanding the fundamental inherent risks to some of these AI use cases. And there should be many use cases where we just as a general society, as an accounting group, as an organization, business senior leadership group, optimize those use cases like the very low risk. We should be trying to utilize them as much as possible and we should be giving access to them to our broad employee base and customer base as much as we can, so then we can then focus instead the governance and controls on those more medium to higher risk.

Neil Morrison: Right. That makes sense. I'm wondering, thinking of how the accounting profession is regulated, it's highly regulated obviously. How has it been responding from a regulatory standpoint? How has it been responding to the rise of generative AI?

Cathy Cobey: So I think that the most regulated part that accountants work within is in twofold. One is in our auditing of financial statements or if you're in the corporation, your preparation of financial statements that is governed by the security exchange commissions in each of the various countries and their regulators. And in that case, be honest, they've been very inquisitive in trying to understand where auditors might be using AI as part of audit procedures, where we're starting to see AI being used as part of preparing the financial statements. And right now it's still pretty limited, and I think it is because of such the high expectation of the reliability of that information and financial statements that it has somewhat made the entire accounting profession and finance departments more conservative in utilizing these technologies. But I think that generative AI now being available, for example through Office 365 and large general ledger ERP systems, it's going to provide a lot more of opportunities to finance departments that I think it's going to be very hard to ignore.

I think that it's going to be quite compelling for the efficiency gains, the quality gains, the ability to get much more faster insights that it's going to tip the balance where getting comfortable that you can rely on the information, knowing how far you can rely on it and what you might have to supplement it with will finally, I think, hit that balance where it is worthwhile to actually start really in a big way utilizing AI in these different applications. And then once more and more organizations are using it and can demonstrate to

the regulators that we do have a good handle on the risks, we do know how reliable it is. We do have supplemental control practices around it. I think everyone will get comfortable.

It's just it's new and there just hasn't been enough organizations to be able to demonstrate to the regulators yet that it can be used in a reliable way that doesn't impair the ability for us to rely on financial information. But I do think right now so far, we're seeing AI used more for management decision-making, so that is less regulated and it is really management taking on more of the risk if they make the wrong business decision. And so I think that a lot of organizations and accountants are getting some experience to AI through that, and they're going to be able to start applying that into the more highly regulated financial reporting area.

Neil Morrison: So it sounds like the key here is going to be having good governance and good controls around its use, but I'm wondering, we've seen how quickly AI can evolve in the past. What is it? It's just over a year, right? Since generative AI came out?

Cathy Cobey: Yeah.

Neil Morrison: How difficult or even is it possible for governance controls to keep up with the speed with which things are changing? Since it's moving so fast, are there general principles that people that accountants, organizations, businesses should keep in mind to just make sure that however it evolves, they've got a handle over the governance of it? Are there general principles that should be always directing things?

Cathy Cobey: Well, I think the good governance practices are certainly by starting with having an AI policy that really outlines, first of all, what is your definition of AI and what is your expectations if any group within your organization wants to start utilizing AI as to what are the principles and objectives that they need to meet? What are the risk assessment, boundary conditions and triage? And then based on that risk assessment, what then are the control practices that you need to have in place? And when you do have those medium and higher risk, I think there's a pretty good understanding as to some of the leading practices, like having an independent model validation team that can test it if it's a very dynamic model that is going to change or could potentially change quite rapidly once in production, having automated real-time monitoring, not only of the technical performance of the AI model, but also its behavior.

So I do think that there's quite a lot of understanding now about what those leading practices can be. The challenge right now is that as we come up with new AI classes like generative AI, it's how to define each of those for this new class of AI. So I think we understand what to do, it's now how to do it. And so whereas a lot of organizations that were leading in developing the more classical traditional AI, they had built a lot of those control functions. Then when generative AI became much more prevalent, now they're all needing to look back at it and say, "Okay, how do we have to modify it for generative AI? Do we have all the right testers that understand generative AI risks and functionality and how to test it before we put into production? Do we have the automated monitoring for generative AI risks?"

And so that's where I think a lot of organizations are right now grappling to do the how. But what they're doing is they're starting with the lower risk categories of uses because they don't need as much control

practices for those and using them as a stepping stone to understand the technology more, get used to it, and then hopefully they're building their control practices as they build out the more riskier use cases using these new technologies.

Neil Morrison: So that's how to stay on the right side of that risk reward ratio, taking advantage of it, not stifling innovation, while also not getting yourself into dangerous territory. It's almost like an iterative process that you're describing.

Cathy Cobey: Yeah. Some organizations had responded to ChatGPT and some of the other generative AI applications by prohibiting their use completely. So putting in firewall filters, not allowing off corporate devices to even go to those sites, which might've been the right approach, the very beginning stage when you didn't have a policy, you hadn't a thought, but by now, I would think most organizations are starting to relax that because they do want as an organization to start being able to use those lower risk use cases, right? If you prohibit it for too long, you lose that opportunity as the low risk use cases than trying to transform into medium and high risk cases. You don't want to just jump right to the medium and high. You want to have that opportunity to use the technology, get used to it and build alongside of it. And so I do think most organizations are hopefully now starting to think about how to relax those prohibitions and start to use the technology more.

Neil Morrison: When we started this conversation, you sounded quite optimistic about generative AI and about its power to really transform in a good way transform the accounting profession. As it continues to evolve this rapidly, are you optimistic that we will be able to stay tilted towards the reward and manage the risks, or are there things that keep you up at night about it?

Cathy Cobey: I would say I'm probably more realistic than optimistic. And by that I mean that I think it's coming, whether we're comfortable with it or not. I do think that its benefits are so significant that for competitive reasons, for efficiency reasons, we're going to need to utilize this technology. It's one of those general purpose technologies. There are some people that go off-grid and don't use electricity, but the majority of the society loves what electricity gives to us. And I think AI will eventually become that type of general purpose technology that it just needs to be used because it is just so efficient and productive. But I am realistic that we're going to have some failures along the way, and some of them are going to be quite big impactful failures, but we're going to work around them, right? If you think about automobiles, and the first automobiles didn't have great brake systems.

They didn't have street signs. We didn't have seat belts. But over time, as we wanted to go longer and faster in these vehicles, we had to build safety mechanisms to allow us. A lot of people say that we have brakes so we can slow down, but actually someone else told me, "No, I've got brakes, so I feel comfortable going faster." And that's somewhat where I have some optimism about AI is that the reason I focus so much on governance and controls is because I'm actually really excited about the technology and I want people to feel comfortable to adopt it and accelerate its use. But the only way to do that is for them to trust them. So that's why we need risk and control mitigation messages.

Neil Morrison: That's great. What a great place to end the conversation. Thank you so much, Cathy Cobey. I've really enjoyed it.

Cathy Cobey: Well, thank you, Neil. It's always great to talk to you.

Neil Morrison: Cathy Cobey is the EY Global responsible AI co-lead. On our next episode, we are going to be looking at AI adoption by small and medium enterprises, and we'll be speaking with Jean-Sébastien Charest. He's the chief information officer for the Business Development Bank of Canada. And, in our discussion he talks about how important it is for businesses to begin adopting AI. Canada is way behind other OECD countries in this regard. However, he adds an interesting caveat.

Jean-Sébastien Charest: All types of AI can help businesses. And this is where I'm going to put a bit of a caveat, right? AI, it's not a solution looking for a problem. Of course, AI can help if it's aligned with the objectives and the goals of the company. So, it's really about thinking about the business goals for the next year or two and where are the most pressing needs, and then how can digital technologies address them.

Neil Morrison: That's Jean-Sébastien Charest, the CIO of the Business Development Bank of Canada, speaking on our next episode.

And that's it for this episode of Foresight: The CPA Podcast. If you like what you heard, please give us a five-star rating or review wherever you get your podcasts and share it through your networks. Foresight is produced for CPA Canada by PodCraft Productions, and please note the views expressed by our guests are theirs alone and do not necessarily reflect the views of CPA Canada. Thanks so much for listening. I'm Neil Morrison.

Disclaimer: The views and opinions expressed in this podcast are those of the guest and do not necessarily reflect that of CPA Canada.