

# Foresight: The CPA Podcast

## Season 6 Episode 5: Meet the tax advisor's newest partner: AI

**Neil Morrison:** Welcome to Foresight: The CPA Podcast. I'm Neil Morrison. One of the things we wanted to highlight this season is case studies, concrete examples of how AI can be, or more importantly, how it is being used to support CPAs. And on our last episode, we looked at how AI can analyze a company's alignment with various sustainability reporting standards. This AI can also compare that company's performance against other comparable companies. And this targets a relatively new reporting requirement for CPAs. Yes, yes, sustainability has been developing for more than a decade and it's now codified in regulations and standards, but sustainability is still new compared to something like, say, tax. Tax advisory sits at the core of the profession. And so when AI expands into this territory, it's inserting itself into the heart of what CPAs do. And it's happening quickly. So quickly, in fact, it might make some tax professionals wonder if AI is really here to support them, or is it here to take over?

**Benjamin Alarie:** If somebody feels that nervousness, it's right to register that nervousness and notice that nervousness and be mindful of that nervousness.

**Neil Morrison:** Benjamin Alarie is the CEO of Blue J, an AI company based in Toronto. One of their products, Ask Blue J, offers tax research in natural conversational language. Ben also teaches tax law at the University of Toronto.

**Benjamin Alarie:** But what I would also say is this technology is really a boon for professionals who provide tax advice. It puts a floor under them. It gives them new abilities to see things, to notice things, to understand things at a deeper level. And the underlying complexity of tax is going to continue to increase to keep pace with our technological ability to understand and interpret and apply it in new circumstances. And so I'm not particularly worried as somebody who's very much in the thick of it and soon maybe there won't be a role for me here in providing tax advisory services, but I think that's the wrong way to connect the dots, at least for the foreseeable future. Perhaps in some far away time that will be the case, but I think for now, everything that I'm seeing points to tax advisors being supported and helped and enabled with this technology.

**Neil Morrison:** Let's get into your particular platform because we want to use it as a case study.

**Benjamin Alarie:** Sure.

**Neil Morrison:** Can you describe what part of tax advisory is your AI able to perform?

**Benjamin Alarie:** So Ask Blue J is a platform that really it's a research platform, so it's incredibly simple as a user interface, Neil. You log in and it's a text box. In the text box it says, "just ask," and it invites you to ask any tax research question you may have.

**Neil Morrison:** What's an example of a tax research question that say a CPA would put in that would be good for it?

**Benjamin Alarie:** Oh, virtually anything. But something that's not super technical, but often gets asked is something like, suppose you're at a small or medium-sized practice and you have a client who comes in. Maybe it's a child of one of your clients, and your client says, "Can you help out my son? He's been playing a lot of online poker, and he's won hundreds of thousands of dollars over the past few years. We want to know what we should be doing with this. He's got great records, but he hasn't been filing and we want to know what's the right way to do this."

You can ask Ask Blue J and just report very high level, a taxpayer is in his 20s, has been playing poker online for the past several years and has winnings of X thousands of dollars over those years. Should this be reported as taxable income and under what circumstances should it be reported as taxable income? And click enter and the system is going to go read... First of all, it's going to find a whole bunch of stuff relevant to that, and then it will synthesize all of that stuff.

It'll take the facts as you've given them, remix it with all of the additional information, and then go and do research based on the case law, based on technical interpretations from the CRA, based on some commentary and produce an answer that addresses the question and gives you several paragraph long answer, kind of a short memo outlining what the answer is and really importantly will point you to the sources that it used in order to generate that answer, so that you as the professional can click through to those source materials, read them, examine them, satisfy yourself that the answer that it's giving is correct. And the really cool thing is you can continue to carry on a conversation. It's very similar in terms of its interface to something like ChatGPT, where you ask an initial question, it's going to go away, think about it, produce an answer. The difference with Ask Blue J versus something like ChatGPT is it's tuned specifically for tax law. So it's going to be drawing on all of the relevant materials in order to produce that answer. It'll give you links to those materials. You can read the full text of those materials from within Ask Blue J. But like ChatGPT, you can carry on and ask follow-up questions, refine it, ask it to turn it into an email if you're going to cut and paste that into an email and send it on to your client, or you can turn it into a memo. You can ask follow-ups. You can say, "Can you clarify this part of the answer? Clarify this? What if this were different?" And you'll get additional analysis. And so it's a very natural way to conduct tax research. It's a way to save time and accelerate that first 80% of the research process to get you all the materials, get you a candidate answer. And then you as the professional need to...

**Neil Morrison:** Go validate it.

**Benjamin Alarie:** Dig in, validate, make sure that... Because these generative AI systems can make mistakes, it's important to make sure that they haven't missed something or mischaracterized it. But the cool thing about them is they often will make mistakes that... When they do make mistakes, and they're getting quite good. Like our internal benchmarking on Ask Blue J is it's producing clean and accurate answers close to 90% of the time.

And in the other 10% of the cases, often somebody who is familiar with the area will say, "Oh, well that doesn't quite seem right," and they'll be able to identify that and flag that, because the kinds of errors that it makes are not necessarily the same kinds of mistakes that human experts would make in the area. And so you still need the professional judgment to be a sophisticated consumer of this kind of tech, but it really is a huge time saver even for the most sophisticated.

**Neil Morrison:** And to anybody who has used ChatGPT, it sounds like it would be very familiar to them. It sounds like the interaction is basically the same as the interaction that you have with ChatGPT. How is it different in terms of what it's accessing or the information that it has or how it's trained than ChatGPT?

**Benjamin Alarie:** Yeah, I think there are a number of differences. The biggest one, of course, is that this system is specifically dedicated and focuses on providing tax analysis and tax research for tax experts. And so so much of the material that Ask Blue J is drawing upon is not available on the open web, so this is proprietary content, not necessarily exclusively available to Blue J. For example, we're subscribers to the proactive disclosure by the CRA. So the CRA makes available to publishers, including Blue J, its own materials that are available to be republished by publishers. And these are all the technical interpretations, all of the things that you can access through other content like Thomson Reuters or Wolters Kluwer. They're also like Blue J subscribers to this feed of information from the CRA. But this is not available on the open web, so it wouldn't be available to ChatGPT as training materials.

**Neil Morrison:** Or to Google for that matter.

**Benjamin Alarie:** Or to Google or to anyone that is not a subscriber to these materials. And so that's one big category. Another big category is Tax Notes. So Blue J has an arrangement with Tax Notes, which would probably be really well-known to your listeners who are tax people. They're publishers of Tax Notes Federal, Tax Notes State, Tax Notes International.

Tax Notes International has a lot of content that relates to Canadian taxation contributed by Canadians, and all of the Tax Notes material going back to the back catalog for the past 10 years and now on a go forward basis is available to Ask Blue J. And so that commentary is available to inform the answers that Ask Blue J is creating. And that's unique to Blue J.

**Neil Morrison:** I don't fully know how to ask this question, but is there something that is not within the scope of what this AI can do when it comes to tax advisory?

**Benjamin Alarie:** One of the things that generative AI for now is not great at is math, Neil. So if you ask a very complex question involving a multi-stage calculation, I would every time double check the math on that and make sure that the math is done properly. Large language models are tuned on tokens, and so their training data are basically words and strings of words in context. And through a whole bunch of algorithmic jujitsu, we produce these large language models that have billions of parameters that are tuned in order to produce the output from a given input. It amounts to prediction, the next token. They're not as good at math as systems that are specifically designed to do that analytical work. So I would trust a spreadsheet for calculations before I would trust a large language model. So that's a part that I would definitely want to double check.

**Neil Morrison:** Who is at the end of the day liable for this? I'm an accountant. I'm a CPA. I use your system. I get what I think is pretty good tax advice. I've gone and checked the sources that it provided. Turns out that those sources were just five of 15. And if it had read the whole 15, it would've gotten a different answer. Whatever the case is, it led me in the wrong way. I gave the

wrong advice. It's a problem for my client. Who is liable in that case? Is Blue J liable? Am I the CPA still liable? Is it shared?

**Benjamin Alarie:** At the moment, the answer, and this may sound self-serving, but ultimately it's the professional who's signing off on the advice who's liable for the advice that's given. It's always been this way. This is why we have regulated professions with mandatory professional insurance requirements. Lawyers are regulated in this way. Lawyers are responsible for the advice that they're giving. Accountants are. Engineers are. The engineer looking at the calculations provided by maybe an engineering technologist on the staff, and the engineering technologist has run all the calculations maybe for the kind of load that a bridge should be rated for and should be able to bear. Ultimately, it's that engineer who puts that PN stamp on those drawings and vouches for those calculations who's taking on responsibility for that professional advice. And of course, they're implicitly expecting the technologists who are doing the calculations to provide the right calculations to them, but it's their name, it's their stamp, it's their reputation on the line for that advice. It's the same way. It's like a book, right? You buy the book, you read the book, and you do your best analysis given what you've read in the book. But ultimately, you're the one responsible for the advice to the client.

**Neil Morrison:** Right. I want to look at the security of the information that's provided to Blue J. We know that the stuff you're putting into ChatGPT is later used for training of ChatGPT itself. So there's proprietary information. Tax is very sensitive information that's flowing through the system, and you're trusting it to be anonymized, but that's part of the whole system. That's part of what you're giving up. Is that a reason for people to be concerned if they're going in and giving all sorts of details that there's a security risk there, there's a privacy risk?

**Benjamin Alarie:** Yeah. This is a huge source of interest for people who are evaluating Blue J and looking to subscribe, and then ultimately subscribing. We have so many protections in place, Neil. We are SOC 2 compliant as a company, and so we have all the appropriate data privacy, data encryption safeguards there. We don't use the queries that users are running directly to train Blue J. We do ask for feedback. And if users provide feedback voluntarily, we will use that feedback that they're providing in order to improve the system and train the system.

Most of our larger clients opt out of any use of their data for training our systems. Some don't though. Interestingly, some of our larger clients have good controls on their side, and so they say, "Our users are not going to be providing sensitive information into your system, but what we'd like to do is understand what are the research questions that people are asking through your system?" Unless they opt in, we can't see their questions. So if they opt in, we can see the questions. We can see the interactions that their folks are having with the system. They're under strict instructions not to include sensitive information in those queries. And then they're asking us for guidance on how do we ask better questions? How do we use your system better to gain an edge over other firms that are looking to adopt this kind of technology? How do we become better question askers of a system like Blue J? And so there are a number of firms like that that are looking for our feedback and our industry best practices on how to ask questions.

And so we will go in subsequent. We'll have an onboarding period of a month or two, and then we'll do an analysis. And then we'll come back and do a lunch and learn or a training session

on these are the best practices, “Here are some examples and we’ll show you how. This is what somebody asked. Here’s a better way of approaching this research question, and here’s the difference in the output that you get from the system. You can see this way of addressing this is much more effective.” And so there’s some of that happening. But by and large, we’re dedicated to being transparent about how we’re using information. Folks have choices about how they want to interact with Blue J. And anyone who’s concerned about data privacy can opt out of any of those sharing things. And so it’s really something that people, if they want it, we can offer it. And if they’re nervous about that, then they’re under no obligation whatsoever and we’re not able to see their questions either on our side.

**Neil Morrison:** This is my closing question here. Can you imagine a time when clients no longer need CPAs? Instead, they interact directly with AI. I’m wondering about say myself. I’m running a small business. Is there a time where I don’t go to my accountant who uses Blue J, but I go to Blue J?

**Benjamin Alarie:** That’s a good question. I think a more likely scenario would be, do I just go to the tax administration directly? Will this technology become so powerful that governments say, “We will observe. You need to give us the appropriate data permissions to your banking information, to whatever information we require, and we will just assess your tax liability directly?” I think that’s the most likely end state. Because the same thing that...

So in this example where if you’re a small business owner, you don’t go to a CPA, you just go to a service provider like Blue J, and then Blue J presumably, I guess, interacts with the tax authorities on your behalf, I can imagine the true end state being tax authorities just go directly to small businesses, just go directly to taxpayers as individuals and just say... And it might not even be annual filing at that point. It’ll just be like we just settle up every day. It’s just like it’s an exact withholding system that’s calibrated to whatever the facts are that are being observed and adjusted in real time. And I think that’s the ultimate end state. Now, with that end state in mind, the question is, okay, how far away are we from that? And I think we can all agree we’re quite far away from that. I think so many things have to go right for us to get to that end state, including politically, economically, socially. I don’t know. We always overestimate how quickly things can change in the short run and then underestimate how much they can change in the medium to long term. So I think if that’s an imaginable end state, and I think it is, it would require so much visibility into the finances of businesses and individuals. You’d need almost complete transparency. I think there will be people who are very worried about the concentration of power in the state if we were to move in that direction. And so I would expect there to be a lot of political resistance to that just for reasons linked to the potential prospect of a totalitarian state, for example, and the abuse of power. People will paint dark dystopic pictures of what that kind of world is. And then techno-optimists will say, “Well, look how efficient that would be. We could really deliver benefits to those who need them. We could collect taxes in real time from those who have the ability to pay. We could calibrate the fairness of the system. It would be totally transparent, totally game free in its ideal form.”

**Neil Morrison:** I think all of this is fascinating. It's going in a fascinating direction, and it just shows the nature of the conversation we had at the end here, it just shows the scope of where it can potentially go as it develops. Completely fascinating. I really appreciate you taking all this time to talk with me.

**Benjamin Alarie:** My pleasure. Thanks for having me.

**Neil Morrison:** Benjamin Alarie is the CEO of Blue J, an AI company based in Toronto. On our next episode of Foresight we speak with Simon Dermaker. Simon is an associate professor in accounting at HEC Montreal. He's created a course on AI that is very popular with students, in fact it's so popular, they're expanding the number of classes that offer it. The course is called "Audit and Big Data." And I said it is a course on AI which is mostly true, but really it is about the the stuff that feeds AI. The stream of data that makes AI run. Simon thinks CPAs of the future are going to have a central role in auditing the quality of that data.

**Simon Dermaker:** Data is the fuel of AI, as it is for data analytics, and most of the aspects related to the fourth industrial revolution have data at the heart of it. The adage "garbage in, garbage out" has become more important than ever, and being able to assure that the data input in models, algorithmic models, machine learning models, is quality data. Ensuring this quality data will have an enormous effect in how we are able to give authority to the outputs given by those technological systems. And once we're able to have a good confidence in the data that's input in our different systems, auditing the actual systems will just be a natural extension of the auditor's work. Looking at the algorithms set up, looking at the models put in place, looking at the smart contracts that are being brought to different processes of organizations.

**Neil Morrison:** That's Simon Dermarker speaking on our next episode. But 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. It really helps us. Foresight is produced for CPA Canada by Pod-Craft 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.