

The Citi Institute Podcast

Title: Ronit and Prag

Guests:

Prag Sharma:	Global Head, Artificial Intelligence Centre of Excellence (AI CoE)
Ronit Ghose:	Global Head, Future of Finance, Citi Institute

Ronit Ghose (00:00):

Welcome to the special episode of the Citi Institute podcast series.

Wrap up Quote 1 - Dr. Prag Sharma

What is Agentic AI for your organisation? How are you able to identify? Do you have an inventory that can capture Agentic AI systems and associated information? From a use case perspective, everything looks like an Agentic AI use case. So you need to be sure that the Agentic AI approach is the right approach for the use case, uh, use case at hand.

Ronit Ghose 30):

We are gonna dive into Agentic AI, the Do It For Me economy, with our special guest, Dr. Prag Sharma. Prag is head of our AI Centre of Excellence, which he runs out of Dublin. Prag, thanks so much for joining us today.

Dr. Prag Sharma (00:50):

Thanks for having me, Ronit.

Ronit Ghose (00:51):

Let's define the terms. Uh, we've all heard about generative AI for the last couple of years. I think we know what AI means, but let's hear it from the expert. What does it mean? What is AI? And specifically, define Agentic AI for us?

Dr. Prag Sharma (01:05):

Yeah, in one sense, that's a very difficult question to start with. As you, as you know yourself, artificial intelligence itself is difficult to define because it is an amalgamation of many, many different, um, areas of study, research and applications and so on. And then we came, come up with generative AI, which is an interesting subsection of AI, where your, where the language models are specifically focused on generating new content. And now we have another term called Agentic AI, which again, is in the same ballpark, but slightly different. So let's talk about Agentic AI. I think the simplest way to define it is, uh, to look at large language models or generative AI, uh, artificial intelligence techniques that generate content, plus take action. I think it's that amalgamation, like you said, it's, it's taking what the language models produce, plus giving them the capability or the ability to take action on those outcomes. And hence the name agents or Agentic AI.

Ronit Ghose (02:06):

So generative AI is AI that shows us things, words, pictures, videos, Agentic AI is AI that puts it into action, that does it for us. So where does it make sense? Where does Agentic AI make sense in finance and economy? And where does it not make sense? Let's dig into that, Prag.

Dr. Prag Sharma (02:26):

Yeah. I mean, clearly for large financial institutions and many, many organisations across, uh, the globe, you know, consuming information is just one aspect of the, of the work we do. So large language models, were very good that, uh, taking that information, consolidating it together, and giving it to an end user, a person to consume it. But actually, uh, we know that consuming is just one part of a, of a process or an activity we do. The other is to take action on that information that a human has just been presented to. So with that respect, Agentic AI opens up even more use cases.

The Citi Institute Podcast

So if you can think of large language models as, uh, doing a particular task for you, which is summarising content, you know, translating content, uh, content, extracting right amounts of information from huge corpus of data, the Agentic AI system is now taking away that specific task and completing it from an end-to-end perspective. So it is moving it, moving from a task to a process-oriented approach. And in large organisations like ourselves, we have many, many processes across the globe. We have many functions that have processes, and these processes have many steps. So Agentic AI is ideal to replicate some of these processes that have multiple steps, um, that you can then put these tools between those steps and look at the final outcome that a human could probably then review.

Ronit Ghose (03:53):

So anything that's repetitive, that's a workflow-based process, and obviously all, call it, white-collar organisations such as ourselves or most of our clients. We have lots of processes that we usually do with a combination of people and software, that over time is gonna get changed. Now, where's the boundary of the borderline? Obviously the, the science is gonna change dramatically in the next couple of years, but where is Agentic AI not useful? Because let's, let's draw some sort of boundaries as to say, you know, for the foreseeable future, these are areas in finance or the economy where Agentic AI is probably not the best use.

Dr. Prag Sharma (04:33):

Yeah. So before I dive into what's not useful, you made a very good point, and you used the word workflow. To many people listening, they've always been dealing with workflows and automating workflows. So we're back to the question then, what is Agentic AI bringing to the fore that's different from workflows, you know?

Ronit Ghose (04:49):

Yeah.

Dr. Prag Sharma (04:50):

We have workflows. You can write code today that combines a number of steps together and you produce an outcome. You could take a, take a complex reporting task, where you have to take data from multiple sources. You may want to compile that data, you may want to produce charts. You may need another person to then review that data and do further analysis, and then produce a final outcome. You may then have another team that would trace back that outcome to the original source, for example. So all of that can be done today without Agentic AI or even with just a simple AI techniques.

(05:20):

But what Agentic AI is bringing is you can now do this with unstructured data. That's the wow factor. You know, we don't always just deal in numbers. If that was the case, this would be a solved problem years or decades ago. But actually, it's the unstructured data, the text, the audio, the video you can now incorporate into a process, into a task that you can then join together, uh, to make a final product. That's really, wow. I, I've said it easily, but you know, I've said it, "Oh yeah. Audio, video, text." But actually, most of our data is audio, video, text, as you know. And that means many more processes have now, uh, come into the flow.

(05:58):

Okay, so what? That's great. That's already an improvement. You're now taking multimodal as we call it, uh, in the jargon, putting it together using the Agentic AI systems, maybe in a workflow pattern. So in a strict, uh, step-by-step process. Actually, what Agentic AI is also bringing to the flow, uh, to the fold, is that it does not need to be a step-by-step process. It does not need to be a workflow that can benefit from this technology. And what I mean by that, is we can make, we can use Agentic AI's, uh, thinking and reasoning capabilities to actually, uh, let the models or let the Agentic AI system make the decision on what step to take next.

(06:40):

So it could be in a research scenario where our analysts are, uh, researching a particular topic. You may not want to describe those steps. You may want the model to decide next steps based on the output. And I think this is the real wow factor that hasn't been thought about before. So now we know really how Agentic AI is distinguishing itself from other workflow-based approaches, or in fact, even generative AI approaches. 'Cause it's not just about taking those next steps, it's about actioning them using many, many different types of tools, such as connecting to your SharePoint, connecting to a Python environment to execute code, connecting to a particular library to draw

charts and graphs, connecting to a third-party system to get that data and execute. So it's really becoming a very, very powerful tool.

[\(07:30\)](#):

Now, where is it not that, uh, good of a use case? I think that answer is also in my previous comments, where there is high risk of you getting the exact answer. As we know, generative AI is the [inaudible 00:07:12], it generates content based on what it's seen before. That generated content does not always have to be the truth. It can be something that sounds like the truth or is very close to the ground truth, but it isn't. In scenarios where that's important, think about reg reporting where you're going to a regulator to produce a very specific, um, report with very specific numbers and figures that have to be exact. In that scenario, you need very robust controls in place, governance in place, humans in the loop, uh, to ensure that these, uh, uh, that these reports are accurate. It doesn't mean you shouldn't be using, uh, generative AI or Agentic AI in that space, you just need to have the right [inaudible 00:07:57].

Ronit Ghose [\(08:28\)](#):

Now, when we talk about Agentic AI, we're really talking about what's coming in the future. Some of, um, some of our clients, some of our partners already trialling, maybe putting it into production. But walk us through, if you were advising senior management of ABC company or the board, what kind of questions would you be suggesting to them to think about when it comes to putting Agentic AI into practise? What are the, I think you've told me Prag before, you have these five magic questions, right? Do you wanna share with the audience what they are?

Dr. Prag Sharma [\(09:04\)](#):

Well, yeah, I think, um, exactly, everybody today is looking at Agentic AI and how to implement it. One thing to note is that the pace of change is so quick that early in the, in the year, if you were having this conversation, Ronit, we'd say, "Oh, maybe in six months time you'll be utilising Agentic AI systems." Actually, that's not true, we are utilising Agentic AI systems today, and their capabilities are increasing at a speed, which means they can do more and more things, uh, within a short period of time. So I think from a, uh, from a senior perspective, the key things to focus on are, one, if you're using Agentic AI, are you able to identify Agentic AI applications within your organisation? Because if you don't identify them robustly, accurately consistently, then you are unlikely to be able to put the right risk and controls across those.

[\(09:54\)](#):

So have a clear definition, think about what that means for your organisation. What is Agentic AI for your organisation? How are you able to identify? Do you have an inventory that can capture Agentic AI systems and associated information? That's step one as you look to, uh, develop and deploy this. From a use case perspective, everything looks like an Agentic AI use case. And as we were just discussing earlier on, um, you know, you can use workflows, you can use robotic process automation, you can use many other technologies, uh, to fulfil the needs of a particular use case. So you need to be sure that the Agentic AI approach is the right approach for the use case, uh, use case at hand. That also makes sense, obviously. But I mean, nowadays there is such a focus on Agentic AI or generative AI that you need to make sure that the use case, um, uh, is the right one.

[\(10:44\)](#):

Also, it's not about flashy use cases, it's about the right use case. Uh, potentially starting with maybe low risk use cases to test out the technology and then work your way towards more high risk use cases. So many of the use cases today are about efficiency gains and Agentic AI supercharging that concept of efficiency gains and hyper personalization, which means in doing something better and faster, specifically for the use case at hand. And that use case could even be for a particular person within a group, as opposed to a particular organisation or a particular department. That's how much of a lead way you now have with these kind of tools. So use case again, um, is becoming, uh, a key focus area, and it's slightly different from generative AI, because use cases now are not tasks, they are end-to-end processes. And you will need to deploy this technology in multiple different ways within a process, um, to get the best results.

[\(11:38\)](#):

So I think, um, as a, as a senior leader, you should be asking, "Is this the right approach and is this the only approach? And I need to combine this with other approaches," such as what we mentioned, traditional AI approaches potentially, or RPA type tools or other audio video tools. "Uh, and look to combine that with our existing tool set that we have." So that's a key consideration for the... "Without the right use case, we're in

The Citi Institute Podcast

trouble." I think that's the, that's the bit the senior management should be asking. And then of course, governance and risk. You are not going to go far in a, um, in a highly regulated environment without the right, uh, ownership who is ultimately responsible for the deployment of the use cases. Um, do we have the right governance and patrol mechanisms in place? And I think many organisations are taking their current processes and, uh, expanding them to include Agentic AI processes as well, which has worked really well for us, and which will work really well for many other organisations.

(12:35):

Even if you're not a highly regulated, um, industry, you should still be conscious of the Agentic AI governance [inaudible 00:12:14] risks. And you should be conscious, because it can have an adverse effect without you really getting to know it at a later stage. So I think that's another very key focus area um, uh, from our perspective. So we've talked about use cases, we've talked about, you know, defining the technology clearly and having an inventory of it. We have talked about governance and risk, uh, related activities. And finally, you know, third party vendors, whom should you be working along with? Who, whether you should be building or buying the underlying large language models or small language models? What is that mix of resources, technologies, third party vendors or otherwise? And then also finally, you know, the technology stack, the architecture components, the hardware, the software, all of that stuff also needs to be considered. So this is a nice holistic view.

Ronit Ghose (13:33):

Let's, let's come back to the technology and infrastructure and data piece in a second and dive into that more. But let's start with the topic that I think all leaders and all of us have top of mind. How will Agentic AI impact people? How will it affect the workforce? If I'm, if I'm a leader, how am I thinking about training my, upskilling my employees or maybe re-skilling my employees? And if I'm an employee, uh, and we are looking to carry on working rather than retire to the beach together, Prag, um, what do we need to, not you, your obviously helping run our AI Centre of Excellence, but the rest of us, what do we need to do to stay employable, relevant and all those good things in the future?

Dr. Prag Sharma (14:23):

Yeah, I think whenever we talk about, you know, the impact of this technology on people, I always find it to be overhyped, uh, and not in the right way. I think there's, there's this discussion around, you know, um, will AI take over our jobs? There's discussion around will we, will we all be sitting on a beach, because it is going to get too good? I think the future is, uh, is not going to be like that at all. For us, it is about upskilling our employees to use this technology. Ultimately, it is a tool, it is a very sophisticated tool, and this tool continues to increase what it can do well. Even so we need humans, uh, in the loop. We need skilled people who have a deep knowledge of their processes, of their technology, uh, of their business, to combine that with generative ai.

(15:10):

So the idea is this is a symbiotic relationship where we should expect generative AI to be doing, uh, more of the heavy lifting, but humans are required, uh, to clearly you know, supplement that know-how, review it, and then produce results. But I think the follow-on question, Ronit, there could be, you know, does that mean you need less humans, you know? Or you need more humans? Really, that's really where the question is, you know. But that's a different question in the sense that depends on the people, depends on the organisation and how you implement it.

Ronit Ghose (15:39):

Sure.

Before we get onto the question of quantum or the number of people doing a particular role and task, can I just go back to my previous question and make you sort of drill further in?

Dr. Prag Sharma (15:50):

Yes.

Ronit Ghose (15:51):

If I were, if I am in the workforce today, I'm in my twenties, thirties, forties, or even dare I say, in my fifties.

The Citi Institute Podcast

What, what should we be as an employee or as a manager of that employee, what are we thinking about in terms of skills? Is it, um, listening to this podcast, spending time hanging out with you, reading your blogs? I don't know what.

How do we, uh, going to your lectures in Dublin, how do we, how do we stay, stay on top of all of this?

I was obviously, reading our GPS reports.

Dr. Prag Sharma (16:18):

Correct. I think I would start by saying read Ronit's GPS reports. But, uh, the point is we should always stay up to, up to date on technologies anyway. I think people in the, in the industry or otherwise that are always interested in innovation and new technologies, we are up to speed, up to speed with this. Agentic AI is again, accelerating that trend as well. So while I'm immersed in this space, I have to read new papers that come out on a daily basis to be interested. So what should individuals be doing? Look, uh, the first thing is, this is not a tech passing phase that, you know, you can sit out, relax. "I think it's gonna, uh, it's the hype is gonna die," and you'll be back to the way you were working before, I think that is unlikely to be the case. What is more likely is that you'll have to get used to using this technology. And the differentiator is how well you use this technology.

(17:06):

To do that, I would highly recommend, uh, staying in touch by reading, uh, things like the GPS reports or other deep dives that really go behind the scenes to tell you what's happening in the industry, what's happening with the technology, what's happening with the use cases. So be familiar with it. And then the way to distinguish yourself is you have lots of know-how that just doesn't vanish 'cause we now have a large-language model. What you've learned over time and your expertise is now even more valuable. It's that combination of what you know from your work experience or life experience that you can use, uh, for your day job when working with a large-language model that is gonna distinguish yourself. So one, read up on it as many times as you like. And you don't need to be a technologist for this at all.

(17:48):

Do not think you should be writing Python code. And if you don't know how to write Python code, you're out of the loop, that is not the case. Get stuck in, read from your perspective, from a risk perspective. There are numerous opportunities to know how to mitigate risk. Controls, governance is a very hot topic, especially in the, in the business circles. And there's a lot of know-how to be gained from that. Finding the development perspective, there's a, there's a lot going on. So pick the topic you like, get stuck in and, and keep up to speed. That's one thing. The second thing is, use this technology both at home and in the work. Take every opportunity to get hands-on with the technology. By that I mean, if you're a developer, get stuck into development. But if you're an end user, use these large language models for your use case because that is the only way you're going to find out what it actually cannot do. And the key here, is find out what the limits of this technology are, so you can exploit it to the max. I think that would be, that would be good advice and one I follow myself.

Ronit Ghose (18:48):

Now, let's switch gears a little bit. We've talked about this topic from the perspective of ourselves as employees or managers. Let's talk about it from the point of view of the client and the product. How does Agentic AI in particular, you could maybe more put in the context of broader AI, how does this change products? So in financial services or related areas, what new capabilities, uh, what new features are gonna come around for clients? How does the life of a client change because of Agentic AI?

Dr. Prag Sharma (19:21):

Yeah, I think one, clients, our clients are very interested in using these technologies themselves. So they often come to us to find out how we are using it. So hands-on experience, again, is key as a first step to really changing your own life yourself or changing, you know, or working with your industry partners to get the most out of this, uh, out of this technology. But today, the, the focus very much is it's a fast evolving field. There are lots of risks. We don't know what all those risks are. So we want to be careful in, in how we can offer this to our clients or how we can advise our clients on using this technology. So like we were briefly saying, first, the, uh, right now we are very much interested in seeing what the benefits of this, uh, technology are, are from a efficiency perspective. Can we make the lives of our clients more efficient by doing certain processes that may have taken longer, um, to do them in a shorter period of time? That's really the first thing, and that's how everyone is looking at this.

The Citi Institute Podcast

(20:16):

The second phase quickly that matures into with the Agentic AI, is what can we offer to clients that they can directly interact with them, buying data, processes, services when we are confident that the risks associated with technology has been mitigated, um, uh, when we know what's happening under the hood, when we can work with our second line of defence and so on, to, to get all the right controls, um, in place? So I think that's where the, that's where this mix is. As the technology evolves, there are going to be more and more use cases. Some of those use cases are going to focus on self-service, self, self-help, hyper-personalization. Again, the term meaning that we will be able to develop products and services that meet the requirements or, uh, of an individual or multiple individuals. So that's where this technology is going. Um, and that's where we see, uh, we will continue to expand using specifically Agentic AI, which has opened up even more than the generative AI use case that we have developed with that.

Ronit Ghose (21:15):

So Agentic AI hopefully be all about making our clients' lives simpler and easier and more efficient, taking the friction away from their lives. And particularly when it comes to financial services, most of our clients, and most of us don't want to spend huge amounts of time filling our forms, redoing KYC, looking at product renewals. And if Agentic AI can make some of those workflows simpler, that should hopefully, give client satisfaction.

Dr. Prag Sharma (21:46):

This is the promise. I think this, the promise is exactly what you point out.

Ronit Ghose (21:51):

Yeah.

Dr. Prag Sharma (21:52):

The promise is simplification, uh, efficiency gains, um, potentially and then, you know, hyper-personalization.

(21:58):

I think those three or four terms capture at a very high level, uh, what we hope to get.

Ronit Ghose (22:04):

And delivering the promise, obviously, and this is the hard bit, requires a lot of changes around or augmentation and building on technology, infrastructure, and particularly for us as FIs, data and data strategy. So let's kind of put all of that together and discuss that. Um, what are the changes in, whether it's technology stacks, whether it's technology governance, data strategy that we need as large organisations that have been around, in our case, for a couple of centuries to deliver on that promise of a simpler, faster, better life for our clients?

Dr. Prag Sharma (22:42):

Yeah, I think this is one of those unique technologies that affect every part of the business simultaneously. So if you look at a large organisation ourselves and we have discussed most of this, there's the five pieces around people, which we need to upscale and we need to make sure, uh, they can use this, uh, technology. We have processes and that we need to review, um, and align with Agentic AI use cases. People, process, technology, infrastructure, governance, and finally data. So these verticals continue to be, uh, uh, simultaneously impacted by this technology.

And I think we have to focus on each one of them and integrate them across the board, so that we can get the results that we want from this technology. But yes, um, fundamental changes to the process, the processes, it's not about doing business as usual and trying to slot in the capabilities of this technology. The successful, uh, organisations are those that will look at the capabilities of the technology and look to put their processes around that. So we are flipping, um, this narrative somewhat from processes first and what technologies can support, uh, to a phase where these technologies are getting so good that we want to see how we can slot our processes to get them started on this technology.

The Citi Institute Podcast

Ronit Ghose (24:01):

Now looking ahead, um, Prag you talked about how Agentic AI and AI and emerging technologies in general are moving at the speed of, speed of light, you know, just moving really fast.

Um, when we were working on the report, uh, with your help in December, January, the GPS and Agentic AI, we were probably thinking 2025 even maybe 2026 is proof of concept, science project land.

And maybe we land this in 2027, 28, and yet we're seeing, you're seeing, I'm seeing, uh, enterprise size companies, enterprise clients putting this into practise already in Q1, Q2, 2025. Um, what's the timeline, if you had to think about the sort of realistic timeline of adoption of Agentic AI across the economy, across, uh, the business sector?

How, when, when should we start seeing, you know, this is becoming a, a real, real thing rather than just a experiment here and a proof of concept there?

Dr. Prag Sharma (25:03):

Yeah. No, I think this year, um, already experiments have begun on Agentic AI systems, including, um, ourselves.

We want to make sure that we align our governance and risk functions to take the most out of this technology.

That's the factor that will maybe slow, uh, the adoption down from today to maybe a month or two months from

now as we, as we build those, uh, things out, robustly extend our existing, um, risk and control frameworks to

ensure that we have, um, uh, mitigated other risks, uh, related to Agentic AI. But yes, by the end of this year, we should have Agentic AI systems. It depends what we are imagining when we talk about this, Ronit.

(25:42):

We started the conversation with what is the definition, you know.

And many people will walk away with very different concepts of what Agentic AI is.

Today, if you have a large language model that connects to a tool, such as a Python environment to execute some code or write some SQL code to get data out of, uh, uh, your, uh, database, that's actually Agentic AI in one sense,

that's happening today. But actually what we are thinking when we talk about this, is multiple large language

models, say 10, 20 large language models all working together to produce end results, connecting to tens of

different data sets, uh, producing outputs that are reviewed by, uh, several different humans and then end

finished product that really might take months. We're imagining that a click off a button might do that.

(26:28):

That's sometime, that still has a way to go. But different flavours of Agentic AI, one or a few models working with a few tools are very much a reality today. It's about that complexity in this. How much autonomy, is a word we

haven't used yet, how much autonomy we should be giving to these Agentic AI models? How much can we let

them run free? What oversight is needed? That's still a little bit down the line, but that doesn't mean you can get

the benefits of this technology today with simpler models, systems, connectivity, uh, with tools. And that's exactly what we're doing.

Ronit Ghose (27:05):

So before we wrap up, Prag, give me one last focus on misconceptions, or as you said, people often have very different ideas of what AI and Agentic AI actually means. In your conversations with colleagues or with the external

world out there, what's the biggest misconception people in the business world or the banking world have about

Agentic AI? What are they, what are they getting probably the most wrong?

Dr. Prag Sharma (27:34):

Yeah. Well, I think the most... Uh, one thing we need to be careful of is it's not, you know, uh, a solution for every problem that you have in your organisation. It is a solution for some problems.

(27:47):

There are many other things that need to be fixed to get the most out of this technology. It's not that this technology will fix those problems so that you can move on. I think this is a, is a concept we need to understand.

As an example, you know, you need to make sure you have good quality data that you can ensure that your large-scale models have access to.

(28:05):

Large-language models are unlikely to produce good quality data for you to use unless you have, you know,

specially trained them to do something like that. But the idea is if you provide contradictory data to a large-

language model, it's unlikely to resolve that for you, uh, without you providing enough background information. So

The Citi Institute Podcast

I think the key thing people are getting wrong is a jack-of-all-trades, this is going to solve problems that I haven't been able to solve before.

(28:30):

That's not quite the way to look at it.

Ronit Ghose (28:32):

Yep. Data, data, data, it's a bit like focusing on diet when you're trying to, you know, keep your weight under control.

Dr. Prag Sharma (28:39):

Yes.

Ronit Ghose (28:40):

Just talking about myself here, Prag, not you.

(28:43):

Um, now just sort of wrapping up, where can listeners, where can the audience learn more about this topic? Um, do they follow you on LinkedIn? Listen to your lectures in Dublin? What's, uh, are there certain podcasts? I mean, obviously this podcast, uh, our GPS series?

What's the best way to stay on top of this very fast-moving area?

Dr. Prag Sharma (29:02):

I think certainly follow me on, uh, on LinkedIn. I post material there are, no problem. But I think, uh, like you mentioned, without being biased, I think the Citi GPS reports are a great starting point. They go into real detail on some of the topics that are of interest, but there are other publications from others as well that people might be interested in. Instead of telling them exactly what to read, what I should say, is keep reading. And more importantly, keep doing, keep doing.

(29:26):

I think this is the bit that, this is where that gap will emerge between people, um, who are interested and people who are interested and are, you know, getting stuck in. So, um, and I finished previous podcasts like this, we should really, uh, finish by saying, get stuck in.

Ronit Ghose (29:42):

Get stuck in. Just do it. Keep playing. Stay curious.

Dr. Prag Sharma (29:48):

Yes.

Ronit Ghose (29:48):

Dr. Prag Sharma, it's always a pleasure to chat to you. Thank you so much for joining us on the Citi Institute podcast.

Dr. Prag Sharma (29:53):

Thank you for having me.

- Disclaimer

The speakers' views are their own and may not necessarily reflect the views of Citi or any of its affiliates. All opinions are subject to change without notice. Neither the information provided, nor any opinion expressed constitutes a solicitation for the purchase or sale of any security. The expressions of opinion are not intended to be a forecast of future events or a guarantee of future results.

(30:47):