S05E01: Are universities ready for the age of AI?

HOST

Richard Westcott (University of Cambridge)

GUEST SPEAKERS

Jonathan Grant (Bennett School of Public Policy and KAUST), Jean-François Bonnefon (IAST and TSE), François Poinas (IAST and TSE)

Richard Westcott

Hello and welcome to Crossing Channels, I'm Richard Westcott. Are Universities Ready for the Age of AI? That's the subject of the latest in our podcast collaboration between the Bennett School of Public Policy at the University of Cambridge and the Institute for Advanced Study in Toulouse at the Toulouse School of Economics.

Since our last series, both institutions have marked important milestones: the Bennett Institute has become the Bennett School of Public Policy, and in Toulouse, IAST has strengthened its ties with TSE through the creation of a new department for social and behavioral sciences, building on IAST's tradition of groundbreaking interdisciplinary research.

As ever, we're going to draw on the interdisciplinary strengths of both institutions to explore a complex issue — how artificial intelligence is reshaping higher education, whether universities are prepared for these changes, and what they must do to adapt for the future.

So, to explore these issues today, we have Jonathan Grant - formerly affiliated with the Bennett School and now Vice President for Strategic Initiatives at KAUST in Saudi Arabia. Jonathan, start us off, what does your research focus on?

Jonathan Grant
So Richard, it's a good question, I reflect on that quite a lot. Over the past 25 years, think there are two golden threads that come together, which is one around purpose and the other one around impact. I've applied that originally to research, so research impact and research purpose, and then subsequently to higher education. I think when you reflect on both purpose and impact, you start to think about the future of higher education. And I suspect that's why you've invited me to join this podcast because obviously AI, HE come together and AI on research and AI on education come together as well.

Richard Westcott
Fantastic. Joining us from the IAST and TSE, have Jean-François Bonnefon and François Poinas. Jean-Francois, remind us of your main research interests.

Jean-François Bonnefon
Hi everyone, I'm Jean-François. I'm a psychologist by training and I do behavioral science for the ethics of AI. What it means is that I run experiments to understand the kind of bad thing we do with AI and the bad things that AI may do to us.

Richard Westcott
That sounds ominous. And Francois, what's your research interest?

François Poinas
Hello everyone, I'm Francois, I'm an economist in TSE. I'm working on economics of education and labor in few words. So I'm just looking at how people make their schooling decisions and how these decisions affect the labor market outcomes.

Richard: Now, AI is already reshaping many parts of our lives and universities are not immune. So generative tools are changing how students learn, how academics conduct research, and even how institutions are managed. In the UK, for example, a 2025 survey by the Higher Education Policy Institute found that 92 % of undergraduates now use AI, and nearly nine in ten have used generative AI tools in assessments. But only around a third have received formal training from their university. So what does this mean for higher education? How should universities adapt their teaching, research and social role in the age of AI?

Jonathan, in your work on the future of universities, you've argued that universities must rethink their role and resilience. So looking at AI today, where do you see the biggest opportunities and risks for universities?

Jonathan: So Richard, I always reflect on that quote by Roy Amaro about that we overestimate the short-term impacts of technology and underestimate the long-term impacts of technology. And when we apply that to AI, it's quite scary because if you think the short-term impacts of AI on our societies in the last two years have been profound, what does that mean the long-term impacts are going to be?

And if you then apply that to higher education in universities, I do, without sounding alarmist, think that there is an existential threat to universities. That's the sort of glass half empty part of me. And then the glass half full says that if we look at the long history of universities over millennia, universities have successfully evolved from both social but also technological shocks, over that time. So I guess I'm an optimist. I do think the universities are going to adapt. But if they don't engage with this debate and they don't change their fundamental operating model and be that around education, be that around research, then I think there is a real risk that we won't see the university as an entity as we see it today and other bits of society will be delivering on that research and delivering on that education.

Richard: So I'm intrigued how that looks, you know, an existential threat. I mean, what kind of timeframe are you talking about? What's going to be the kind of killer blow?

Jonathan: Well, so I think you can quite quickly see why do we need to deliver education through the entity of a university as currently constructed. You could have an online platform that is providing AI tuition, delivering online lectures. Students can acquire knowledge quite easily nowadays. I use that word acquire very specifically.

And the sort of skills, if you like, of universities from an education point of view is to then provide that critical appraisal. You know, how do we assess whether this is true or not true? How do we apply this new knowledge to this context? But do you need to do that for a university? Maybe a big global brand, a BMW decides, well, actually, we don't need graduates anymore. We're going to hire young people at 18, we're going to put them through this AI driven knowledge acquisition process, and then when they come and work in our company, we will then coach them through the critical appraisal bit. So that's just one example.

Richard: Jean-Francois, in your research, you examined fears about AI across 20 countries and different domains of application, for example, in medicine and law and management and so on. And you found that concerns vary depending on the context. So where did people think that AI could be a positive force? And where were the areas that worried them where they thought AI couldn't replace humans?

Jean-François: Well, typically people don't want AI in very high-stake domains like justice. And they are much more willing to let AI be journalists, for example. But the key thing is that there are quite a lot of variation. And one problem here is that people sometimes use bad heuristics in terms of what they trust the AI for, because they project some kind of effect on the personality vibes of the AI on its competence. So to say it simply, if you think that typically professors sound

confident and career advisors sound friendly, then you might think that an AI that has a confident vibe might be able to give you good medical advice. And you think that an AI that has friendly vibes might be able to give you a good career advice. But these things are totally uncorrelated when it comes to AI. And I guess a corollary of that is in the same spirit of people applying the wrong heuristics when judging what the AI can and cannot do is that typically with humans, you extrapolate that someone will not be able to do something difficult if the person can do something simpler. The problem is it completely breaks down when it comes to AI, which has this kind of jagged frontier of performance, which means that sometimes it fails at something trivial, but that doesn't mean it won't be able to do something much more complicated if you give it to it. So what we observe is that there's a range of anthropomorphic biases that people apply that make them make bad predictions about what they should use AI for and what they should not use AI for.

Richard: Francois, your research has explored how education systems shape early career outcomes. How well equipped are universities to provide the kinds of skills graduates will need in an AI driven economy?

François: So the role of universities is basically to train students and to train graduates to be able to develop AI tools, but also to use AI tools in their everyday life and their work. So basically they have this objective, in fact, to train students on this. And they are in very good position to do that at universities because, they have, like Jonathan was saying, they have done that for centuries. So this is their job, in fact, to adapt to new technologies, and to teach to the students a knowledge that is useful in fact when the new technology comes out. So what is a challenge here is the speed at which AI tools are developing and basically so there is a risk of skill obsolescence and the fact that the skills that you give to your students might become obsolete at some point but in fact, universities are not just learning very precise skills. They have always learned their students how to learn, basically, and how to adapt to a changing environment. For the graduates, I think that, in fact, there are some risks and opportunities. You have on the risk side, basically, the risk that, well, you need less people trained because AI is going to replace some workers on some tasks. This is a job destruction aspect.

On the opportunity side in terms of job creation, when you have a new technology that arrives, then people who use this technology then increase their productivity, which creates in fact value creation, and then it has some positive effect in fact on the labor demand. So it has positive effects on obviously in fact on sectors in which we are using AI, but also in sectors in which we are not using AI, because you have this increase in fact in terms of growth.

So there is those these two effects, so risk and benefits, it's hard to tell which effect is going to be at stake here. But then at the end, you know, you have these two different effects, risk and opportunities for the graduates.

Richard: We're talking, Francois, about universities in isolation, really, aren't we? But do they need to not be working with the schools, even down to very, very young ages, so that when they arrive at university, there's a better understanding of AI?

François: Well yeah, that's a very good point in fact. So then if you want to develop AI skills, learning is cumulative, you know, so you learn skills based on the skills that you have learned

before, in fact. So the connection between primary, secondary and tertiary education is going to be key in this process. Yes, for sure.

Richard: Jean-Francois, you've studied how people respond to advice from an AI. So if universities begin using AI to support students, and we know that there's lots of stress on universities, financial stress, and basically when you provide support, it becomes very expensive. So it's quite tempting to bring in potentially support services that are AI driven, for example, by offering feedback or guidance, et cetera. So what does your research suggest about the conditions under which students will trust and benefit from AI advice and what challenges should universities be ready for?

Jean-François: Well, my work is usually oriented towards bad outcomes. So here I'm going to provide more warning than advice or guidance. But what we found, which is a little bit worrying, is that students and actually everyone might say, show a tendency to cut ethical corners when they are using AI in a way that they're not maybe quite realising.

So imagine that, for example, a student would ask an AI tutor about general guidance about how to write a given paper. There is nothing wrong with that. But then the AI might spontaneously volunteer the full paper without being prompted to do so.

Or maybe the AI is going to suggest an outline and then append a question like, would you like to see how the full introduction would look like? And the problem is here that psychologically it does not feel like actively cheating. When the AI is actually taking the initiative to offer all these things that you've not actively asked for. And so this makes it likely for people to just go with the flow, and accept help from the AI that they would not have dared to ask for, or that they would not have considered okay to ask for. So we have to be careful about how we use this AI tutor so that they stick to responsible requests from the students and do not go overdrive into actually pushing the students to cheat.

Richard: Jonathan, in The New Power University, you called for new approaches to teaching and learning aligned with civic and social missions. So could you explain a bit more about that? How should universities rethink teaching and assessment to stay true to those missions?

Jonathan: Well, so I think Richard there number of dimensions to this. First of all, I think there's this sort of delivery dimension, let's call it that. So for quite a while now, university sort of commentators have talked about ideas of the flipped classroom, quote unquote, where students are expected to acquire, I use that word again, their knowledge before they enter the classroom.

And then the classroom is actually where you do the critical appraisal elements in small groups. I think AI accelerates that type of thinking. And in a slightly ironic sense, given the location of the Bennett School, you could actually, you're reinventing the Oxbridge model of tutorials through AI. And if universities get this right, then actually they will be releasing resource from the ‘lecture’ in a very classic sense, because that will be delivered in an automated way and developing that critical appraisal skills in a much more discursive human to human type approach. So I think that’s potentially quite exciting. I think the second thing then it gets to what Jean-Francois was just talking about to a degree is this sort of assessment and appraisal type

element of education. How do we differentiate sort of first class students from 2:1 students or whatever grading system we want to use?

And I just I teach a module for a master's course at the University of Pennsylvania around the New Power University. And this year for their final assessment, I told them basically use AI to answer this question and then critically appraise the response coming out of AI. And I was marking this assignment this weekend. And what I learned by marking it was actually the quality of the AI response was driven by the quality of the prompts that the students were putting in. And the students who were putting in what you might call more thoughtful prompts were giving better critical appraisal analysis of the outcome from the AI. There were only 10 papers, so small numbers. But I found that really useful. So from a sort of educational point of view.

We probably need to spend a bit of effort thinking about how, as users of AI, we can get the best out of it by creating the right responses as well.

Richard: That's interesting, a new way of assessing people is how well they can use the AI, how sensibly and cleverly they can use it. Francois, if we think about labour market outcomes, do you think AI can help reduce skill mismatches between what the universities teach and what employers need? And we keep hearing about these mismatches. Or could it risk widening that gap?

François: Well, in fact, universities are used in fact to assess what the labour market is needing in terms of skills. I mean, so they have done that for centuries because they are well connected in fact to the labour market. There is no risk, in fact, that this time will be different, you know, with the way universities are going to deal with the labor market needs.

Even though here, so we have to be careful because of course the technology is going to is evolving very fast, but also the economic model is not, this is not a technology on which there is a mature economic model. So it's very difficult to say now, in fact, which AI tool is going to win the game, I would say, you know, in 10 years. And basically, so there is here a risk on which, you can, if you specialise yourself, you know, in adopting a certain tool or certain model, in fact, to teach your students with an AI technology, then there's a case that in 10 years from now, or even five years from now, either the company is going to collapse or the model, in fact, is not going to be effective anymore. And so this is something that is something on which we have to be careful as universities to make sure that not to make some choices that are going to be difficult to reverse and to engage students into wrong solutions.

Richard: Well, bring Jonathan in on this. So what kind of policy and leadership decisions do you think are going to be most important for the universities to make sure that they stay ahead of this? Like Francois says, they, you know, because it's so difficult to predict where this technology is going, how could they set themselves up to be able to adapt quickly in the right ways?

Jonathan: So can I just go slightly back there, Richard, and just comment on Francois, because there's a line of argument I haven't heard often coming from universities. So we are appropriately as a sector focused on skilling the next generation of workers. And as Francois

articulated, that demand from industry and the public sector changes every generation because societies progress.

Universities do that well. If you read the literature around AI and its economic impact, there is a scenario that in five, 10, 15 years time, we will have quite high levels of either unemployment or underemployment, projections running at 20%. It seems to me what we at universities need to be thinking about is how do we educate our students so they can be fulfilled and have purpose in their lives when they are working much less. The classic at least model in the sort of western democracies is you get purpose through work. We object to that quite rightly, but that's kind of the model. But if you're no longer getting purpose through work, how are you going to get purpose? And actually that's a shout out for the humanities who always feel slightly sort of nervous about this conversation. But actually I think there's a real case for universities investing in the humanities because we're going to need graduates who are not only job-ready but also sort of socially ready to be able to work in a different society and feel fulfilled, have that sense of purpose.

Richard: I wonder, Francois, do you have a view on that, a view on how you get that sort of human fulfillment from life if jobs are going to start sort of disappearing off? Where do you get your sense of purpose, your sense of satisfaction, bearing in mind a career could be 40 years.

François: No, I think it's a very interesting point, in fact. And so this is a scenario that Jonathan is emphasizing is one of the scenarios. There are other scenarios on which, well, there will be more place, in fact, for labor, depending on if productivity effects of AI are going to create, in fact, value and increase the demand of labor on other sectors on which AI is not going to be used. And there is also another effect, in fact, that the economic literature is putting on this, it is a kind of enhancement effect. So the fact that you can have new tasks, new jobs, in fact, and new ways of producing wealth, basically, in a world where AI is going to enhance labor market work. And basically, so this is a sum of all these effects that is going to determine if there will be still jobs, in fact, and how many later on.

Richard: I'd also, I'd love to see something that enhances humanities as well as a guy who did English literature and got a lot out of it and really value it. So, Jean-Francois, we're talking about opportunities for universities here. So you've contributed to global debates on AI governance. Should universities position themselves as ethical leaders in AI adoption? And if so, how should they go about it? Because ethics is clearly a running theme through all AI now, isn't it?

Jean-François: True, well it would be so great if that would happen, if a university could take the lead here, because it's true that there are places where knowledge is created in a transparent way, they're full of people who take ethics very seriously, and you know they're relatively free from commercial interference, so they would be like ideal places to lead the debate.

But on the other hand, there's something a bit strange happening, which is that I'm seeing, well, significant proportion of academics who have seemingly knee-jerk, negative reactions about AI, generative AI in particular.

And sometimes this reaction do not seem fully informed or rational. So in the range of speculation, for example, it's not uncommon to hear that generative AI will prevent students from acquiring critical thinking skills.

And I don't think there is so far a single piece of evidence that would support that. But you commonly hear that as a complete red flag, deal breaker with using AI. Or you might hear that because AI is likely to hallucinate its sources, you definitely cannot use it for research. I mean, that was probably a problem two years ago.

But to be honest, by now, this problem is certainly not the kind of deal-breaker that it's sometimes announced to be. So I'm kind of worried, actually, that attitudes toward AI might be showing some signs of polarization within academia where what matters is where you stand about AI and not really why. So that would mean that the university might lose this incredible opportunity to be the place that leads the ethical debates if they cannot avoid the common problem just arguing for something without really investigating why.

Richard: It's interesting that something just two years old can be so out of date as well as that just shows how fast the technology is moving. What do you think universities could do then in terms of the leadership to bring staff up to date to make sure that you do take this opportunity and lead in a debate and perhaps other people who might have monetary interests at heart aren't taking over the ethics, for example?

Jean-François: Well, first it's true that companies, AI companies do contribute to the debate. They have red teaming and risk assessment and all those things. And they're increasingly encouraged to share the results of this ethical work with the public and with their competitors.

But as you can imagine, this kind of sales disclosure is going to be selective, right? And bound to some commercial timelines and pressure. So we just cannot rely on the AI company to win the ethical race for us. So we are going to have to, in universities, contribute to that debate. One of the problems, of course, is that those companies are miles ahead of us in terms of access to the models and the kind of experiments they can do with the models. I absolutely don't have access to frontier models by OpenAI to investigate what they can do, so I have to trust them. So the problem we have is that we don't have the access.

The rate at which things are changing is just insane. I computer scientists have sort of a culture that accommodates a fast pace of research. Other fields like the social sciences don't have a cycle that can accommodate a research object that changes every other month. So we're going to have, perhaps, to change the culture of the social sciences to align better to that of computer science if we are to study that kind of object. And the problem is that even then, even someone like me would spend his entire research time on studying these things. Well, I am overwhelmed. And so I cannot even imagine how my colleagues who were not studying that for a living might find the time to keep abreast of what is going on. So I'm sorry that I don't have really good solutions about this, but I just wanted to frame how difficult that problem is.

Richard: Okay, we always like to finish by looking ahead into the future because obviously it's so easy to predict the future, isn't it? We always, always get it right. But if you could highlight one priority for universities over the next five years in responding to AI, what would that priority

be? And also, what do you think the risk is if they don't do it? So let's start off perhaps with you, Francois.

François: Yeah, so I think that the best challenge in fact, yeah, so the most challenging part in fact for universities will be to train students, producing AI, using AI and to train them well. Meaning that in fact, and refers to what Jean-Francois was saying, so first it is a public role in fact to promote an ethical use of AI, to train well our students on what it is capable of, what it is not capable of and what are the limits in fact of AI and like biases and hallucinations. And the most important to my opinion in fact will be for universities to reflect about what skills I will need to teach to my students in the future.

So, if you take an example of coding, instance, then there is some debate about should we continue teaching coding to our students? One argument that can say, well, the machine is coding well, in fact, then so why do need to learn coding? And then on the other side, so then you have something like, well, but you need to learn your students to code for two reasons. First, because you need to be able to have some knowledge of coding in order to be able to assess possible hallucinations, for instance. And the second aspect is, well, basically, AI is able to code well into situations that are more and more complex. But at some point, so then you need to sort of train some people who are going to solve more complex problems than what possibly AI is going to do. And then so the key here will be to train well, in fact, students in coding in order to have the skills that will be complemented well, in fact, with AI. And the key here is to reflect, you know, about this complementarity between what AI is going to give you and what skills, in fact, you need to train your students and if we don't do that, so then there will be two consequences. First, we are going to miss this massive technological transformation, so AI is just a way in fact to increase, to create value, to increase wealth.

Our societies need to create wealth in order to face challenges that are facing in terms of global warming, in terms of spending. We need to have some spending, really need to defend spending, for instance. And basically, the risk is that, well, if we don't train well, you know, graduates, then the society as a whole is going to, you know, not to benefit from this, to this technological revolution that can just be the revolution that will permit, in fact, to increase prosperity.

Richard: Jonathan, you've articulated the risk very well earlier on in the podcast talking about an existential threat potentially to the universities. So what do you think their priorities should be in the next five years to go the right way and not the dangerous way?

Jonathan: Well, Richard, I think they actually need to fundamentally redesign how they deliver and assess education, because if they don't, others will do it and they'll do it at half the unit cost. Now, at one level, if you take a sort of national economy viewpoint, does it matter that universities aren't doing that and somebody else is doing that? And I think that does get back to my opening comments about purpose. I think we have had universities in societies across the world for millennia, and there's a reason for that because we need to have a protected space that is not driven by profit motives, not driven by politics, to educate the next generation and explore new ideas and new knowledge and drive progress. So we need universities, and it is absolutely essential in my view that leaders of universities actually stop taking an incremental

approach to the future and start taking a revolutionary approach because the downside of not doing that is that these activities will be delivered by other entities and that will be in the long run damaging for societies.

Richard: And Jean Francois?

Jean-François: Yes, it's true that one risk is that if we don't adapt to AI, we may lose some credit from the population that is currently trusting us. The fraction of the population with an understanding and trust in higher education and research and to them we might feel like dinosaurs if we don't move on. But one opportunity we have is that AI could actually be very useful in getting credit from the fraction of the population that is currently not listening to us and forgive the self-promotion, but I have this article with my team coming out soon where we showed that there are people you might call science outsiders, people, for example, who don't want to hear about climate science, they don't want to hear about vaccine science, they have no interest or trust in scientific research and they prefer the alternative facts.

But AI, generative AI, can be extremely good at reframing scientific knowledge in a package that these science outsiders find palatable. And suddenly, they engage with the science facts the same way they used to engage with the alternative facts. And that's half of the battle. So what I mean here is that, yes, we might lose some credit from the portion of population that used to be our consumers. But there is a great opportunity to actually use AI to reach out to people who have never listened to us so far. So on balance, you know, maybe we'll be okay.

Richard: Well on that note, that's all we've got time for on this episode. Thanks to Jonathan Grant from KAUST and Jean-François Bonnefon and François Poinas from the IAST and the Toulouse School of Economics. Thanks for listening to this episode of Crossing Channels. We'd love to hear what you think. Your reviews help us shape future episodes and make it easier for new listeners to find the show. And if you enjoyed this one, why not explore some of our other episodes too? And please join us next month for the next edition when we'll be talking about the health and wealth gap.