Bird & Bird

Regulating AI – Perspectives from Europe and China

Podcast transcript – Wilfred Ng (Partner, China), Shima Abbady (Associate, Netherlands) and Toby Bond (Partner, UK)

November 2023

Update: This conversation took place in October 2023 and is accurate as of January 2024.

Toby Bond:

I'm Toby Bond, a Partner in Bird & Bird's IP Group in London and I spend far too much time talking, thinking and writing about AI. Now, regulating AI is pretty high on the agenda at the moment. It's been getting significant airtime at the G7, the UK hosted the first a global AI safety summit in November 2023, and the legislative process for the EU AI Act which began in April 2021 reached an important political agreement during the final political trilogue from 6 December to 8 December 2023. Today we are going to be discussing two jurisdictions which are really leading the way in AI regulation — the EU and China. One of the pleasures of working at Bird & Bird is if you need an expert on any tech subject across Europe, the Middle East and Asia Pacific, it doesn't take long to find one. I'm really pleased to be joined by two of those experts. From Bird & Bird in Hong Kong, we've got Wilfred Ng. Wilfred, do you want to say a few words about yourself?

Wilfred Ng:

Thank you Toby. Hello everyone. My name is Wilfred Ng. I am a Partner in the Commercial Department of our Hong Kong practice. I practice transactional and regulatory matters for our technology clients in the region, and I'm really excited to be given the opportunity to exchange views and learn from my colleagues who are experts in the AI space.

Toby Bond:

Brilliant, thanks so much Wilfred. And joining Wilfred, from Bird & Bird in the Netherlands we've got Shima Abbady. Shima, do you want to introduce yourself?

Shima Abbady:

Thank you Toby. My name is Shima Abbady, I'm a lawyer in Bird & Bird in the Netherlands in the Commercial and Data Protection Practices. I have been at Bird & Bird for about seven years specialising in AI, data regulation and data protection mainly, and I'm doing both advisory and contentious work, and apart from that I also undertake interdisciplinary PhD research in the field of AI. I'm very excited that we have so many great experts in Bird & Bird and to be talking about this today.

Toby Bond:

Great, thanks so much Wilfred and Shima. And Shima, with you doing a PhD in AI Regulation, you probably spend even more time than me thinking about AI Regulation and Technology so it's great we're all together today. So, from my perspective as a lawyer based in London who is very into the AI space, I've seen a lot of articles discussing the EU AI Act. We've also seen some on the sort of Chinese regulation on generative AI, but the one thing I haven't really seen was people coming together to talk about the two in comparison and to actually see what each regulation is trying to achieve, as well what themes we see emerging between the two. Are there commonalities that we can learn from? That's why I thought that it would be great to bring you both together to talk about that, and I think maybe the way to get into it is to start by talking about what the context of both of these proposals

are. What's the background? Where have they come from? Shima, let's start with the EU, and the EU AI Act. Can you give us an idea of where has this come from? What has the EU been thinking about before and why has it got to the point where it's trying to regulate AI now?

Shima Abbady:

That is the best question to start with. It came from the desire of the EU to be the first to introduce horizontal AI regulation in the world. The EU realised in about 2007 approximately that AI was going to be this huge thing in a fairly short term, and they wanted to capitalise on that. Regulation is good for legal certainty and it's good for the protection of EU values and EU rights and therefore it is good for fostering trust in AI within the Union. Then you can leverage that trust to create innovation or encourage innovation, growth, you can also use it to try to be a standard setter for the rest of the world, just like we tried to do and maybe even accomplished with GDPR. So, overall, I think you can see it as an investment in AI from the EU. If you're looking at the wider legal framework, in terms of how it fits with a whole host other types of new digital regulation like the Digital Services Act (**DSA**), the proposal for a new AI Liability Directive, the Machinery Regulation, Product Liability Directive Update, etc., they're all meant to be pieces to this giant puzzle that is supposed to streamline the 'digital decade' — the decade that we're currently in, according to the EU. And the AI Act is mostly the puzzle piece that is supposed to provide the general product safety perspective for horizontal AI regulation, and as you mentioned earlier, a political agreement was finally reached on the AI Act on 8 December.

Toby Bond:

Yeah, interesting. Obviously, the EU wants to be first and move quickly but it seems the technology is moving pretty quickly as well. During the process we've had ChatGPT laws, concerning AI, so could you just tell us a bit about how generative AI came into the picture with the regulation?

Shima Abbady:

Of course! It's the thing that everybody is talking about. These EU regulatory processes for drafting new regulation can take a lot of time, especially in a completely new field such as AI which has never been regulated before, so the EU started undertaking this process years ago — the first draft proposal from the commission came out in 2021, more than two years ago, back when we didn't have ChatGPT. So when the ChatGPT craze started was exactly the moment when Parliament was working on their own position — the Council and the Commission had previously respectively already done so — Parliament decided, "*well we have to also include something for generative AI specifically"*. The parties subsequently spent a lot of time negotiating over how exactly to achieve this., There was a lot of disagreement about whether providers of general-purpose AI-systems/models (including foundation models, such as the one ChatGPT is based on) should be regulated. Ultimately, in the political agreement, the parties landed on regulating most of this group to a limited extent and regulating a very small part of that group more extensively — namely, the providers of very powerful models which are deemed to entail 'systemic risks', which will very likely include OpenAI as the provider of GPT4.

Toby Bond:

So although the EU wanted to be first, I think it's fair to say that China got there first on regulating generative AI, with the CAC's Interim Measures for the Management of Generative Artificial Intelligence Services (Interim Measures) coming into force on August 15, 2023. But it wasn't the first thing. Wilfred, what had been happening before the new Gen-AI regulations in China?

Wilfred Ng:

Because of the jurisdictional and cultural differences, I suppose it's fair to say that the starting point was quite different — but you're absolutely right, Toby, as it actually traces back to September 2021, when a definitive and cross-government departmental policy statement made a pledge to establish a regulatory framework for algorithms used in internet information services of a specific nature. Later that year, ministries led by the Cyberspace Administration of China (**CAC**) jointly published a regulation on algorithm-based online recommendation technologies in December 2021, which cover a wide range of services that can filter, recommend, and rank content for individual users, such as the usual typical functionalities that you would see in an app. China has also focused on other aspects of regulating technology, and in September 2022 announced regulations regarding "deep-synthesis" technology — if technologies can automatically generate audio, visual, and textual content, such as vary a piece of online content for example, they would fall under the ambit of this regulatory direction. So, although the Interim Measures for the Management of Generative Artificial Intelligence Services has been most talked about and most

often compared with other overarching AI regulations in various jurisdictions, it was actually preceded by other regulations on AI. Even more recently, in October 2023, the National Information Security Standardization Technical Committee (**TC260**) first published the draft Basic Security Requirements for Generative Artificial Intelligence Services. In addition to the various AI and algorithm-specific regulations, the overarching trio of the Personal Information Protection Law (**PIPL**), the Data Security Law (**DSL**) and the Cybersecurity Law (**CSL**) also form an important backdrop to the regulatory framework for generative AI.

Toby Bond:

Thanks Wilfred. Shima, I know that recommendation algorithms are something which has been one of the hot topics for the EU's way of looking at regulating AI, so it's interesting to see that this is definitely an issue which both jurisdictions are looking at.

Shima Abbady:

Yes, although it is worth mentioning that there is a slight difference with China because the main focus on recommender systems in EU regulation is really in the DSA, whereas it is not really a main topic for the AI Act. During the AI Act negotiations, there was also a proposal for Parliament to regulate recommender systems of very large online platforms as AI systems of high-risk, but in the political agreement, the negotiating parties agreed not to do this, as they are already regulated under the DSA.

Toby Bond:

That's probably a good segue into a second area which we wanted to cover. We've now talked about where this sort of regulation has come from and what was happening before. We should probably talk about what is actually being regulated and what each jurisdiction is trying to cover. Maybe if we just start with what types of AI systems are being classified in different ways. Shima, do you want to set out where the EU's thinking is around classified systems?

Shima Abbady:

Of course. For the EU, the definition of an AI system will be fairly broad — in the political agreement, the negotiating parties adopted a revised version of the OECD's definition, namely, a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that [can] influence physical or virtual environments, additionally noting that different AI systems vary in their levels of autonomy and adaptiveness after deployment. However, despite the broad definition of AI systems, only those that are considered to be 'high-risk' and those qualified as general-purpose AI models — especially the most powerful models which are deemed to entail 'systemic risks' — are going to be extensively regulated. Additionally, there is going to be another category of 'prohibited systems' which are going to be banned entirely. Among these include databases based on bulk scraping of facial images, systems for general social scoring, predictive policing, as well as systems that exploit vulnerable groups such as children or the disabled. There are also systems which are going to be regulated to a very limited extent, meaning that only transparency obligations will apply to them. These are systems intended to interact with people like chatbots, as well as systems producing deepfakes. So where it's really at, especially for us lawyers, is the high-risk and 'systemic risk' categories, because I think those are the categories that most of the regulation is going to apply to. And which systems qualify as high-risk systems and 'systemic risk' systems will be able to change over time, because the EU intends to take a flexible risk-based approach where they can designate systems or models as high-risk or 'systemic risk' on a case by case basis depending on the purpose of the system (and, in the case of 'systemic risk', depending on how powerful the model is). With regards to high-risk, the thing to look out for is whether you are operating a system that is likely to have fundamental rights implications for individuals, such as in the areas of recruitment, public benefits and law enforcement — these areas are likely to be regulated as a high-risk area. The same logic applies for systems likely to have potential safety implications (both on the individual level and societal level, such as critical infrastructure). But overall, only a minority of all AI systems/models which are used in the EU will be regulated as high-risk — and an even smaller minority as 'systemic risk' — so this is something to bear in mind, since a lot of people do not realise that most systems/models probably won't be regulated at all.

Toby Bond:

Yes, the EU AI Act isn't a regulation for every possible AI, it's for particular types of AI. Wilfred,

what about China? What sort of scope and categorisation are we seeing coming out there?

Wilfred Ng:

I think the interesting comparative angle here is that the Chinese measures are more focused on the generative nature of artificial intelligence. So, if you juxtaposed the two, you could see that China's is less of a classification system; both the CAC's Interim Measures and the draft TC260's Basic Safety Requirements only apply to the use of generative AI technology or services, with a similar definition of generative AI being strictly in relation to models and related technology that have the ability to generate text, images, audios, videos, or other content-specific output. Here we can already see from a legislative or policy perspective that the measures are targeting a very specific use of AI logic or AI application, as compared to what Shima has comprehensively explained on the EU front, which is more of a wide spectrum of application. That probably has to do with a lot of the existing PRC content-related regulations in other laws which are not specifically limited to the AI context. For example, if you provide an internet service, there are already existing content regulations that you would need to comply with. And as we get into other obligations, you would also see some parallels in the regulatory and the legislative intent.

Toby Bond:

Your comment there about it fitting into an existing content-related regulatory context is quite an important insight. One of the key things people need to understand about these regulations is who is actually going to be regulated, and which parties are going to have to comply with these regulations. Shima, from the EU perspective, who are the primary targets of regulation in terms of the actors in the AI supply chain?

Shima Abbady:

The primary targets are providers and deployers of regulated systems. Providers would be the parties that are either developing these systems and bringing these systems to market in their own name or the importers of such systems, whereas the deployers would be the parties that are using the systems for their own purposes. For instance, a bank which has a creditworthiness checking system with AI input would be considered a deployer of the system and therefore would likely be subject to regulation as the deployer of a high-risk AI-system. Lastly, as mentioned before, providers of general-purpose AI-systems, including foundation models, will also be regulated.

Toby Bond:

Wilfred, tell us about how that is in China — do we have the provider and the deployer, or is it more focused on one than the other?

Wilfred Ng:

Yeah, I think on this particular point it is quite similar, so the primary target is what is defined as a service provider, and it is very clear that this will catch both the developer of the AI platforms, as well as the organisations who are going to incorporate the AI technologies through APIs into their apps, for instance. Therefore, both deployers and developers will be within the scope of these measures.

Toby Bond:

What about the jurisdictional scope — how do you fall into these regulations? Where in the world would you need to be to get caught?

Wilfred Ng:

It will be very similar to the EU, and will come into play as long as you are providing generative AI services to the public inside China. Thus, there is an extra-territorial element if, for example, I am hosting and providing an AI service outside of China, but I am actively targeting the audience inside China — I would still be caught under the measures.

Shima Abbady:

Yeah, that does sound similar to the EU's scope, but in classic EU fashion, we're going one step further in terms of extraterritoriality. So not only are you caught by the EU Act by deploying these systems within the Union or by providing these systems by bringing them onto the Union market; you would be caught simply by *using* the output of the system (and Parliament even wanted to go as far as triggering regulation by the mere *intention to use* the output of the system) in the EU. Thus, it's very easy to fall within the scope of regulation.

Toby Bond:

The approach to territoriality is definitely a theme we've seen come up before. We have worked out where this regulation has come from, and who's going to be caught by it, but what do you need to do if you are falling within the scope of the regulation? There seems to be this focus on ex-

ante assessments. In other words, before you make these systems available, before you deploy them, you have to go through some assessment processes to make sure that they comply with the regulations. So what sort of risks and main areas do you need to be assessing for compliance? Shima, what's the EU's thinking about? What do you actually need to be doing to put these systems into practice in the EU?

Shima Abbady:

Indeed, as previous mentioned, the AI Act is akin to product safety regulation, so it's primarily written from a 'product safety' angle rather than another angle such as data protection. Hence, we're looking at very traditional product safety compliance obligations, but of course, these are applied in a new context — the context of AI. For providers of the systems, these obligations include maintaining a risk management system, maintaining quality systems and monitoring these systems, having up-to-date technical documentation, record keeping, logging, and so on and so forth. For deployers of systems, there will also be obligations, for example, to use the systems in accordance with the instructions of use, ensuring that the data you are using is also compatible with the intended purpose of the system, or monitoring the operation of the system. For deployers, most obligations are in the implementation stage. However, ex-ante, providers will have to comply with a lot of the more traditional product safety obligations that are going to be filled out further by standards that have yet to be published. Notably, the risk that the EU regulators are really concerned about is 'product safety' specifically from a fundamental rights perspective, so you have to look at the risks to health or safety of national persons, and their fundamental rights, including potentially equal access and equal opportunities. In this process, Parliament has especially mentioned concerns along those lines, which have also been echoed by the other parties, the Council and the Commission, but Parliament wants to also look at the risk to democracy and the rule of law or the environment, to give a few examples. In other words, there could be quite an extensive list of risks to consider.

Toby Bond:

That's a pretty extensive list. It is interesting to see you take the concept of traditional types of product safety harms to people such as physical harm and extending it into new types of harm, both to the individual as well as perceived potential harms to society, so it's all part of a broader shift we've seen, especially with the discussion around online harms.

But that's how the EU is thinking. Wilfred, what do you need to do if you are falling within the scope of the regulation? What is the Chinese perspective on the type of assessment which needs to take place before these systems can be made available?

Wilfred Ng:

Thanks Toby and Shima. In a nutshell, the regulations require service providers to be responsible for the content generated by generative AI, including ensuring that the content is accurate and truthful, non-discriminatory, and in line with the laws, regulations etc. in China. Additionally, service providers are required to 'earmark' to the public any content generated by AI in a conspicuous manner. Where non-compliant content is generated, service providers would have to prevent the generation of such content, or optimize the training model, and report to the competent authorities, strictly speaking. Attesting to how quickly things have been moving, the Interim Measures took effect in August 2023, and by October, the draft version of the Basic Security Requirements for Generative Artificial Intelligence Services was published by the TC260, which are somewhat of an ex-ante nature as well, and which supplements the regulatory framework made up of the regulations mentioned earlier. So, for example, if the generative AI technology is specifically applied for services which have public opinion attributes or the capacity for social mobilisation (which are not new obligations, as a separate regulation exists for internet-based information services capable of creating public opinions or social mobilisation; however, these are particularly elaborated upon in the Interim Measures for the generative AI context), then you will have a filing requirement for your algorithm, as well as being required to undertake a self-security assessment, which needs to be filed alongside the algorithm. The supplemental standards on the security of these algorithms really set out what kind of assessment you need to do in detail, so, for example, you would need to satisfy certain prescribed rates of filtering prohibited content deemed to be 'security risks' (such as discriminatory content, content violating the legitimate rights of interests of others, infringement of intellectual property rights etc). There are very specific thresholds and a test which you need to run with your algorithm before you're able to file with the regulator. Thus, in a sense, China's ex-ante considerations are not too dissimilar to the EU, and ultimately, the security assessments go towards ensuring the legality of

the training data and ensuring that the model generates safe and appropriate content. Also, in the supplemental standards, there is a very detailed section on the kind of content filtering of the training data, so, for example, you would have to proactively ensure that the training data does not contain the prohibited content deemed to be security risks so that they are not disseminated through your algorithm. You would also need to demonstrate that any IP infringing content will not be used as part of your training data, and the same also extends to personal data privacy considerations (i.e. if personal data is involved here, whether consent has been acquired). So, in practice this might not be couched as an ex-ante assessment per se, but it only makes sense for the developer (as well as the organisations who are going to incorporate this technology) to fulfil all these obligations prior to rolling out their product, due to the sheer collaborative effort which would be required within their legal teams and their product teams. I think this is also quite specific because the nature of the Chinese regulations are quite content driven, and also have a very specific focus on the underlying algorithms of these technologies.

Toby Bond:

Yeah, absolutely. It seems to me that quite a lot of the action is actually going to be at the level of these technical standards. Shima, that was something you were mentioning for the EU, as we still need to wait and see what a lot of these technical standards are actually going to say to implement these broader principles in the acts. So, our discussions would likely continue at that technical level going forwards.

Shima Abbady:

Yes, the technical standards setting bodies will have a lot of power, that's for sure.

Toby Bond:

Indeed, we see that this applies to all of the existing EU product legislations. You have these harmonised standards which you can comply with. You get a presumption of conformity if you comply with them. You can go about doing it other ways, but it is generally a lot easier to follow the standards and show you are working that way for them. We will have to wait and see on how the process will work. Wilfred, you touched on registration requirements briefly. Can you explain a little bit more about how that works?

Wilfred Ng:

Sure, in essence, if a generative AI service has public opinion attributes or the capacity for social mobilisation, it is subject to a filing requirement under the Interim Measures, which, apart from a self-assessment security check list, an extensive list of information regarding your underlying algorithm is also expected as part of the filing. Apart from details about the functionality and nature of the algorithm, there should be sufficient demonstration of how the algorithm meets the prescribed security thresholds. So, for example, the draft security standards require your algorithm to demonstrate at least a 90% success rate in the detection of prohibited and unlawful contents or key words through various tests. This will be part of the filing requirements to be submitted within 10 working days upon service commencement. In practice, given the volume of requisite information to be included in a filing, an AI service provider would likely want to 'front-load' the obligation in advance of the 10 working days window, making it essentially an ex-ante obligation in practice. I am also very interested to know Shima's view as to how the current EU legislators may look at what China is doing in terms of these all very extensive filing and registration requirements.

Shima Abbady:

I think that China and the EU are definitely somewhat aligned on that. First of all, I think the EU also is very much focusing on the ex-ante process, the ex-ante documentation and assessment and all these product safety requirements that have to be met. In the EU you would have to get a CE marking ('Conformité Européenne', French for European conformity) which means that you have to go through a conformity assessment (which would likely would be a self-assessment for most high-risk systems) Once that is complete, it would mean that all of the obligations under the AI Act are met and the system can be marked with a CE marking to designate it as a 'safe' system. After that, the system has to be registered in an EU database which will also have to contain information about the system — we'll have to wait of course for details of what exactly will have to go in that registration. But, it is likely that the system will be publicly registered, meaning that there will have to be some public commitments about how the system works, what it can do, as well as its limitations, potentially. The EU also wants to go as far as to potentially also oblige particular types of deployers of high-risk systems to register their use of these systems. For example, public authorities are now very likely to be subject to this obligation. The parties have also agreed in the political

agreement to impose information provision obligations with regard to the workings and limitations of general-purpose AI-models, which will, again, be most stringent for general-purpose AI-models deemed to entail 'systemic risks'.

Toby Bond:

Indeed, we have registration in both jurisdictions. There is obviously a lot happening in both jurisdictions in this space. If you were an organisation, who is going to be impacted by these rules, and what should you be thinking about? How do you go about approaching this sort of regulation — what are your top tips? Where would you start, Wilfred, from a China perspective?

Wilfred Ng:

I don't think it will be significantly different from your usual internal guidance or external guidance for your employee's usage for these technologies, or even for your vendor to incorporate, for example. The reason I say that is because some of the issues are very familiar to many lawyers already, such as how developers can ensure the legality of data sources. Beyond data protection, this could mean that content-specific considerations involving infringing IP contents and contents that could potentially violate national laws are all relevant — for instance, considering how organisations can pre-empt and provide redress when there is a violation. The classic 'take-down' mechanisms will be a common recurring potential cure, same with how you respond to data subject rights requests. This would hopefully all be part of your existing data subject request (**DSR**) policy. The only thing I would highlight is the definition for service provider in the measures, because again, it interestingly catches both the developer and the enterprise user who would incorporate the AI technology, and on that basis the measures specifically state that the service provider in the broader sense has primary liability to be responsible for any content violations, and will also be deemed as a data controller. Now, one might query that service providers can have two permutations here, as it can include both the developer and the organisation who is incorporating the technology — so this is again a very classic issue of how you contractually allocate the underlying risks, particularly when you procure the AI technology from the developer. Additionally, when you have your contractual negotiations in this context, a lot of the effort and time of that negotiation will likely be spent on determining who is ultimately going to be responsible for the content. This also gives rise to a little bit more nuance in the AI context: For instance, upon the development of a particular application of generative AI technology, how far is the enterprise customer able to optimise it? Can the customer tweak it? And if the customer tweaks it, how should liability issues be allocated? Hence, I think we will be having a very familiar set of conversations, but in a very different context.

Toby Bond:

Shima, what Wilfred was talking about was just ringing so many bells with me with regards to some of the work you and I have been doing with contractual arrangements from the EU perspective, in terms of how you manage that regulatory responsibility, and how that gets managed in the contracts. As an organisation, where else do you need to be looking if you are potentially going to be touched upon by these regulations?

Shima Abbady:

I think Wilfred put it beautifully when he said that it is a lot of familiar obligations in new contexts. I think that is very much true here as well. We have obviously been dealing with digital regulation with data elements, with data, with content regulation, with IP; in other words, with basically all the areas that are going to be caught by AI. When it comes to this regulation, I think the most important thing to start with is to just be aware within the organisation and to ensure that you have on your radar what is happening with AI. That is already the point where I see it going wrong in many organisations at the moment. If you have an overview, then you can also assess the risk, including the likelihood of being caught by the AI Act. Then, you'd know what time you should be starting, and you'd know when you might be caught by regulations. It is possible that there is going to be an exception, or an extra-long grace period for systems that are already on the market before the AI Act enters into force or even into application, and so that might mean that not all of your systems would *immediately* be caught; Nevertheless, it would still be smart from say, a liability perspective to look into those systems as well to make sure that they would also comply with at least the legal norms that we can see from the AI Act, breach of which could result in torts, for example. With regard to how your organisation should start your compliance process, of course, it is currently somewhat difficult to do everything immediately, given that we are still waiting for a final text of the AI Act, not to mention the fact that after we have a final text, we will be waiting for standards to be released. Nevertheless, there are already a lot of standards for AI ethics and AI risk

management in existence which would be a good basic starting point. You could start, for example, with the assessment list for the Ethics Guidelines for Trustworthy AI from the EU high level expert group set up by the European Commission. That is a fairly basic one, but if you want to go more into technical specifics which is recommendable, then you could look into standards that are already existent, such as those from the National Institute of Standards and Technology (**NIST**) or the International Organization for Standardization (**ISO**) norms to name a few. If you are a deployer of systems, then, in addition to preparing for the AI Act yourself, it is good to look into whether the vendors that you are purchasing from are also looking into all of these things demonstrably, as that would likely mean that you are safer purchasing from them than from others. Furthermore, one recommendation for all parties would be to ensure that you have your contracts in order and that you have internal policies that correspond to your risk management efforts. It is in most cases also recommendable to perform a Data Protection Impact Assessment (**DPIA**) and a more general impact assessment (e.g., regarding fundamental rights). I hope that provides some framework to start with.

Toby Bond:

To draw to a close, the two key messages for me are that there is going to be a lot of work that needs to be done here to get yourselves compliant with the legislation — get ready for it, think about your contracts, think about your polices and government. So there is a lot to do, but there is also a note of hope here which I think both of you expressed with that, which is that this isn't entirely new, this has not come totally out of the blue. There is a lot of work and a lot of thinking that has been done already which organisations can draw upon and leverage. Thank you, Shima and Wilfred, that was an absolutely fascinating discussion and it was really interesting to hear both your perspectives brought together.

twobirds.com

Abu Dhabi ● Amsterdam ● Beijing ● Bratislava ● Brussels ● Budapest ● Casablanca ● Copenhagen ● Dubai ● Dublin ● Dusseldorf ● Frankfurt ● The Hague ● Hamburg ● Helsinki ● Hong Kong ● London ● Luxembourg ● Lyon ● Madrid ● Milan ● Munich ● Paris ● Prague ● Rome ● San Francisco ● Shanghai ● Shenzhen ● Singapore ● Stockholm ● Sydney ● Warsaw