

THE **NEW** BAZAAR

DECEMBER 19, 2024

AI and the Global Battle for Tech Supremacy

Why diffusion matters even more than innovation

CARDIFF GARCIA: Jeffrey Ding. Welcome to the New Bazaar.

JEFFREY DING: Great to be here.

CARDIFF: Before we get into your theory about General Purpose Technologies and the importance of diffusing them throughout the economy, I want to actually talk about some individual technologies themselves. Here's where I want to begin. You co-authored an essay with Helen Tomer and Jenny Xiao called [*The Illusion of China's AI Prowess*](#). What is that illusion?

JEFFREY: I think the illusion there is specifically this belief that China is poised to overtake the US as an AI superpower. And so in that essay, we were showing how China is still very much a fast follower in artificial intelligence. Specifically, we looked at 26 different Large Language Models that Chinese labs were building and showed how they were still pretty far behind the US, and some of the challenges that China has with deploying some of these Large Language Models, specifically as it relates to the need to have information controls on sensitive political content.

CARDIFF: Is it that the technology itself is lagging, or is it something else that's preventing them from being able to present to the world true competitors to the American models?

JEFFREY: I think back then, when we were writing that report, there were a few different factors. One was, at that time, there were still some major data limitations in terms of the high-quality Chinese-language internet data. There's just more English-language data like that out there. If you think about academic papers, even some of the best Chinese scholars, they'll publish their work in English.

So there's just a data limitation that's since faded a little bit because Chinese labs are now training their Large Language Models on both English and Chinese.

Another factor we were looking at were compute restrictions, and that was partly informed by the export controls that the US has imposed on semiconductors, so limitations in terms of accessing those high-end chips that are needed to train AI models.

And then the last set was related to censorship and this obsession with controlling the spread of information that sometimes just doesn't jive well with these freewheeling Large Language Models, where it's hard to even predict what they're going to output.

So the Chinese government has struggled to regulate this space too, and there have been some overly stringent regulations that have imposed some very high costs on Chinese Large Language Model providers to ensure their models are not spewing out any sensitive political content.

CARDIFF: I see. So it's a combination of burdensome regulations, lagging technologies, and some of the kind of economic warfare that the US and China are waging against each other that actually does seem to be slowing down Chinese advancements in AI. Is that right?

JEFFREY: Yeah – it's a combination of all those, but that gives you a sense of the range of factors.

CARDIFF: It's interesting to me that China does seem to have surpassed the US and Europe and Japan and everybody else in the world on a couple of technologies.

Specifically, I'm thinking about electric vehicles and the batteries that are used in electric vehicles (and in other things). And I find this kind of startling because of a point that you also make in a lot of your research, which is that, as of right now, the average Chinese worker is still much less productive than the average worker in the US.

In the past, whenever one country would overtake another or come up with a technology that shoots past everybody else, the productivity differences were a lot closer. China's still way behind. Are you surprised that China's been able to leapfrog everybody else in a couple of these technologies?

JEFFREY: In one sense, it's pretty remarkable what Chinese companies have done in some of these fields, like electric vehicles. I think on some level though, I'm not really that surprised. Just to give some context to what you were saying, in some of my research, I've looked at, "When does one country surpass another

to become that leading economic power?" And that's oftentimes based off of having a more productive economy, right. Labor productivity or an indicator economists use called total factor productivity —

CARDIFF: Which is like a measure of underlying technological advancement.

JEFFREY: Yes, and so, in the past, when two great powers have been competing, and, as you were saying, that gap in productivity levels was much smaller... One interesting case, though, is Japan, which started out at only around 30 percent of the US's productivity levels and actually shrunk that gap to just 5 percent difference at one point.

And so for Japan, they had a lot of different industries that were world-leading like —

CARDIFF: This was in like the late 1900s, right? Second half of the 20th century?

JEFFREY: Exactly. Think 1980s, 1990s. Similar to what China is doing with electric vehicles right now, Japanese consumer electronics were dominating the world stage. But the issue was, with Japan, you had these really remarkable industries on the one hand, but then some other industries, like in the services sector of the economy, that weren't doing as well.

So, it became a study of contrasts. And so, I think, without that general widespread productivity growth, we might see the same thing with China, where you do have some pretty impressive sectors, like electric vehicles, but you're also going to have that contrasting sector. You're going to have those other lagging sectors too.

CARDIFF: Yeah, and one interesting thing about this, which is very much a central part of your theory, is that there's a difference between General Purpose Technologies, which are technologies that can influence every industry across the economy (or most industries) — versus some even very important technologies that are a little bit more self-contained, have some positive spillover effects, but they're not going to necessarily make every single industry better, or even most industries, right?

JEFFREY: Yeah, I think electric vehicles is a good starting point as an important technology but not one that's general-purpose, because electric vehicles really only have one use case, which is transportation. Of course, transportation has some spillover effects, to your point, across the entire

economy, but not as much as something like a General Purpose Technology, such as artificial intelligence or electricity in the past.

I think a good test of this is: Think about how advances in AI will fundamentally transform the electric vehicle industry in the future, whereas advances in electric vehicles are not transforming the field of AI. So one technology is more enabling than the other. And so in the past, economists and economic historians have called these General Purpose Technologies “engines of growth” because they often precede these huge waves of economy-wide productivity growth.

CARDIFF: What are some good examples of past General Purpose Technologies that really had that kind of effect? Electrification, I'm guessing, is the big one.

JEFFREY: Yes. The big three are usually electricity, the steam engine, and the computer. Oftentimes, those correspond with past Industrial Revolutions. The steam engine with the First Industrial Revolution, electricity with the Second, and the computer with the Third.

CARDIFF: If you think about a computer, it's not like that stayed contained just within the IT sector. *Everybody* uses a computer now. *Everybody* became more productive because of computers, and then eventually some innovations [arrived] that were partly dependent on the computer, like the internet and that kind of thing, right?

JEFFREY: Exactly, and that's one part of what makes a GPT unique: it has this ability to diffuse and spread throughout the entire economy. There's this pervasiveness to it, and as with the computer, it wasn't just contained to one industry. All these other industries that started using the computer had to come up with their own complementary innovations too, right?

We had to reorganize the way businesses were run. We had to reorganize the way we trained engineers across all these different sectors.

CARDIFF: Even the way factories were run. This affected the physical world as well. It didn't just stay within the realm of bits.

JEFFREY: Exactly, exactly. And that's actually one of the most illustrative cases, from electricity. At first, when electricity replaced the steam engine in terms of powering factories, they just substituted that central steam engine that was driving all these shafts and belts to power different machines.

They just substituted it with an electric dynamo, but that didn't really bring the productivity and efficiency gains; only when they completely restructured the factory layout so that each individual machine was not powered by these shafts and belts connected to a central driver, but [rather] each individual machine could be powered by electricity. It was a much more decentralized layout. That's where you got the productivity gain so there's all this organizational adaptation that has to occur, too.

CARDIFF: One of the points you make in your book repeatedly is that we should be humble about our ability to identify what the next General Purpose Technology is. And you have this great example where 20 or 25 years ago, people thought nanotechnology was going to be this incredibly transformative thing.

You don't really hear too much about nanotechnology anymore. It hasn't gone away or anything like that. It's there, I'm sure there's improvements being made and all that, but now you're more likely to hear about AI. You might hear about things happening in the biotech sector. And I want to kind of zoom in on some of the technologies that exist now and see if we can identify some of the things that might suggest that they are a future General Purpose Technology.

So why don't we start with the case for AI? What do you see as the components of AI that might lead it to become an amazing General Purpose Technology and not just some clever thing that we use to interact with on ChatGPT or whatever?

JEFFREY: I think I would emphasize three things. The first is AI seems to show this scope for continual improvement. So, oftentimes with a General Purpose Technology, there's an entire research paradigm around it. And so it's not just a very concrete application. It's almost like a foundational transformation.

And so we see that with AI, with the growth of the deep learning field and these continual improvements that are happening with AI. I think the second component would just be all the different industries that are using AI. We're starting to see a diverse range of use cases, not just in image generation or perception capabilities in automotive vehicles, but also when it comes to coding assistance, writing assistance. And now you're starting to see the broadening of use cases.

And then, the last thing is there are some numbers we can use and some quantitative indicators we can use to try to identify General Purpose Technologies early on.

So for example, if you look at patent citations, for a set of machine learning patents compared to a set of biotechnology patents, the patents that are citing machine learning patents are coming from a much wider range of technology classes than the ones that are citing the biotechnology patents.

And so, when someone is citing a patent, that might be an indicator that they're starting to use that technology or they're trying to build something off of those technologies. I think that is the case for AI as this era's GPT. But as you were saying, we should be humble about this. Maybe nanotechnology is coming back. Maybe we just haven't been paying close enough attention. Maybe biotech has these general-purpose capabilities as well, but those are some of the indicators you can use to identify.

CARDIFF: Yeah, it'd be great if the nanotech evangelists from 20 years ago just turned out to be super early, but right instead of wrong [CHUCKLING]. That can happen, right?

JEFFREY: Totally.

CARDIFF: You made the point earlier that it took a long time, decades, before electricity had that widespread effect on how we do essentially everything.

It seems likely that if we end up using AI across the entire economy that that also would require a lot of reorganizing how we do work. Is that also your view?

JEFFREY: That is my take. If you look at these past precedents, it will take multiple decades for a General Purpose Technology like AI to make its mark on the economy. Now, other people, I think reasonably, have said maybe AI is different in that now we're starting to see recursive improvement, in the sense that now AI systems might help AI labs do better AI research and develop better AI systems. And then you get almost this virtuous cycle —

CARDIFF: All of these artificial intelligence systems will begin communicating with each other and the improvement will be instantaneous, or it'll feel that way.

JEFFREY: That's the feeling, that's the conjecture. For me, I still think, in terms of actual commercialization and landing, it will still be a little messier than that. And you're going to need humans to monitor those systems, to be able to deploy these systems. You're going to have to do the sort of organizational adaptations we talked about earlier.

So for me, I still think it's going to take a prolonged time period for AI to diffuse throughout the economy.

CARDIFF: You said something interesting earlier, which was that artificial intelligence can help make possible, or at least transform, the way electric vehicles are made, produced, and used — same thing with autonomous vehicles — but that it doesn't really go in the other direction.

EVs can maybe provide some extra data inputs or something to AI, but beyond that, not gonna make a huge difference [to AI improvements] and that that's a good way to identify what might be a General Purpose Technology.

Using that logic, I guess I don't want to rule out the biotech stuff just yet, because if that makes us all healthier and helps us live longer or whatever, that could feed back into the economy because we'll be better workers, clearer thinkers, and so forth. Right?

JEFFREY: Yeah, biology is so fundamental to everything. So if you can make someone a more productive and healthy worker across all these different fields, then that has that same sort of generalized productivity boost.

CARDIFF: Let's now go to the theory in your book, General Purpose Technology theory, and what it essentially means for the possibility that a new country can become the global economic and technologically dominant power, the new great power, possibly even leapfrogging the current great power.

So, a General Purpose Technology has to be the thing that makes this possible, or multiple General Purpose Technologies — not just leapfrogging with some new, even very impressive innovation. So we were talking about this earlier, but that's step one. Is that right?

JEFFREY: Yeah, it's identifying, first, what are the technologies that have the most potential to shift the economic fortunes of nations?

And I think General Purpose Technologies are that prime candidate.

CARDIFF: And you note that a lot of previous scholars, international relations scholars, often forget that point. They'll point to some amazing new innovation and say, “Oh my God, we're done for. China has the new EVs. Pretty soon they'll have everything.” And that's not necessarily the case, right?

That's one big mistake the traditional thinkers make.

JEFFREY: Exactly. And we went through a little bit with people thinking Japan would become the number one economy because they were leading in certain industries like consumer electronics and semiconductor components —

CARDIFF: Cars, for a while.

JEFFREY: And cars — exactly. But they weren't leading in the General Purpose Technology, which was computers and computerizing all these different application industries. Not just in manufacturing, but also in the service industries.

CARDIFF: The second step is that an economy has to have the ability to widely diffuse that General Purpose Technology so that it can be applied to a whole bunch of industries and not just stay within its own realm. Is that right?

JEFFREY: I think, Cardiff, if you look at all these past cases, when you're looking at competition between great powers and rising powers in fields as broad as AI or electricity, no one country is going to be able to monopolize innovation in all these sectors.

No one country is going to have the only frontier firms in AI or electricity. So then it goes beyond just the capacity to pioneer new fundamental breakthroughs and to this capacity to diffuse them across the entire economy.

CARDIFF: So far it's: You gotta be close to the bleeding edge on a breakthrough. That breakthrough should be a General Purpose Technology. And you should be able to diffuse that technology throughout the economy.

The next part is where we get into some crucial differences between the US and China.

What does the ability for a technology to be diffused throughout the economy depend on?

JEFFREY: There's a bunch of different factors that we can go into. What I focus on in my research is institutions that build human capital because the skill base of the country feeds into everything else. It feeds into your specific technology policy. It feeds into your patenting systems and intellectual property systems. It feeds into all these other factors we could talk about.

And so, my argument is that in this current competition between the US and China, successful adaptation to AI, the successful ingredients to AI leadership,

will rest on the ability of that country's institutions, education, and training systems to widen the engineering skill base connected to AI.

CARDIFF: Here's a quote from your book. You write, "When widespread adoption of GPTs is the priority, it is ordinary engineers, not heroic inventors, who matter most." What do you mean by that?

JEFFREY: I think as I was going through all the historical cases, the conventional wisdom would really valorize the heroic inventors, the James Watts of the world.

CARDIFF: Tesla!

JEFFREY: Tesla. Yes. And today even, we celebrate the technologists. The pioneering innovators. And actually, when I went back through and poured through a lot of the new work that has come out, in terms of work by economic historians and historians of technology, they've showed that Britain in the First Industrial Revolution, for instance, didn't really have an advantage when it came to heroic inventors.

France was producing a lot of the major breakthroughs at the time too. But Britain's advantage was just in this average technical literacy across the people who needed to be familiar with applied mechanics.

And it was the same thing with the US too, in the Second Industrial Revolution, 1870 to 1913. At the time we were sending our best and brightest chemist students to Germany to complete their PhDs. We weren't at the scientific frontier, but the US was able to win out in that Second Industrial Revolution because they were able to train a wider base of mechanical engineers.

CARDIFF: You also emphasize, and I'm drawing from a recent essay of yours, institutions that incentivize technology transfer, trade openness, along with human capital. How important is openness to the rest of the world in this process?

JEFFREY: I think part of the reason why trade openness is so crucial is that no one country is going to be able to go it alone in this. In this current age, even among the most advanced economies, the majority of their productivity advances are going to come from international sources. And so, openness to other economies allows your firms and universities and market actors to be plugged into those global innovation networks so that you can quickly adapt an innovation that gets pioneered somewhere else.

CARDIFF: The American university system, for all its flaws, and I know there are quite a few of those, is still the envy of much of the world, and it's quite big and varied. It seems like, based on your theory, that is just a massive advantage over everybody else. Is that right?

JEFFREY: 100 percent. For me, one of the reasons the US has been so successful in adapting to all of these General Purpose Technologies is because it has a decentralized approach to its university system. And so, you have this diverse mix of all of these varied types of universities. They can set their own curriculum and adapt to the local conditions that they need to foster — to spread different technologies.

And so, in the context of US-Japan competition, for instance, Japan's approach to shaping its software engineering curriculum was very much top-down, and having their education ministry try to centralize the control of that curriculum. And they tried to foster these centers of excellence.

CARDIFF: So it's like by telling people what to study, what to emphasize, instead of letting people choose and letting it be a little more organic.

JEFFREY: Exactly. And we see that with China's attempt to institute an AI major across their entire university system. You would never see something like that in the US system, where the Department of Education is saying, “Okay, all these schools have to develop an AI major.” You would let schools, for some schools where it makes sense to maybe develop a separate school of AI or some schools would fold it into the computer science curriculum.

And I think that that's part of the advantage, the adaptability of the US system.

CARDIFF: You mentioned before that it's really helpful to have a lot of human capital and, in particular, a skill set that makes it easier to apply the General Purpose Technology. And you mentioned AI engineering in the case of AI.

One question I had was, what if the skillset that's really useful for applying AI is a little different from engineering or the skillsets that in the past were helpful for applying previous General Purpose Technologies? What if some element of creativity or other kinds of thinking, or other skills generally, are the ones that are really important — or even some in combination?

JEFFREY: I think that's a fascinating thread. I think it could very well be the case that the skills needed to adopt AI are not necessarily only those that map onto engineering-based skills. You might need to foster creativity. You might

need to foster a lot of knowledge about the humanities too, in terms of trying to understand how to interact with these systems and communicate with these systems because a lot of times, especially with the development of these AI agents, maybe adoption is more of just an effective communication exercise.

And so there's a lot of different skills that need to be learned. I think [the reason] why I focus on engineering skills is because engineering disciplines systematize and standardize the best practices associated with these past General Purpose Technologies.

And that lowers the threshold for adoption. That makes it easier for someone like you and me to, for example, use an engineering protocol to fine-tune a Large Language Model and apply it to the particular industry we're involved in.

CARDIFF: And maybe even what we refer to as engineering itself could change in response to the development of this technology. I remember like 10 years ago when the advice was “everybody needs to learn to code”. And now what I keep hearing is: “Well, AI is going to make coding obsolete.” You need to be able to use the AI or be good at prompting or whatever, as opposed to being able to code itself.

That took no time at all, to go from one to the next [LAUGHS].

And I know that that's just what I'm hearing anecdotally. I'm sure coding is still an extremely valuable skillset, or maybe even it just helps you think a certain way and that itself could be really valuable. So I'm being a little bit facetious here. But it's just interesting to me how the skills that are necessarily helpful, useful, and that could help in the context of your theory to make it easier to diffuse technologies more easily — that stuff could change. It could evolve over time.

JEFFREY: Exactly. And I think that's a great example of why it's tough to make technology policy in this area. Because with AI, if we were talking two, [or] two and a half years ago, we might not be talking as much about Large Language Models. We might be talking about computer vision and facial recognition, right? So if you—

CARDIFF: Like Google Maps or Google Earth or image recognition? That kind of thing?

JEFFREY: Image recognition, yeah, using—

CARDIFF: AIs in cars, and all that?

JEFFREY: Exactly. Using AI techniques to identify defects in the production line, early on in the manufacturing process — “computer machine quality inspection” is the term for that.

So that might've been our conversation, but if you had institutionalized a policy focused on training talent for those particular subdomains, you might have missed this huge explosion in the natural language processing field — all these Large Language Models that we've been talking about.

CARDIFF: You cite some research showing that patenting and research & development levels are not as good at predicting long-term productivity growth and technological improvement as are these indicators you just mentioned of diffusion capacity.

Were you surprised to find that? Or, based on your research, were you like “Nope, this fits in perfectly”?

JEFFREY: It was surprising to me that these science and technology indicators we often gravitate towards, these patenting indicators, these R&D indicators, were not as predictive —

CARDIFF: Still important though!

JEFFREY: Yeah, still important...

CARDIFF: Just not as predictive.

JEFFREY: Not as predictive as we would have expected, and not as predictive as indicators of diffusion capacity.

So, some of the indicators that were more predictive were like: number of books published on a particular technology subject. And that being more of an indicator of this technology actually being used and diffused. And there's strong communication channels sharing information about what's happening in this space.

I guess it makes sense on some level because adoption is when a technology actually makes its mark, when it goes from the lab to the shop floor and to actual usage.

CARDIFF: Here's a concrete example you give. It seems like this idea helps explain why we won the Cold War against the Soviet Union.

So I'm just quoting you here: "By 1970, the Soviet Union was the global leader in R&D spending as a percentage of gross national product, outpacing the United States on this metric. And yet, they could not widely diffuse any of the findings that came out of this R&D research, or any of the technologies that came out of that."

JEFFREY: Right, and I think maybe the most prominent example of this is — think about that Sputnik moment, right? The Soviets are the first to pioneer the satellite, but who actually developed a robust economy based off of satellite technologies and adopting some of these satellite applications across the entire economy?

It was not the Soviet Union. The Soviet Union struggled to access these fast-acting, almost organic diffusion processes that the US had, in part because their economy was just overly planned. They could meet these production targets, but it wasn't something that small- to medium-sized firms in the Soviet Union naturally wanted to adopt.

CARDIFF: Flexibility, bottom-up, decentralization, market-driven system as opposed to central planning. These are all themes that have repeatedly come up just in this conversation that you and I are having. And it seems like that's a big deal for diffusion.

JEFFREY: Yeah. I think that's definitely one piece of the puzzle. I think, on some levels — we'll get into this later, I'm sure — but on some level, there are some industrial policies and strategies that states can implement to try to correct areas where there are market failures.

For example, in skills training, sometimes private firms might underinvest in training talent in these General Purpose Technologies, because they're worried that they'll train up their workers, and they'll leave and go benefit another firm.

So that's an example of a market failure where you can have more coordinated and directed action. But, in general, I do think that the US benefits from its more decentralized, bottom-up processes, at least compared to some of these planned economies.

CARDIFF: Yeah, I'm front-running the part where we talk about some of the policy implications. But because we've talked about human capital so much, the

university system, it also seems like high-skilled immigration is a natural way to go here.

JEFFREY: Yeah, and I think that was part of the US's advantage over Japan. Japan just could not access a wider base of software engineering talent because they had such restrictive immigration policies. And they were relatively isolated from these global talent flows, whereas the US was able to attract so many software engineers during that time period.

CARDIFF: Yeah. Fascinating. We can also start to see the outlines now in this conversation of some key differences between the United States and China. So you make the argument that China has, and this is your phrase, a 'diffusion deficit,' which is this idea that a country can have very strong innovation capacity but that it is lagging when it comes to diffusion capacity.

So what's the argument that China has this diffusion deficit?

JEFFREY: Yes, so, when you look at, for example, let's just take the domain of AI, when you look at highly-cited publications, patents, even strong frontier labs like Baidu, ByteDance, Alibaba. When you even look at some of the performance on benchmarks for Large Language Models, it does seem like China is getting pretty close to the US on a lot of these indicators, if not even surpassing the US on some of these innovation-centric indicators.

But when you flip the script and look at China's performance on diffusion capacity, its ability to adopt other information and communications technologies — like cloud computing or industrial software — across the entire economy, it trails the US significantly on those indicators.

So that's what I was trying to get out there: that China is almost a bit top-heavy in its science and technology ecosystem, and its innovation capacity far outpaces its diffusion capacity.

CARDIFF: You also looked at the capacity for each country to train average artificial intelligence engineers. And you've got this statistic here that I want to read from, which is also really interesting.

There's this measure that looks at a university's ability to train an AI engineer based on whether at least one researcher at that university has published a paper in a leading AI conference. And then you looked at it for the years 2020 to 2021. China had 29 universities that met this standard. The United States had 159. Pretty good, right?

JEFFREY: Yeah, I think what I was trying to get out there is, you set a very low baseline for quality — you just need to have one researcher affiliated with your institution that's published in an established AI journal or conference — and that's trying to show, almost like, how deep is your bench?

What is the range of institutions that can train average AI engineers? And on that question, the US is really the world leader. China, once you get past the Tsinghua Universities, the Peking Universities, the top tier — not as much of a wide and deep bench of universities.

CARDIFF: A couple of previous guests on this show, Caleb Watney and Heidi Williams, have a lot of really interesting ideas for how to make academia more relevant to industry, how to improve the communications between the two, and how to pursue ideas that could be more applicable, more practical.

You say this is also a profound problem in China, maybe even a much bigger problem there than it is here.

JEFFREY: Yeah, I think sometimes it's hard to compare across different systems, but one indicator I was looking at there is the number of AI papers where there's at least one co-author from industry, and one co-author from academia. And the US leads the world on that metric. China is at less than half the number of the US.

So I agree that there are issues in the US, but I think we're still doing a lot better than most other countries on this front.

CARDIFF: I mentioned earlier that there seems to be quite a lot of economic warfare between China and the US. This is not a surprise. This has been going on for years. A lot of this was accelerated in the first Trump presidency. But under the Biden administration, same thing, if anything it was ramped up. We have a new Trump administration coming in.

And, most recently, there were export controls that the US put on China for semiconductor manufacturing tools and for these advanced memory chips that you need for AI, for example. So the US made it harder to export those kinds of things to China.

And then China recently retaliated by prohibiting to the US exports of gallium, germanium, antimony, super hard materials. I don't know what any of those things do. I just know that they're needed to make semiconductors and some military equipment and batteries.

So there's all this economic warfare happening, but you kind of make the argument that this is based on, and again, using your words, a “mistaken assessment” that China is on the verge of becoming a scientific and technological superpower, that there's a different risk with China.

Can you take us through that?

JEFFREY: Yeah. So I think a lot of the first Trump administration's approach, and the current administration's approach, is very much based on the assumption that technological leadership is about which country can monopolize innovation in this new fast-growing industry, like AI.

And so, with the export controls, it seems like we are trying to build up Fortress America and prevent any of these crown jewels from leaking out. Because that's how other countries are going to overtake us.

Whereas, what I'm trying to present in my research is: actually technological leadership is not about cornering all the innovations in the space, especially with General Purpose Technologies. No one country is going to dominate innovation.

It's about who can play this long game, who can run this marathon competition to adopt AI across the entire economy, this run faster model.

And so, that's where I disagree with what seems to be a bipartisan consensus actually, that Fortress America is the way to go on AI competition with China.

CARDIFF: Very interesting. You also look at the possibility that China is what a scholar named Michael Beckley refers to as a “peaking power”.

And he makes the point that if you look back over the last century and a half, roughly, a peaking power is one that had a long-term economic boom, but that boom is now slowing down. And so he makes the point that [the boom] has given these countries the capacity to shake things up on the global stage because they have improved so quickly, they're more economically important, and so forth.

But the slowdown has also prompted them to become more aggressive. And you kind of make the argument that that's the risk we should be looking at, and not necessarily dominance in any one particular technological sector, as important as that sector might be to parts of the economy, or even to the military.

JEFFREY: I feel pretty confident that the US is very well positioned to win the AI race in the long term. So this other risk that you're talking about, to me, that should be more front of the agenda, which is Michael Beckley's point. I think one way to express this point is that a weak China might be more of a threat to US national security interests than even necessarily a strong China or a stable China.

And so, if the export controls actually do what they're intended to do — it's almost a form of containment. If they actually work, if you are able to close all the loopholes, potentially you are really hindering China's economic growth overall.

We know that peaking powers potentially could be more dangerous for the international system. And so, it seems counterintuitive in this environment — especially in D.C., which is where we're sitting right now — but it's actually been US diplomatic policy for decades, which is that a strong and stable China might be more in the US national security interest.

CARDIFF: How do we balance that idea, though, with the notion that we want to be as secure as we can be when it comes to national security, when it comes to the military, and frankly when it comes to a lot of economic sectors as well, and in particular if China increasingly sees the US as a geopolitical adversary?

It almost seems like there's no way to avoid some version of this [containment strategy]. I'm not saying that the particular way we're going about weaponizing our economy against China is the right way. I don't know. But it seems like some of that is inevitable, because we're trying to make ourselves, the US, less vulnerable to a country that views us as a threat, or at least as an ideological and geopolitical opponent rather than as an ally.

JEFFREY: Totally. And I think that is where the current Biden administration's small yard, high fence idea is really useful. There is going to be some competitive friction, to your point, but can you almost bracket that off and insulate it in a few areas without it affecting the entire economy? Like we had with the block systems in the Cold War?

CARDIFF: It seems so complicated, though. I asked Paul Krugman a version of this question as well in an earlier podcast, and he said, “Look, if we import a lot of toys or shoes and things of that nature from China, things that have no national security implications, that's totally fine. We should not have tariffs on things like that just to hurt China for the sake of hurting China. But on things like industrial capacity and national security, we have to pursue this.”

And because of that approach, it means that something like semiconductors is going to get tied up in this. And yet semiconductors are obviously quite crucial to both countries when it comes to staying on the technological frontier.

JEFFREY: Yeah, I think it's tough to implement this in practice, to your point and to Paul's point. My starting point though is just being clear for AI — what is different about AI than toys? Or what is different about AI compared to even semiconductors? And then trying to craft an effective technology policy from there.

And if you think of AI as this General Purpose Technology, that no one country is going to bottle up, then it makes more sense to adopt that run-faster model.

CARDIFF: Let's talk about some policy implications of your work on General Purpose Technology diffusion theory. We mentioned high-skilled immigration. That has shown some clear advantages in the past, and you've said already that it would be clearly advantageous for the US to pursue now and in the future.

What else? Are there too many regulatory burdens when it comes to diffusion in particular — to that part of it and not just to the innovation part? When you look across the US and the economic and political landscape, what else could potentially help make the US better at diffusing new technologies?

JEFFREY: One of the things I emphasize in my work is: how do we broaden the base of AI engineering skills in this country? And one interesting thing is that the CHIPS and Science Act actually had already set up some of the building blocks for this. The plank of the CHIPS and Science Act that gets the most attention is the “CHIPS” part — all the investment in semiconductor production facilities.

But for me, the more important aspects of that act were the “Science” part, which invested a lot into STEM workforce development. Now I'll give you one guess as to which one of those planks got a lot more money. That money was allocated much faster and implemented much more quickly.

CARDIFF: [CHUCKLES] It all went to the “CHIPS” I'm guessing.

JEFFREY: Exactly, and I think that shows the tension between these two approaches.

There's no constituency, no concentrated interest group, that's saying, “Hey, let's make STEM workforce education a little bit better for everyone,” right? That's a

diffused, dispersed interest, almost by definition. And so those are the types of policies I would push — fully implement that part of the CHIPS and Science Act.

Another potential recommendation would be to provide greater backing to community colleges as another pathway to developing AI engineering skills. So those would be some policies on the skill formation front.

Also, as we've been talking about, you don't just want to develop a lot of AI engineers. You want people to have the practical skills and be connected to industry, right? These university-industry linkages.

There's some cool policies I think the Canadian government has implemented that encourages internship programs jointly funded by industry and universities to get those sorts of practical skills for students, those experiential learning opportunities for them while they're still in school.

And so that would be another set of policy ideas.

CARDIFF: A lot of your research seeks to overturn long-standing thinking on these issues, and on what makes it possible for one country to either catch up to or leapfrog another country. How's it been received?

JEFFREY: Well, [LAUGHS] I would say that there's been some resistance from the US national security community. I think the main counterargument they pose is that AI is different from these past General Purpose Technologies.

And in particular, I think a lot of folks have this view that there will be some sort of fast takeoff scenario with AI. And by that I mean, what we were talking about at the very start of this conversation: AI systems getting recursively better, building other AI systems, then coordinating with other AI systems —

CARDIFF: Like a sudden spike in exponential growth or something?

JEFFREY: Sudden spike in exponential growth, and voila, you have an army of a hundred million programmers that can execute cyber security attacks and boost some sort of decisive military strategic advantage. And there, that deviates pretty greatly from the General Purpose Technology diffusion picture.

And so, for me, it's hard to say, right? We're all just kind of monkeys throwing darts at the wall. And I think the way I approach it is: I'm trying to research and make predictions based off of thinking about what is AI like. It's like a General

Purpose Technology. Let's look at how competition over General Purpose Technology has unfolded in the past.

For these folks, they're very plugged into what's happening in AI today, in the frontier labs. Very much in that sort of inside view.

And so, I don't think that one is definitively wrong, and one is definitively right — I think we should just have a greater mix of both perspectives, and see what comes out in the wash. I just think that we hadn't really had this alternative perspective yet. So that's what I wanted to insert into the debate.

CARDIFF: Very interesting. Has your work made you more optimistic or pessimistic about the economic and technological future of the US?

JEFFREY: It's definitely made me more optimistic about the economic and technological future of the US.

CARDIFF: Really? Why?

JEFFREY: Yeah, I think, as I mentioned in that last chapter of the book, I've come to the conclusion that the US is very well positioned in competition over AI with China. Some of the statistics about universities that we cited, the ability to broaden the base of AI engineering skills, and sort of the starting point that the US already has with its lead in productivity across the entire economy.

I think what's made me more pessimistic is the trajectory of the current administration and future administration towards missing the boat on what actual effective competition looks like.

And then that means you can't address other risks, like the ones we talked about, sort of like a China lash-out scenario. Or you can't address AI safety risks as well, and trying to manage accidents from these powerful AI systems.

And those are things that might require a lot more cooperation, not just competition.

CARDIFF: What's something about your work that you think people often mistake? It doesn't have to be about your work or your papers specifically, but about the topics that you cover that people get wrong, and that you'd love to correct the record on.

JEFFREY: In general, one interesting — not necessarily misconception — but I think one really interesting thread is just this idea of technological determinism. How much weight do we assign technology in shaping global politics?

So maybe one example that's a little bit off the beaten road that we've been talking about is, think about all this conversation about China exporting digital authoritarianism, right?

So if Zimbabwe imports a Chinese facial recognition technology, that will make Zimbabwe more likely to be more authoritarian. And so that's this idea that technology has almost this kind of independent force that can transform a country's politics. And so, I think rightly, people have said that is a misconception.

Society, people, we have the ability to shape technology for its own purposes. Technology is just a tool. For me, in a lot of my work, I'm trying to get into the middle ground of those two perspectives. Technology doesn't determine everything, but there are certain technological trends, like General Purpose Technologies, and how they make their mark. And we should see some regular patterns in that over time. And then, society has to adapt to those regular patterns.

CARDIFF: Jeffrey Ding, thank you so much for being on the New Bazaar. This was a really great chat.

JEFFREY: Thanks for having me.