ECONOMIC INNOVATION GROUP

THE NEW BAZAAR

March 23, 2026

IDEAS FOR A POST-YIMBY HOUSING FUTURE

ARPIT GUPTA ON WFH, HOUSING AFFORDABILITY, AI, AND MORE

CARDIFF GARCIA: Hey everyone, Cardiff here. Our guest today is Arpit Gupta, a finance professor at NYU, but more importantly, something of a polymathic scholar whose work spans across a whole bunch of topics that not only fascinate me, but which really could not be more relevant to the moment we're in right now. A moment in which we're all trying to figure out new ways of working, what AI might mean for the future, especially how it'll affect our work, and what can be done to end the housing affordability crisis that we find ourselves in.

So, work, AI, housing, finance, urban economics — Arpit's work ranges across all these themes, and we're gonna try to hit on as many of them as we can in this chat. Arpit is here with me in the studio. Arpit, did I leave anything out?

ARPIT GUPTA: That sounds great. Thanks so much for having me. Yes, also regular podcast guest.

CARDIFF: Yeah. A little too wide ranging, to be honest, for us to fit into one thing. We'll probably have to skip something, but let's start with something really fun and really new, which is remote work. You've got a new paper out on remote work, and I'm fascinated by this topic. I have been for a while. I interviewed my colleague Adam Ozimek, our chief economist at EIG, about this not too long ago.

But your paper sheds some light on something I've wondered about, which is all the criticism that's come from the heads of these big companies, and who I've been quite quick to criticize for what I considered to be their archaic views on remote work. But what you find actually shows that maybe they've got a point.

ARPIT: Yeah, absolutely. So taking a step back, I've done analysis on remote work in different areas. Started with some of my colleagues looking at the impacts on migration and residential real estate, moving on from there to commercial real estate and offices. And the question in the back of my mind through all that was the issue of is this gonna stay? Is this gonna be a sticky technology? Which ultimately boils down to are people as productive working remotely as they are showing up to the office?

CARDIFF: That's an open question still, right?

ARPIT: I think it's still kind of an open question, especially if you think about firm heterogeneity, right? Surely it's not gonna be one answer for every type of worker, every type of firm. So we wanted to get a sense of not only the point estimate, but a little bit also the distribution and what are the factors that help remote work function or not for different organizations.

CARDIFF: Yeah. And I should say also that to this point, it looks like remote work is a lot more common than it was before the pandemic — which, oh my God, that's six years ago now. But it is lower than it was at what you might consider the height of the pandemic. It has settled in at a kind of higher equilibrium overall.

What your work did was to look at how different kinds of firms, and in particular firms of different sizes and different ages, how it works for them and how different it can be. So yeah, let's go right into your findings. What'd you discover?

ARPIT: So, we need to measure productivity for this at an individual level. And this turns out to actually be a really hard problem. We don't have great ways of measuring white-collar productivity the same way that we look at counts of the number of widgets to calculate manufacturing productivity, for example.

So we turn to one source that's been increasingly used by several scholars, which is the quantity of GitHub commits. This counts for coders — and these coders, by the way, can be multiple different professions. It's not just people in the tech industry. It can be people that are coding in many industries — but it counts for these coders how many commits, or how many software pushes, they're doing in any given day.

CARDIFF: So for a normal person who doesn't know anything about coding, a GitHub commit means what? How many updates you make essentially, and then it tallies that?

ARPIT: Right. So imagine, if you're not coding, think of a Word document. Imagine that every time you open up a Word document, turn on track changes, make some edits, and send that back to your collaborator, right? That's gonna count as one document edit, or in our context, is gonna be a GitHub code commit.

CARDIFF: Right. And then you use that as a measure of how much work somebody's actually getting done because it's something that actually can be quantified, can be measured.

ARPIT: Exactly. And it's something that, having talked to some firms, they actually use internally as a way of trying to measure the success of different workers. And it's something that we correlate in our analysis to measures of individual career progression. So people that are getting more commits have more career progression.

Firms that are doing more commits seem to be more successful. So it kind of correlates and associates the way you'd want it to as a proxy for output. And that's gonna be your measure for productivity or output for different firms.

And the key finding is that firms that are larger, that are older — that are older than 10 years of age — see a decline in quota productivity after they go remote. Whereas these younger firms, these firms less than 10 years, actually see an increase in productivity.

CARDIFF: And here's where I want to go back to the framing I used earlier of how you have these CEOs of big companies like Jamie Dimon — they were saying that essentially they want their people to get back to work as quickly as possible, back to the office as quickly as possible. They're being quite rigid about it. I think some other Wall Street firms have also essentially said, we're coming back to the office.

And again, my own response to that was, maybe these old-timey CEOs should try to figure out what's going on here, because it seems like such a useful thing to at least be able to offer people. But if it really is the case that by going remote at these older, bigger, more established companies, the workers are becoming less productive, then I kind of can't argue with it. What I can argue with is the idea that their conclusions on remote work apply everywhere.

ARPIT: Right. Because we do find that it's for those younger companies that it does seem to work for them.

And then we did a lot of work to try to figure out what is it that's different about the large organizations, the older established organizations, versus the younger organizations that contribute to that. We have two things we think are going on. The second is actually brand new, so it'll make the paper soon, I hope.

But the first reason is pretty straightforward. It's that these smaller companies, younger companies are growing and scaling more. So they're taking advantage of the ability for remote work to help them expand and overcome hiring frictions. For larger companies, it's just less of a big deal because they already have offices all over the country and the world, so they're already able to tap into talent across multiple geographies, whereas startups are more constrained on that dimension. So being able to hire talent all over the place is a bigger deal for them and helps them gain productivity, to the extent these new hires help spill over their productivity advantages to other workers.

CARDIFF: It's interesting because at least that specific explanation is not about how being remote suddenly makes an individual worker more productive because they get to skip a commute or they have more energy or whatever advantages you might get from working from home.

It's more about just the fact that you can hire more productive workers at this company because you have access to people from the whole country instead of just somebody in your town, if you're in a town where there's a shortage of the kind of workers you might wanna hire. Right?

So it's a different explanation than remote work makes you more productive. It's more that productive workers who are gonna be productive anyways now can work at these smaller places.

ARPIT: Well, we do find this spillover effect, so even the existing workers get more productive when my firm hires these other additional productive workers from elsewhere. So there is a team benefit to this larger organization having more productive team members, such that everyone is able to perform at a higher level.

The second finding, which is again brand new, is it seems to be also something related to the stock options that smaller startups are offering. So the way stock options are gonna work is if I am at Google, I may get some stock options with my compensation, but is my effort really going to change the needle on the amount of money that I get for these stock options? Probably not. I can't really change Google's value as a company.

But if I have stock options and I'm at a startup, it's a really binary thing. Either this company succeeds in the next round or it doesn't. My effort can make a big difference in how that works. And so I'm naturally motivated by this aligned financial incentive such that my manager doesn't necessarily need to check up on me all the time, which is very different from a large corporation where, again, maybe I do have some stock options, but it doesn't really matter for my ultimate performance. And so you may need to have the manager actually looking at people and making sure they're working.

CARDIFF: One of the conclusions you make in your paper is that working from home, because of these very reasons, is likely to lead to more business dynamism — more companies that start up and are successful and then also more companies that fail. Right? Dynamism hits you from both sides, basically. There's more competition for businesses.

I also imagine that it would lead to more labor dynamism, right? Because you have better matching between people and the companies where they would be the most productive. And that seems to be a logical conclusion to what you're saying. How big of a deal would that be if it is in fact something that leads to more business and labor dynamism?

ARPIT: Yeah, so we find a couple of results consistent with that. So one is some of these new workers who are more productive at remote firms were not as productive actually at their last firm. There is a difference in their productivity when they join

this remote company. So that suggests there is a possible matching explanation that the worker becomes more productive when they're matched to a particular firm.

The other finding we have is that at the larger companies, when they go remote, we expected to find benefits of retention, which is something that people have looked at in the past. But we actually find that for these larger firms that go remote, they actually lose some existing employees. They grow a little bit, they get some more workers when they go remote, but they also lose some existing employees. And it seems like some of those actually go on to start companies of their own.

So that suggests that we're gonna see more churn, better matching, maybe new firm formation, especially because you can go remote first and already scale to get workers all over the country or indeed the world. So that should have benefits in expanding economic opportunities, especially to underserved geographies.

CARDIFF: Yeah, I was wondering if that was the flip side of the coin of what you were saying earlier about how smaller, newer companies can now hire people from anywhere around the country, which also means that they might be hiring some of those people away from the big established companies that were those workers' only options in the past. You know what I mean? And that could be another channel through which this is harder for the biggest companies. Is that right?

ARPIT: Yeah, exactly. And another finding that we had looked at — I don't think it's in the paper — but just one descriptive thing we had seen is it looked like there was some narrowing of the productivity gap across geographies.

So it used to be the case that coders in California were extremely productive and coders in, say, North Carolina were less productive compared to the ones in California. And it looked like there was some narrowing of that across the pandemic, which you could imagine — that it's these new opportunities to work at different firms that might lead to a convergence in productivity.

That's something we saw in manufacturing, which led to the spread of manufacturing jobs, including to the Sunbelt, and it's possible we're seeing a similar phenomenon happen with services as we allow the entire services sector to become a little bit more tradable and spread out.

CARDIFF: Yeah. This is fascinating because it essentially suggests that when those big company CEOs complain about remote work, they're not really complaining about the fact that they're now losing workers to less productivity because those workers are going home and not working and goofing off or whatever.

Essentially, what they're complaining about is that they now have to compete harder to keep those workers, to keep good workers, right? They're complaining about the new competitive pressure that it represents. I'm not really sure still why I should care

about that complaint. This strikes me as a good thing if it leads to more dynamism, but it does sort of make you look at those complaints from a new angle, right?

ARPIT: Yeah, the most favorable interpretation I guess I would have for the big company CEOs is there may also be attributes of the company's culture that's enhanced by being in person, which is a positive retention force behind why people would wanna stay at a given company. But for sure we have this other phenomenon as well.

So for example, it was Amazon and Microsoft that moved to Seattle to be a little bit further away from the lure of workers leaving for other tech companies. Right? So a little bit of a monopsony story might explain why some of these large companies are unhappy about the labor churn and more national, indeed global labor market we have with this relative remote work.

CARDIFF: How do you think this is gonna affect economic geography? The places that rise and the places that fall throughout the country and how they're distributed. What kind of an effect is it gonna have on that?

ARPIT: I think the results we're seeing for economic geography are a little Goldilocks, if I had to describe them.

In the sense that it suggests there is a durable advantage for these large cities such as New York, because we have here in the city so many large firms that, as our research suggests, do seem to have challenges going remote. So it suggests that while these large firms might stay in these large cities, that should be an advantage for urban areas that want to retain workers to come into the office.

At the same time, it suggests we have another new set of firms, younger firms that have a different model based around being remote first, hiring people all over the place. And those firms should be able to create more economic opportunities for workers no matter where they are. So if you're a worker in a less dynamic environment, maybe you're able to collect new opportunities working for these new firms.

If you're a worker in an existing metro city, then you have other opportunities, particularly working in larger firms.

CARDIFF: The way I've always kind of interpreted remote work is that it's another dimension along which companies can compete against each other for workers, right? And so I just assumed it would lead to better matching. I think your paper offers some suggestive evidence for that.

And I understand the point that bigger companies are gonna face bigger competitive pressures than new startups that now have this abundance of workers that they can

hire, which is great. But on net, I still have never seen a good argument that remote work is bad for the economy, for the country, or anything like that. Where do you fall on that?

ARPIT: I personally feel that it has to have really massive benefits for the economy overall in allowing us to have better choices over where to live and where to work, right? Historically, we've had a very close tie between these two things. You have to live exactly where you work, and that tie, I feel, has led to a lot of negative political economy consequences.

Which is to say, if you're stuck in a location because you have to be there for work, you have fewer opportunities to move and express your residential location, which means that the local government can extract a little bit more value from you. The quality of local services might be a little bit worse.

So I think that is actually one of the long-term benefits of remote work from an economy-wide standpoint, which is to say allowing greater mobility should free up more responsiveness when it comes to corporations competing for laborers. It should free up more responsiveness for local governments that have to maybe compete a little bit more to make sure that they attract enough workers locally.

CARDIFF: All right, great. Sounds good to me. I'm in favor.

Let's switch now to a topic that I don't see mentioned as often anymore as maybe it was a few years ago, which is the urban doom loop. This was originally the idea that in the years after the pandemic, and starting during the pandemic, a lot of people would leave the more densely packed cities.

Now they have the option of working anywhere in the country, partly because of remote work, that it would therefore lead to essentially falling property prices for, I think, both commercial and residential areas inside of big cities, downtown cores in particular. And that because of that, you'd end up with less tax revenue and then services would get worse, which would lead to more people leaving and so forth.

Doesn't really seem like that happened. I think some places have suffered more than others. The place that everybody worried about so much is our city of New York City. It really doesn't seem to have happened here, right?

How do you characterize what ended up happening over the course of the last few years, and why it was that this worst-case scenario just never came to pass — unless you think it still might?

ARPIT: So first of all, our paper — this is with Stijn Van Nieuwerburgh and Vrinda Mittal — was about trying to quantify the commercial real estate office impacts of remote work.

So we are finding these large reductions in value of commercial office buildings, and I believe that has still held if you look at the stock prices of various REITs that are estimating the market's assessment of the value of this office real estate. They're not doing well.

CARDIFF: You're talking about across the country.

ARPIT: So first of all, that's true kind of across the country, but then especially in some of the hardest hit areas like San Francisco or New York, those values are still very impaired. If anything, there's a more recent narrative around AI taking out white-collar work, which has further depressed the values of these office stocks.

So that was the main thing we were looking at in that paper, is the office value. And I think those values are still pretty in rough shape.

We then explored the possibility for the other consequences of that shock, and we mentioned this urban doom loop as a possibility or risk that cities would have to think about. We were drawing an analogy to the urban doom loop that did happen from the sixties and seventies, which affected many cities in the Midwest and Northeast as a result of de-industrialization.

So we were kind of having in mind this whole cycle of you lose employment, you lose industries that result in some population loss, which adversely impacts the quality of public services provided in ways that lead to an ongoing spiral.

So I think the phrase “urban doom loop” is a very nice label to describe that whole cycle, which previously we all knew about but maybe it didn't have the same label.

CARDIFF: Is it too simple a label? Like does it try to capture too many different things essentially?

ARPIT: So it's an analogy to sovereign debt doom loops, right? So the phrase “doom loop” in economics kind of originates from thinking about some of these circular dynamics that happen in international markets.

So we're trying to draw a domestic urban analogy to that, and I think it's a reasonable phrase to describe a cycle of downward spirals that, again, has happened in the past. We pointed to it as a possibility that would happen as a result of remote work or other shocks, and I agree, it's an open question to what extent that is gonna really materialize.

I think what we've seen is a combination of factors. One is kind of related to the first paper we talked about. It turns out that there are a lot of benefits to being in person, so not all of the remote work has materialized, and there has been some return to office.

Another factor is that the policy environment is really crucial. So how cities cope with these negative shocks is a really important aspect of how they come out of them. So some cities like San Francisco have actually responded with very positive policy responses. They've also had the AI revolution, which has revitalized their urban areas in some ways.

At the same time, other cities have actually taken different policy responses that threaten bigger losses. So Chicago, I think, is a prime example of a city that has had some initial shocks as a result of remote work and the commercial office value losses, has responded to that by lowering some service qualities and raising some taxes in ways that are threatening to lead to further rounds of urban flight and changes in migration, which might amplify these dynamics further and lead to further financial distress.

So I think it's an accurate phrase to think about the dynamics, but it's ultimately the job of politicians and all of us to figure out how to respond to these shocks in ways that might dampen their influence.

CARDIFF: How's New York doing?

ARPIT: There are positive and negative aspects of New York for sure. So certainly the return to office is a huge plus. The success of the A-plus office markets — that was another big part of our paper, separating out what's going on at the high end of the office market versus the low end. The success of that high-end office market is certainly a huge positive.

The city's assessment for taxes, for better or worse, isn't really recognizing large losses. So the office owners are still paying pretty hefty property taxes. So we haven't seen that round of reassessment of taxes that would result in a fiscal loss.

So those are all kind of big strengths and plus points for the city. On the modest point, I think you would point to there have been some pretty substantial losses over the last few years in terms of employment in different sectors. And if you look at the long-term trajectory, New York City has really lost out its share of finance jobs, its share of millionaires. So relative to other parts of the country, it is not the same mecca for high-end work that it used to be.

CARDIFF: I mean, we talk a lot about the projected future deficits in New York City. That's still a problem that has not been solved. And so maybe the urban doom loop you're describing is just something that takes a long time, instead of one massive shock that leads to chaos over a few years. It might just be something that drags out or acts as a drag on the economy for a while.

ARPIT: Exactly. We're trying to identify a series of dynamics that relates to how cities ultimately respond to negative shocks that they experience. So right now, as

you point out, New York City is facing a little bit of a fiscal deficit, and what really matters is how we respond to it.

If we respond to it by cutting essential public services or raising taxes that encourage some residents to leave, that's gonna amplify some of the negative aspects of those shocks in ways that have ongoing repercussions.

CARDIFF: Yeah. The two things you just described are the two things that usually most people look to as, alright, you gotta pull one of these levers or both in tandem. And they both have those damaging effects. So it just seems like a tricky problem unless you grow your way out of it. Right?

ARPIT: Right. It's certainly a tricky problem. So there's either growing your way out of it, which is to find new growth opportunities for the city, or you find ways of cutting wasteful spending. And surely we have some wasteful spending in government that we could target.

CARDIFF: Oh, you're kidding!

One thing you've written a lot about as well, I think, is office-to-residential conversions. And this is an idea that has always seemed quite appealing because especially in a place like New York, where you have a really severe housing affordability crisis, young families leaving on mass, especially after the pandemic, in part because of that affordability crisis — if you end up with more residential supply and you do it in a part of the city where, as you just said, commercial real estate values are down, then that seems like a real win-win situation.

But it also seems we've come to learn over the last few years that it's really hard to get that right — both the actual technical, logistical side of it and the politics side of it. So how are we, not just New York but the country at large, doing on office-to-residential conversions? And how realistic is that as a possible solution to some of these problems?

ARPIT: Yeah, so I've been really surprised at the extent to which we've been pursuing these office-to-residential conversions in the last few years. Many different cities have various types of policy incentives to try to encourage these conversions to happen. New York's program seems to be relatively successful in—

CARDIFF: Oh, it does.

ARPIT: —in helping to convert office units to residential units. There are a lot of benefits to trying to make it happen, which is to say, on top of everything you mentioned, you're typically getting more housing units in the very center of the city where typically all your city's transit lines are running into.

So the same transit that's intended to take commuters into the office core — if you're in one of these buildings, you can take that same transit line out anywhere else in the city. So you're already starting with a pretty well-connected set of locations that might be appealing to households and families.

And we don't need that many conversions to actually put a dent in the level of office vacancy that we have, right? So if a city has, let's say, a 20% office vacancy level and takes 5% of those offices and moves them to residential, that may not move rents very much, but it might make a huge impact on the office market.

So I think we're seeing a lot of success in these office conversions across the country. I've been really surprised at the scale and scope and what the architects are able to make happen in these different buildings. So I think we'll be seeing more of this, and I think more broadly we will see this principle of adaptive reuse — how can we take space, repurpose it, do different things with it — happen as we encounter more shocks in the future.

CARDIFF: The way I've always thought about these conversions is that if you look at one of these massive high rises in midtown, you go floor by floor — none of it looks like residential. It's this huge layout. The bathrooms are in weird places and they're also communal bathrooms. How do you end up carving that out into six different residential units instead of one big office space?

So I always understood that that could be tricky. The politics of it also seems to be really hard. You need to change the laws sometimes to allow for that kind of thing. What do you think is the biggest barrier that remains? Are you surprised more by how the politics have shifted on this, or by how clever the designers are in terms of making this something realistic?

ARPIT: I've been surprised by both. So the politics of it, I think, have been easier than I would've expected because you don't really have as much of a NIMBY problem — which is to say, in your office neighborhood, there's no one around to complain about the additional residents. So you don't have to deal with that issue as much.

CARDIFF: The building’s already built.

ARPIT: Buildings are already there, and you save on embodied carbon by not having to tear down all this steel and concrete and so forth.

So the main challenge has been figuring out the right policy mix of what abatement, what affordability requirement, et cetera, is necessary to get the whole thing to pencil out. But a lot of cities have been experimenting with different types of policies, and many of them seem to be working.

The architecture firms are also quite creative, it turns out. These are companies like Gensler, which does a ton of these conversions. They'll figure out all sorts of wizardry to manufacture cores that go through buildings, cut out sides of buildings to create more sunlight space, and just a whole series of creative solutions to help these things pencil out.

When I first started looking at this, my prior was that it's gonna work in FiDi, the financial district, where you have these tall, skinny skyscrapers that you can convert to residential because the floor plate is relatively narrow. But these architects are showing you can make it work in all sorts of weird midtown buildings as well. And the rezoning efforts are really helping to facilitate those.

CARDIFF: I wanna switch to another idea you have for helping with the housing affordability problem, which is changing how we approach property taxes.

ARPIT: Yeah.

CARDIFF: Give us the background there.

ARPIT: So this one is a little controversial, because if you were to ask people how do we make housing more affordable, probably the last thing on their mind—

CARDIFF: Raise taxes.

ARPIT: —will be raise your taxes. But it rests on a very natural principle of capitalization, which is to say if you're looking at two identical properties and one of these properties is gonna have a higher property tax liability in the future, well surely that property is gonna be worth less than the other because you're capitalizing the value of those future payments, and you're gonna be willing to pay less for an asset that has those higher obligations in the future.

So that's just very normal, sensible capitalization math, and we do some work in the paper to try to estimate how large this can be. But the upshot of that is if you are increasing your property tax liability, you are making at least the upfront purchase price of that house more affordable.

At the same time, what you're doing is you're increasing the — sometimes we call this a carry cost — of how much it actually costs you to hold the house period by period. And so that acts as a little bit of a nudge to make sure that people in the house actually really value the asset and they aren't just hoarding it, staying there just because they bought it some time ago, but they actually really value the asset they're in. So it encourages a little bit of churn in migration in the housing stock.

CARDIFF: Yeah, I wanna stay on that point of how it reduces the upfront cost of the house. The idea here is that if you know that property tax rates into the future are higher than they are now, then you're going to take that into account when you buy

the house — and so are all buyers — and so it acts as a discount, right? You have to lower the upfront cost of the house.

And what that could encourage is younger homeowners who don't have as much upfront money, but they're looking at 30 to 40 years of work because they're still only in their twenties or thirties. Now they can buy a house.

And one of the big problems that we've seen in the housing market, which my colleague Jess Remington just did a huge piece on, was the way in which housing affordability differs for, on the one side, new and prospective home buyers, who are really having a tough time, versus existing home buyers, who have it easier.

For example, seniors tend to be existing home buyers more so than people in their twenties to thirties. And so if you have higher property taxes, it'll make it a little bit more affordable for the youngsters to buy homes. And yes, it will add to the cost of the seniors, but they're also the ones that have a lot more accumulated wealth in the first place.

ARPIT: Yeah, exactly. And one thing to emphasize there is that because we're trading off these future liabilities for the current price discount, it's not like the asset is fundamentally cheaper. We've just changed the time profile of when that expense is borne.

But that kind of tends to work in the favor of younger buyers who are more likely to be down payment constrained, which is to say you're looking at, let's say you have $100,000 saved up for a down payment, you're looking at a $1 million house. If you wanted to do a 20% down payment, you can't make that work.

Whereas if the property tax stream was higher such that the house price fell to $500,000, now you can take your existing saved down payment, put that into the house, and then get that property.

So the property tax stream is working kind of like a mortgage, because what a mortgage does is trade off less upfront expense in exchange for future obligated payments. So the property tax is kind of adding a mortgage to all these properties, which helps the financially constrained population who otherwise can't qualify or get these mortgage products.

And it acts as this additional cost, which, if you place it on these existing properties, now acts as a little bit of a nudge for these older households to maybe reoptimize their housing situation.

CARDIFF: Your paper finds that this isn't all just speculative. There's actually some empirical evidence that, and here I'm quoting you, that areas with higher property tax rates feature more young homeowners, fewer empty bedrooms, because you don't

have as many seniors who are now empty nesters still living in the house where they raised their kids, so there aren't just a bunch of empty — there's no idle capacity there, you might call it. There's empty bedrooms.

A higher percentage of children in the population in places with higher property taxes, and lower house prices, and price-to-rent ratios. So it seems to work. I will say some of this surely depends on how the property tax rates are applied.

I could see a scenario — and I think this also exists in some parts of the country — where if the property tax rates are assessed on the purchase amount of the home, okay, and it stays that way over time, well, people are gonna be reluctant to leave that home to buy another one because if they buy another house at a time when property tax rates are higher, their house, which may have gone up in value over time, will be assessed at a higher amount and the tax rates will be higher and then they're just never gonna leave.You get that lock-in effect, right?

ARPIT: Yeah, absolutely. So we kind of compare some of these property tax instrument impacts, which, as you point out, work best when you're applying that property tax at a true market rate of the house as opposed to some other forces which might lead to lock-in.

So what you're pointing out is one thing that we didn't look at, but will surely lead to lock-in, which is if the property tax assessments are a little bit stale in ways that discourage transactions. Another thing we did look at is capital gains taxes. So if you are — the way that the taxation system works is if you can defer the realization of a housing transaction until you pass away, your heirs can inherit it with stepped-up basis and avoid a capital gains liability.

So that's another force for lock-in. There's no point in reselling the house when you're old, incurring some capital gains, if you can instead hold it. So there are a variety of ways that the taxation system might encourage people to remain locked into properties, whereas an alternative would be to try to find taxation instruments that instead encourage households to re-optimize their housing stock in ways that will support more churn and more opportunities for young people to access housing.

CARDIFF: Here's another idea of yours, and it applies more, I think, to the rental housing market. You're in favor of something called accelerating depreciation schedules. All right, let's explain that for normal people first.

If I'm understanding this correctly, this is essentially the idea that if you are an institution and you commit to building a multifamily rental property, that you can essentially expense the entirety of the upfront cost of doing that right away. Right?

ARPIT: Yeah. And what's key about that is changes in this expensing schedule don't really change the total amount of tax liability received by the government. It's about

changing the timing of when that tax liability is due. So it does change the present value of the total tax paid, but it doesn't change the undiscounted total value of tax liabilities.

What that basically means is for the developer that has a hurdle rate when they're applying investments, right — so they're willing to invest if they receive a rate of return that exceeds a certain benchmark — being able to accelerate the timing of the tax benefit associated with housing is able to lower the hurdle rate by improving the timing of that investment return. So that's enough to get more projects to pencil out that wouldn't have otherwise because of the timing of when the cash flows are realized.

CARDIFF: Yeah. And that would lead presumably to a lot more multifamily construction. And it also, by the way, would be in keeping with the accelerated depreciation that we do allow for other things already. Right? So some forms of investment you can already expense right away. This would just be making multifamily structures consistent with those other things. Right?

ARPIT: Right. And we had this system in the eighties. We paired that system of more accelerated depreciation back then along with some other tax shelter rules that went even further. And so they provided an even further benefit to the tax benefits for owning real estate.

And the collective impact of all that is we went through a little bit of a real estate hangover in that period — which we would love to have today. We were building way, way too much.

CARDIFF: Way too much!

ARPIT: A lot of the office buildings we're now trying to convert to other uses today were kind of coming out of that overall real estate boom from that time period. But that shows you how sensitive these investment decisions can be to their tax treatment and the after-tax return.

CARDIFF: Yeah. There is a funny part of your piece where you were writing about this very topic, and how, in the past, when these depreciation schedules have been introduced, a lot of existing corporate real estate companies resisted them. Essentially, they saw 'em as more competition, and you wrote, and I'm quoting you here, “it's important to fight against the real estate lobby and insist that they receive tax benefits.”

That's the best way to get new construction. Right? Basically forcing these institutions to get these tax breaks because you'll get new construction, but they just didn't want the competition. Maybe they didn't want there to be more units out there 'cause it means that they would not be able to charge us higher rents.

Is that usually where the resistance comes from here?

ARPIT: I think that's basically it. That's my understanding of the industry's pushback when they were considering some of these rules back in the 2017 TCJA. In fairness, I suppose to the industry, there is a reasonable concern around how much this is gonna be a boom-bust cycle to real estate.

So what we would, you know, ideally like to do is have an improved incentive for real estate construction that'll kind of apply across the cycle. We kind of want to have a sustained investment towards more real estate, but what we did see in the eighties is that providing these incentives kind of did lead to a little bit of a boom-bust cycle in real estate.

I.E., during the best years of the eighties, we're getting a ton of new construction. The impact on rents wasn't just a level decrease in rents, but it was related to these kinds of dynamics that ultimately were destabilizing.

CARDIFF: Well, it had a very — it was extremely volatile in the eighties as well.

You had the savings and loan crisis, which was sort of tied to a lot of these tax changes. But, if it also leads to a ton more building, then that's the thing we should be worried about.

ARPIT: Yeah. Yeah. So that's—

CARDIFF: Not worried about it in a bad way. I mean, that's what we should want.

ARPIT: Right. So that kind of comes back to a little bit to our discussion on remote work, which is, "do we care about what the corporate executives feel about remote work or are we concerned about the aggregate benefits?"

So I think the industry is probably right to worry that there might be some boom-bust dynamics that are potentially affected by tax credits. But we need to, as a society, prioritize the housing affordability in my view.

CARDIFF: Final thing on housing is a big one because you've got an upcoming paper that's titled "Industrial Policy for Factory Built Housing."

And I want to get a little bit into this. It's important to note here that you are already, as I understand it, pretty much on board with most of what people refer to as the YIMBY project, the Yes In My Backyard Project, that we do need better zoning codes to allow more construction in a whole lot of parts of the country. And this is

your attempt to add something to that. Not to contradict it at all, but to do something that's complementary with that. Is that a fair thing to say?

ARPIT: Absolutely. And one sort of — this is with Steve Teles — one sort of thought experiment we have in mind in conceiving the piece is: imagine the YIMBY movement succeeds, which is to say we address a lot of the zoning constraints to housing construction. Is that gonna solve the problem? Is that going to lead to housing abundance, or are there remaining barriers even after we address some of the land-use constraints that apply to housing?

And in our view, there still are gonna be some additional constraints that pertain primarily to the construction productivity that we see. So we've had the last several years of declining construction productivity. That means it's getting harder somehow to build homes than it once was. And that requires addressing if we're gonna be able to lower the actual cost of building, because even if we address some of the land-use constraints, we also have to address the cost of building the structures.

CARDIFF: I actually wanna pause to emphasize something here. It's not just in the last few years that we've had lower construction sector productivity; the stagnancy of construction sector productivity has lasted half a century or more by some estimates. It is the strangest thing. It is one of the biggest outliers in all the economic data that we have right now for an entire sector to have declining productivity growth. For this long. For decades and decades. It's weird, but it's a reality.

I wrote about this, by the way, at the Financial Times in 2014 in a post that I'd labeled “The Astonishing Decline in Construction Sector Productivity.” That's how long people have been studying this.

And I was obviously building my work, I was just a journalist, on the work of other economists and other people who study the construction sector. So it's a shocking thing to have happened. There's all kinds of reasons that are floated for why it's happening. And a lot of those are pretty realistic.

We don't exactly know what's most contributed to it, but it is a staggering thing that this exists.

ARPIT: Yeah. And to add to how puzzling it is that we have this decline in productivity, right? Not just stagnation, but we're actually progressing backwards in our ability to build houses.

If you talk to a builder or go to a construction site, it seems like they're doing new things. Like the power tools are different, right? The measuring tape is different. I don't know. We have lasers on site. We are doing—

CARDIFF: The individual technologies have advanced, but not how we use them.

ARPIT: Exactly. Every individual component in how a home is constructed seems like it's had improvements over the last few decades, just as everything else has in the economy. But the composite task of putting together that entire house has somehow gotten worse in productivity.

CARDIFF: Yeah, it's amazing. If I were to tell somebody that we are worse at making calculators than we were in 1976, you'd be like, "what are you talking about?" (CHUCKLES)

You know what I mean? Like, it would be beyond comprehension that we could actually get worse at making something that we already knew how to make. But it applies in the construction sector. So let's talk about your idea for what to do if, in fact, someday there is a YIMBY win.

Or even if there are some YIMBY wins, and it's not a wholesale victory, something that can help. Your paper again is titled "Industrial Policy for Factory Built Housing." And we should describe that first, we should define what factory-built housing actually is. It's two things: it's prefabrication and modular construction. Let's explain what those two things are first.

ARPIT: Absolutely. So I think that the title of the piece may have changed. So I think it might not be "Industrial Policy for Factory Built Housing," and I actually don't know what the latest status is. In any case.

The point of the piece is to talk about changing the way that we build housing away from onsite, which is how housing is traditionally done, towards more factory-based methods. So the modular side of that would be to build entire components that could then be assembled on site. Factory-built, in general, refers to building things off-site, which would allow for concentrating productivity improvements in controlled environments in factory settings, which allow for us to take advantage of all the benefits we get from learning by doing and successive capital investments and productivity growth that have characterized how productivity has grown in other sectors.

So we can borrow from that and try to adapt that in this sector because it does seem like the reason that housing has not had the same productivity advantages is that

you're doing everything on site. So imagine you're trying to build a car, and one day, someone shows up with the windshield wipers. And then, a week later, you're waiting for the electrician to show up to build your car. You would never build a car that way.

CARDIFF: Like you never buld it in your own home garage or something. You never build a car in your garage, one piece at a time or whatever. (CHUCKLES) So what you're saying essentially here is that this is houses or big parts of houses, whole rooms, that are built in factories and then moved to the site of where the house is gonna be, as opposed to building the house at the site of the house, right?

And then, that way you can take advantage of studying how it works in the factory when you're putting those parts together, making improvements. And then over time, you get better and better at it. Just like with cars, right? You started making them in factories, and then people would study it, and then they'd come up with better and better ways over the last century of how to make a car in a factory, and it just gets better over time. You'd see improvements.

ARPIT: Exactly. And if you look at other industries, like the car manufacturing industry or shipbuilding or other industries like that, they really struggled with figuring out how to get those economies of scale and industrial-scale production. And something a bunch of industries had to cope with is that when you hit a downturn, you have declining demand, right?

Because you're producing a durable good and people don't wanna buy durable goods in downturns. So that was a huge issue for all industries, but housing being the most durable good is really badly impacted by this. It's sort of in the same category as planes or something like that, and how durable it is.

However, those other industries, like shipmaking or planes, received a lot of government assistance and help, which allowed those types of very durable, capital-intensive goods to continue production through, for example, the Great Depression. However, the Sears home construction business, which was also around in the 1920s — so Sears was building homes, manufacturing them, and shipping them on site — that whole industry completely collapsed because of the demand destruction, which happens in downturns, completely eliminated that industry.

So housing hasn't had the same level of government support to help it smooth through business cycles the way other goods have, which we think is one of the contributing reasons why the industry never adopted these kinds of factory-built methods.

CARDIFF: Let me raise what I anticipate would be a point of objection here, which I'm sure you've anticipated, but the idea that if you could build a house, a really good house in a factory, and you'd get better at building it over time, it'd be getting cheaper. In addition, it would be getting better and better: more varied, more interesting, whatever, than moving it to the house, to the housing site.

And you'd be able to sell it. It would be more profitable to do it that way if that story were true. And so the question would be, why would you need government intervention to help with that? If that were really a better business model for doing things, wouldn't it already be done? And if it's not being done, what are the barriers — if it's not being done for other reasons, right — what are the barriers that prevent it from being a viable business model now for the companies that would do this?

ARPIT: So I think it's a really typical coordination problem, which is hard for any one firm to break out of because we have a very fragmented housing and construction business. Things are very local, and so you're dealing with this patchwork of different rules, laws, and regulations that limit the natural addressable market that you could try to reach through a better housing construction technique.

And so the first thing that we advocate is for more deregulation, which is to say we should address not just the land use codes, but also another type of regulation, which is also being increasingly targeted by the YIMBY movement, which are building codes.

Part of what makes housing very hard to build through these factory-built methods is the complexity of the building codes, which regulate exactly what can be built in different areas. So having more standardization about these codes could help create a more standardized marketplace.

CARDIFF: So essentially, this kind of building is not allowed right now. In some parts of the country, it's restricted, or there are big rules piled on top of it. That actually sounds like it's perfectly in sync with the YIMBY movement. Like it's a part of it.

ARPIT: Exactly. And we have some success stories, I think. So one big success story is what's happened with ADUs in California.

So what they've done is—

CARDIFF: Tell our listeners what accessory dwelling units are. Tell them what they are real quick.

ARPIT: Right. So ADUs, or accessory dwelling units, refer to small houses that can be placed on your site and—

CARDIFF: On property. Like, if you have a house, you build a smaller house right behind you, and you rent it out or whatever, or sell it.

ARPIT: Right. Sometimes called mother-in-law suites. And the reform that California did was to legalize these ADUs statewide, but they also allowed for certain pre-approved plans. So the idea is as long as your plan conforms to one of the pre-approved plans, you can sell it across the whole state.

And what this allowed is the entry of new firms that created these prefab plans that allow you to construct these ADUs in a factory, where I think many of them are union factories too. And so you can build the ADU on site, put it on a truck and ship it, and then just plop it on your site.

It has helped to create a new marketplace that has allowed for new firms and new types of production methods because of the extent to which you have pre-approved plans statewide.

CARDIFF: And is part of your industrial policy agenda that you're proposing here, in addition to the deregulation, does it involve government funding? Does it involve a clever use of government grants, things like that? Can you take us through what else would be involved?

ARPIT: Yeah, so we lay out several additional, more proactive government steps to try to improve this process as well, in addition to the deregulation. So that includes one component, which would just be research and development. This would be trying to figure out new building techniques and methods. And of course, now that we're all talking about AI.

Maybe there's a scope for robotics, for improvements in autonomous vehicles to deliver on-site, for improvements in AI technology to assist in the drafting and the soft costs that are associated with construction. These are all items where industry will also evolve, but the construction industry has been fairly capital-light, and so there's maybe room for the government to take a proactive role in helping to advance research objectives, which would benefit the whole industry.

So again, that's one component on the research and development side of things. In addition, there is a role for the government to act as this guaranteed purchaser, because, again, these swings and booms and busts in the construction cycle do seem to pose a challenge for firms trying to create a large, consistent marketplace.

So the role for the government there would be as a purchaser to ensure consistent demand for different housing products that helps to create a large enough guaranteed market to enable firms to invest in the capital necessary to make the housing.

CARDIFF: How would that work? Would that be that the government would say, "we commit to buying a certain number of these factory-built houses."

What would the government be committing to purchase exactly?

ARPIT: So there are different ways of structuring it. We saw with the CHIPS Act — various types of industrial policy applied to chips production. So one form would be the government directly acquiring housing units for its own portfolio because the government does build housing in many different ways.

So that would be the government agreeing to purchase different numbers of housing units. You could also have the government provide debt support for facilities. So, lending facility support. You could imagine the government providing subsidies or some amount of funding to support the number of housing units you've produced until a certain amount, right?

So there are different ways this support could take place, but it would all have the effect of helping to create a larger guaranteed market.

CARDIFF: Yeah. And then when we talk about industrial policy, that could mean so many different things, right? Like, this is a very nuanced, textured idea. For this, here's what I'm wondering.

One of the skepticisms of other kinds of industrial policy is that the actual plan might look good before it starts going through the legislative meat grinder, right? But then, when it comes out the other side, people have added all kinds of provisions and requirements.

And in particular, during the Biden administration, in a couple of its big bills, which were considered industrial policy, you had things like Buy America provisions. These are ideas that can sometimes make an idea less economical. So in the case of industrial policy for housing, you can imagine a scenario where the government says, "yes, we'll act as the guaranteed purchaser, and in exchange for the money, you have to use only these factories. Or you have to have these new labor requirements that we add onto it."

Or there's a Buy America provision or something like that. And how do you sort of avoid that or how do you sort of structure something so that it's simple enough to get through without all those things being added?

Is that possible? Is that a reasonable point of skepticism?

ARPIT: It's definitely a reasonable point of skepticism if we ever saw these things becoming proposals. Right now, this is sort of an idea phase for us. We wanna get people thinking about the importance of construction productivity and different ways we could try to address it, whether or not they agree with our particular way of addressing it.

And we think a lot in the piece about the political-economy challenges of pulling this off. And so we definitely recommend finding ways of time-limiting the support to try to find ways of getting the industry to a sufficient scale, so that it can stand on its own two feet, so that it doesn't need further government assistance.

Now, one thing that plays a role in that is that we're dealing, in housing construction, with something that's not really a tradable good, which is to say we're not competing with other countries in the production of American-built homes. So on the one hand, that does relate to why we have a construction productivity problem, which is to say we can't easily import the houses produced in other countries.

On the other hand, we think that this is maybe a reason why the industrial policy might not have as negative effects, because we're not really in the business of propping up a national champion. We're not trying to create an Intel that can compete all around the world, and that needs additional support and subsidies for some national greatness objective.

We're trying to improve the productivity of firms that we already have in our own construction productivity ecosystem. That's one thing to keep in mind. Another important consideration is the broader political-economic benefits this type of policy would achieve.

So one of the things we think about in the piece is the fact that we have this negative type of conflict between urban and rural areas, right? Those two regions don't necessarily feel like they're sharing in each other's progress and productivity. Whereas if we found waves of building more factories, maybe in more rural areas, producing homes that are shipped to urban and suburban areas, that would have other political economy benefits and helping to stitch together the economic portions of different regions.

CARDIFF: And it would level economic geographic inequality, I should say, in some ways.

Very interesting. I wanna close with a discussion of finance and AI. You're now teaching a course on finance and AI, and let me start with a question that does relate to housing, because you came out with a paper, with a co-author or co-authors a couple of years ago, on housing regulations throughout the country.

And one of the points you make in the paper was that that very paper would not be possible in the days before the introduction of large language models. Is that right? What did you do and what did you find?

ARPIT: We were interested in categorizing and understanding housing regulations across the country.

So we built a dataset of thousands of these housing regulations that we had collected and scraped, and that corresponded to these municipal codes. And the challenge is now that we have these thousands and thousands of municipal codes, we have all these regulatory questions we want answered to understand the complexity and divergence of these regulations.

And if we had to rely on humans, we just wouldn't have done that. And so what the literature's kind of done is they've either answered a very detailed set of questions on a small number of regions, or they've sent broad surveys to try to get the general sense of our regulations, tight or loose, across the country.

Being able to answer specific regulations across the country hasn't really been possible prior to these new tools, which have dramatically scaled up our possibility to work with data.

CARDIFF: And what'd you find?

ARPIT: We found a lot of cool things, but one of the things we found is that there were kind of two regimes or two types of housing regulations.

One relates to things we associate with value capture. So it's gonna relate to affordability requirements, things like that. These are tools that municipalities are doing, particularly in high-rent areas, in order to extract a little bit of the value that comes with housing regulations. And so I knew about inclusionary zoning, of course, but I didn't realize that there was a sort of whole cluster of regulations that kind of pertained to similar types of requirements.

And the other type of regulation looks like exclusionary zoning, which is to say things like minimum lot sizes, setback requirements, and so forth. These are a set of tools used by municipalities, primarily for exclusionary purposes, to ensure that housing is costly and expensive enough so the riffraffs live there.

CARDIFF: Fascinating. You've also done some writing on the ways that humans can best work with AI. And here's a line I want to read from one of your posts. This is on your Substack, Arpitrage. Very clever, by the way. Here's what you write: "What we need to realize as a society is that mass deployment of AI has the capacity to dull the mind as much as processed foods and easy transportation led to an obesity epidemic. We'll need to figure out the right heuristics, the mental equivalent of working out, to ensure cognitive discipline in a world of AI slop." What do you mean by that?

ARPIT: I think it was pretty clear. I'm trying to summarize some of the research I've read that I teach on now, which has led to the conclusion that people can use AI as a crutch or substitute for their own thinking.

And it's pretty easy to see how that's gonna happen if people are using the chat window to completely substitute for some line of human thought. It can dull and substitute away for the task of hard thinking, which is typically how we build our brain muscles.

So that's sort of the observation and then the parallel with processed foods is sort of thinking back to, "well, this is not the first time that something has happened that made things a lot easier for us, and we didn't necessarily respond to it great at the time." And imagining the future that you might need to have something like an AI Sabbath or something, right? You might need to have a day or a time period where you say, "I'm gonna avoid the use of all AI tools and make sure I'm really thinking about this from first principles."

CARDIFF: AI Sabbath. (CHUCKLES) And a day of AI breaks. I was thinking about it also in the context of your other findings on AI and finance, which is the idea that AI can be used for very easily codeable things, right? Scraping documents and telling you what the right financial ratios are to look at for a company, but the decision of what to do with that output remains very much in the human realm.

That AIs are still not very good at that, at the more creative or thoughtful or interpretive part of the job. And that's interesting in and of itself, but it's also the case that a lot of people develop those skill sets by doing the hard, easily codeable work first, right? So, in finance in particular, if you spend five or ten years learning the accounting, learning the financial statement analysis, and all that stuff, then on the

other side of that, you'll be able to sort of interpret those results just based on the experience you have.

Well now, if you've essentially outsourced all of that earlier work to an AI, is it really gonna be the case that you're gonna be as good at sort of doing the creative part of the job as if you'd taken all the time to learn the rote stuff in the first place, the stuff that you had to memorize in the first place?

ARPIT: This is a fascinating question, which intersects with what seems to be the labor market signal that juniors are being disadvantaged in the labor market, that young people are finding it difficult to get a toehold. And it's unclear how much AI is playing a role in that right now, but it's very easy to imagine that might become the case in the near term if it isn't already.

Which is to say, companies are able to offload a lot of the grunt work that they're currently asking junior analysts to do. At the same time, they're kind of seeing more leverage for the senior people who are able to take on many more tasks because they have AI as an augmenting tool. So it creates this issue in the skill ladder because you don't need juniors the same way you did before.

You have few opportunities for juniors to build the skill sets by doing rote tasks, and again, which is how we currently build that level of experience. And so the way I'm thinking about it is I'm curious as to how the education system could, for example, build those skills instead. I'm at NYU. Can we at NYU create a senior who's graduating and is as good as a VP at a bank, whereas currently they're joining as an analyst?

And to do that, I think the task will be to create simulations or experiments or games and scenarios that help a person see as many examples in a controlled environment as they otherwise would have doing five years of rote work, so that they can get to a senior point already in their career with that level of domain experience.

So I don't know that we can, but that's the task and challenge I'm thinking about.

CARDIFF: Candidly, I'm pretty optimistic about that. I really am. I think this kind of thing happens all the time. People end up extending what they learn in the first place. I don't know. But you tell me like, how's it going so far? You're teaching students all the time. Do you think they're learning different skill sets already than what you might have taught them 10 years ago, 15 years ago?

ARPIT: I can tell you one thing I'm doing differently in my course now, even though I'm also teaching AI, is that I'm doing many more case discussions.

I am teaching the course differently by incorporating in my class many more business case discussions with the objective being trying to get students to develop a deeper set of skills regarding the interpretation of business scenarios. Being able to articulate to other students and to myself what they view of those business scenarios. To express a point of view.

And I think we'll see more things like that. More chances and opportunities to develop those higher-order skills in the classroom so that you're more prepared to take on more advanced roles in the workplace. And to your point, that has kind of been the role of advances in technology over time.

Someone would have originally started work in a white-collar field, let's say a hundred years ago, as a clerk doing all sorts of tasks we now think are completely routine. There's no point in anyone doing this kind of completely routine work. And so we've steadily upgraded the set of tasks that we asked even junior people to do.

And I think that trend will probably continue.

CARDIFF: So my first job out of undergrad, more than 20 years ago now, was as a junior-level banking analyst. And if I had had that job 40 years ago instead, I wouldn't have been using Excel. It would've been a lot more labor-intensive and manual, all that stuff.

And it's not like I used the time that I gained from being able to use Excel to then goof off for a few hours otherwise. The stuff that I had to do was extended. I don't see any reason why that can't apply in some other, you know — right now, to finance — but also to some other professions.

It just might take some time. It might force disruption in how people do things, and also, in your case, in how you teach things. And if it sounds like you're already doing that.

ARPIT: Yeah. We have a whole AI assessment group at NYU, and they're looking at all sorts of different experiments and ways that people can change what they're doing after AI.

And one part of that that does seem very exciting to me is the opportunity to really increase the level of sophistication we expect from students because the expectation is now shifting to, “Okay, well, you have AI, so fine. You can do your old assignment with AI very easily, but maybe now we can give you a different assignment. We expect that you're gonna use AI to help complete it, we're actually

gonna ask you to do a much more challenging task than what we were previously asking you to do."

CARDIFF: Has the ongoing rise of AI and its use in finance and in other parts of life made you more optimistic or more pessimistic about the future?

ARPIT: Probably more pessimistic. (CHUCKLES)

CARDIFF: Really? (CHUCKLES) Why?

ARPIT: So we have the whole issue of super intelligence, how to align humanity with super intelligence, and questions of alignment and so forth. And I'm certainly not the best person to think about all those challenges and hurdles. But what I can't ignore is that the AI researchers that seem to be most interested and invested in these concepts are very worried about them.

They seem to be taking them very seriously. I have to kind of see there's a signal in that, which is the people that are developing the technology, that are in the position to see what is coming, are worried about some of the broader implications. And you don't even need to get to the point that you have a completely autonomous AI warring with humans.

You just worry about the standard problems of human alignment apply to bad actors that now have more resources to produce viruses, to produce weapons, to connect terrorist attacks, things like that.

The long tail of negative outcomes, I do think, is worrisome. And so on average, that does probably make me a little bit more pessimistic.

CARDIFF: Do you think that some of the researchers who are coming up with this stuff are over-indexing on the part of the economy that they know best, and they see a big disruption coming for the people who do a lot of coding work, programming work, and so forth? At least that's what they're anticipating, but they're maybe not appreciating either the benefits that it could have in other parts of the economy or the costs, frankly, like they just may not have an understanding of how it might affect other things because they're only or too much focused on the parts of the economy that they themselves are most intimately familiar with.

ARPIT: That's something we talk about a lot in the course, which is the question of whether AI is a "normal technology or not." And by that, I mean a technology that has some disruptive aspects. It kind of eliminates some tasks and roles that are completely substitutable for what has happened before.

But new technology typically also augment a lot of people's functions, and they also create new jobs and skills in the economy. And the net consequence of all this is typically productivity growth that leads to advances in income. And this is indeed how we've advance as a species to not live lives of drudgery.

So that is the normal pace of technological advance. And even if AI is very capable on certain narrow technical benchmarks, its ability to impact the world is bottlenecked through all sorts of organizational market and other frictions, which will sort of slow down the pace of AI advances in ways that limit the upside, but also limit some of the downsides.

So that's sort of my modal belief that AI will look like a normal technology. But again, I was very recently talking to some people at the firms analyzing AI's progress. I sort of said, "Look, your data seems to suggest that AI might be a normal technology." And even they're like, "no, I don't think so. I think this sounds crazy." So I can't completely discount those views either.

CARDIFF: What are you optimistic about?

ARPIT: With respect to AI?

CARDIFF: Or with respect to the economy, with respect to finance, the courses you teach, anything.

ARPIT: I do think that we'll see these ongoing technological advances, which will have the benefit of extending lifespans, addressing health concerns, and having a lot of benefits to the world.

I was talking to someone at a party and this person was not even like very pro-AI. Very skeptical, but at some point they stopped and said, "well, actually AI did kind of save my life," because there's some health result that they uploaded and stuff like that.

So, the surprising part maybe from that, even after that, they're like, "I'm not like totally sold on AI, but it has a lot of benefits" and health is one of those things where it's easy to discount a health benefit in the abstract if it comes with cost. But when you have a health problem and there's a tool that can solve it, there's nothing more important to you than getting a better solution to it.

CARDIFF: That's great. It does suggest a very powerful negativity bias for somebody to say, "AI saved my life. I still think it sucks." It's fascinating. Last question. I think listeners have gotten a sense of the breadth of your work, and it's

great. I've been following your stuff for a very long time and everybody who's listening should. How do you choose what you work on? How do you choose which topics to focus on? When you're choosing between urban economics and real estate and housing and finance and AI and tech — give us a sense of what pulls you to something and how you make that choice.

ARPIT: I'm not entirely sure, to be honest. This is, I guess, what Tyler Cowen would call your "production function." How do you decide what to work on? It's really just a combination of what am I interested in, what am I naturally curious about and areas that I feel I can have some leverage that I can add something and find a team of other people that are interested in the same thing with me.

So whenever I can find some combination of factors around things that I am interested in and other people are also interested in at the same time, that's usually it. So I'm probably more bottlenecked by finding collaborators and people who I can work with. Maybe I will change that down the road, you know? We'll see.

CARDIFF: I think that's good advice for everybody and a good place to end. Arpit Gupta, thanks for being on The New Bazaar.

ARPIT: Thanks for having me.