

## **Outrunning AI (Before AI Outruns Us)**

Join Greater Perspectives' John Matze and Neil Chilson, former FTC chief technologist, in a riveting discussion on the rise of AI, its regulation, and its far-reaching impact on our lives and privacy. Discover insights on ChatGPT, generative AI, and how industries are transforming under AI's influence.

**This transcript is edited and condensed.**

### **John Matze, Host of 'Greater Perspectives'**

We have **Neil Chilson**, the former chief technologist at the Federal Trade Commission and currently a senior fellow for the Center for Growth and Opportunity.

### **John Matze**

We have Open AI – what, I guess, everyone knows as ChatGPT – and a lot of these AI companies are kind of popping up.

I don't think the technology is necessarily that new. It's just now popular. And I think open AI really made a splash, you know, going out there with that.

Do you have any insight on what the FTC is doing? What do they think of all this? What government agencies in general might be thinking of with AI?

### **Neil Chilson**

Your point is perfect, that there's not a lot new in this technology necessarily. There is a new form factor that people can use it in.

OpenAI made a huge storm when it dropped ChatGPT and that storm has definitely been raging in Washington. Everybody is concerned-slash-excited.

I guess in D.C., it's probably more on the concerned side than the excited side about what this technology can do. Congress has held a ton of hearings on various aspects of the technology. Generative AI in particular, you know, is a sort of infrastructure level technology. Machine learning is. So there will be a ton of applications in a lot of different spaces.

For the FTC, as a general consumer protection and antitrust agency, they are focused on companies making claims about how they're using AI. One of the FTC's jobs is to ensure that companies live up to their promises when they're material to consumers. I think that has been sort of their focus as far as enforcement actions.

They've talked a lot about what they might do in the competition space or in the consumer protection space. But I think a lot of that is very nascent.

The overall frame that the FTC has brought to this has been one of fear and concern and less about the focus on the potential that this has to disrupt existing players.

We saw that Google had this technology. Yet it wasn't until ChatGPT came out that it lit a fire under them and they started really jumping on this as well. That happened to basically all the biggest tech companies.

That alone should be a sign to the FTC that there is a dynamic space and that they need to think about what competition looks like, not just within single markets or products, but how the disruptive outsiders can come in and shock the big incumbents. I think ChatGPT is a great example of that

I hope the FTC takes those lessons— but I'm not that confident that this particular FTC will take that lesson to heart.

### John Matze

What makes you say that?

### Neil Chilson

The current leadership at the Federal Trade Commission is very focused on bigness and has taken a new tack to antitrust.

I say it's a new tack. It's new in the last 40 years, but it's pretty old tack under the previous regime, the previous approach, which is very suspicious of size and is less focused on what are the outcomes for consumers.

When I say size, it's not the size of employees or anything like that but market cap size. Tech companies, even though they have very large market caps – often some of the largest in the world – they don't have enormous employee bases. Compared to Exxon Mobil or Walmart, [tech companies] are not the biggest.

That focus on market cap size has meant that this Federal Trade Commission has focused primarily on just how they get at that size component, without as much concern about whether or not the consumers are benefiting.

What they could learn from the OpenAI space is that it's not just about size.

There are inherent constraints to size, often. In the software space and in social media, there's a lot of returns to scale, but just bureaucratically, there's a lot of dead weight that comes with being a giant company.

That's why Google basically was super cautious about getting this stuff out the door or, you know, they knew it was going to compete with their existing business model so they couldn't push it out the door, right?

### John Matze

Those big companies are extremely slow and afraid of change, typically.

But sometimes that comes with some upside. For example, the quality of Bard – I would say [Google] thought of a lot of things more than [OpenAI] did.

When they released ChatGPT, they just said, you know, 'this is what we have.' And then immediately people are finding all sorts of awful stuff with it. They're getting instructions on how to make bombs from it and all sorts of stuff. They had to patch it.

Of course, they patched it. They fixed it. That's part of being a nimble startup – making a mistake, fixing it, moving on.

But, you know, Bard was a little bit– Google was a little bit more cautious on that approach.

### Neil Chilson

I think that that kind of shows the dynamism.

You have more mature processes that consider legal or liability risks. And you have the 'Hey, let's build something cool' excitement of the smaller group that spurs action from some of those big companies.

That's a competitive process at work. Looking at those big companies and thinking that they're going to just sit there forever without being disrupted...historically has not been borne out.

### John Matze

They will be at some point disrupted. It just happens once you get too big.

If it wasn't for the small startup – I don't think anyone thinks of ChatGPT or OpenAI as being small at this point, but at the time – they forced everyone else to take action.

They forced the market to take a risk and that normally wouldn't have happened organically.

And they brought popularity to something which I think is tremendously helpful

### Neil Chilson

I think they were shocked by it themselves, which is often the sort of case with hot technologies.

### John Matze

Yeah, well, you can't control it once it's out of the bag. You never know where it's going to go.

That all being said, there's all sorts of cool stuff with AI that you could generate. You could auto-generate videos at this point, images. You could auto-generate scripts. You could summarize articles. You can do so much.

The problem, now that people are looking at it, a lot of what it's doing is not that accurate anymore. And so when Chat GPT is asked to do some basic math, it's sometimes making mistakes. It's only learning from what it knows on the Internet. It can also spew things that are not true.

How is the government going to handle regulating that? And what is the proper moral thing for companies to do when they're trying to use AI tools?

### **Neil Chilson**

I think that your second question is like the key one, right? That's what they're all grappling with.

These tools are essentially predicting the next token or imagining an image from essentially white noise and a text prompt.

People should be clear that these are not search engines, right? Like these are not going out there and finding a resource. These are imagination machines.

And I think...if they're properly described that way, I think people would be less surprised when they give something that sounds sort of plausible, but is dead wrong.

That's sort of what they're built to do.

These companies – they're trying to build products that provide reliable results by creating some blend, but they need to be careful in setting the expectations of the users.

ChatGPT has very prominent disclaimers at the bottom of the screen all the time that says 'Look, this stuff may or may not be true. Double check it.'

### **John Matze**

But I mean, a disclaimer is a disclaimer, right? We're so used to there being disclaimers on everything that they now mean nothing.

### **Neil Chilson**

That is a big challenge with disclaimers. The real test here will be a market test, right?

Are these things useful? Are they generating results that are getting people in trouble?

If you're a lawyer using them to generate case citations, people are going to figure it out pretty quickly – like, that's not good.

But they are very powerful.

I think the moral thing to do here is set the expectations correctly, do your best to make the product meet those expectations. It really will differ on the generative AI possibility.

On the image generation side, people are not surprised that what it's making is not accurate.

Sometimes that's the exact point, like kittens who are astronauts. Nobody thinks that it's generating a picture of a cat in space that's accurate.

On the tech side, I think people still have the sort of search engine form factor in their mind.

Getting people to understand that, 'Hey, this thing is imagining things.' It's kind of the average of the internet, that's how I sort of think of it.

**John Matze**

That's a fair point.

**Neil Chilson**

And the internet is not that trustworthy. There's a lot of content out there that's not great, or not accurate. If you think of it that way, I think it can help.

**John Matze**

I've also heard people describe it as a cheap ripoff of a human. It's trying its best to be a human and it's trying to pretend, but it's still a cheap ripoff.

So when it's been asked to make music or create lyrics or do some creative thinking, it just slams a bunch of stuff together.

And it's like, 'Here, this is what I think a human's like.'

**Neil Chilson**

I find it most useful as a complementary tool.

It's a great way to brainstorm, to get it on a raw blank page, get some text on it and then edit. That sort of process, I think, can make it very useful for humans, but I think it really will continue to be a complement.

It certainly is a level-up for people who are less skilled in a lot of different areas.

The art, the images you can generate, are pretty remarkable on something like Midjourney. I couldn't draw that stuff.

You can just turn your imagination into a picture pretty quickly. That's really powerful. People are going to find that really fun and interesting, and it's going to create some new markets. It's certainly going to disrupt some old ones.

### **John Matze**

My personal favorite with the images are the ones where they take children's drawings and use the AI to make them into real drawings.

So it has this like...boxy, weird looking owl and it comes out and it's this horrendous, disgusting, realistic image of what it's pretending to be.

And they're awesome. I've been posting those occasionally.

### **Neil Chilson**

I heard about a guy who, his kids describe their dreams to him and he types them in as prompts and then generates pictures of them. Usually they're pleasant.

### **John Matze**

It's really interesting to see what people are doing with it, and I think we're going to see a lot of people making some serious mistakes with their products. Relying too much on AI and not fact-checking it.

We've already seen some examples of that with people in the news media space trying to have AI-generated articles and then being tremendously inaccurate.

I hope companies in general realize that this stuff is a tool to help people and not a tool that can just live and breathe on its own.

But we'll find out what ends up happening as people start making these mistakes over and over again.

### **Neil Chilson**

There's a sort of tension in the power that the tools have and the fear that people have of it.

You hear two narratives. One is that it's so smart, it's going to cause existential risk to humanity. The other one is that it's so dumb, it can't do math.

And the truth is, it's just different— I keep calling it an 'it,' – I wouldn't even call it an it for the most part. These are algorithms. They're basically statistical models of text or images and they create an output based on an input.

They don't have will. Most of them don't have agency of any kind. So it's a very specific type of artificial intelligence. One that historically is quite old.

The neural network model is very old, it goes back to the 70s.

But I think everybody's sort of surprised that dumping this much information into a model and training it actually creates useful products.

I think people are just surprised by that.

### John Matze

I think a lot of people think it's kind of like magic.

How does this thing work? It's generating outputs that are kind of impressive and it's just doing it on its own with a question prompt.

It's basically statistics and math. So it's just taking words, facts, anything it sees online, shoving it together, creating kind of an average there and spitting it out for you, if I can try to say it as much in layman's terms as possible. It's going to be inaccurate.

What you'd said is that, some people are afraid that it's so smart, it's going to take over the world. Some people are saying it's so dumb that it can't do math.

I think it's the combination of those two. I think it's just smart enough and it's just dumb enough that it could take something and come up with a ridiculous conclusion and want to implement it.

In my mind, if you trust it too much, you're trusting something that's not smart enough to do it yet. It might come up with wildly crazy conclusions.

It might just decide, like the classic case of...humans polluting the environment.

We need to save the environment, so we must get rid of the humans, right? That's the classic really oversimplified version of it.

But I can see it being that dumb.

### Neil Chilson

The challenge there and those types of autonomous agents – they're not really the generative AI stuff.

So there's a loop around generative AI that creates a goal and trying to execute that goal. That part we haven't made that much progress on as far as creating goals and executing them.

AI researchers have been trying to do that for a long time but it's quite different than the sort of statistical modeling of large bodies of text that the generative AI models do now and so well

A lot of this is influenced by the fact that OpenAI decided to release their large language model in a chatbot form.

They didn't have to do it that way, it could have been a plugin where you provide prompts [instead of] a back and forth conversation – more of a programming interface. But the chatbot interface drove a huge amount of adoption.

But it also lets people think of [ChatGPT] as an entity. That model suggests that it's like us, that these large language models are like us in like having goals or being able to execute.

They don't really have a loop like that. Now, you might be able to write like a sort of loop on top of the large language model that has some sort of logic and sets goals and tries to do things and then uses the large language model to communicate or execute that stuff.

But that part up there is still really hard. People don't know how to do that.

And I don't think that the large language model research – it's not really connected to that goal. That's a separate thing.

Sure, large language models spit out crazy stuff, but they don't set goals and they don't really execute right now. You need somebody else to do that. You still need humans basically to do that.

So I'm much less concerned about that.

### John Matze

Throughout history, there's new technology that comes out. Everyone gets afraid of it. Then over time people learn that, oh, this just helps people be more productive.

It just eliminates one extra thing.

For example, the brainstorming process for how to write an essay: ask ChatGPT to create the outline for you. Then you fill it in as a human. It just saved you a bunch of time and energy.

It's a tool that helped you, it didn't replace you.

### John Matze

The things that it can replace are pretty mundane. I think the most hilarious example is how 'South Park' did their episode on ChatGPT. Did you see that?

### Neil Chilson

Of course, it's great.

### John Matze



Basically all the boys in the school were using ChatGPT to respond to their girlfriends and their relationships all started going really well because it was just doing a good job. Then none of them had any idea what their girlfriends were talking about.

It was really great.

### **Neil Chilson**

[Laughs] Right. A misuse and backlash effect. Very well done as usual.

### **John Matze**

I feel like these kind of tools would be really good at replacing basic courtesy emails, stuff like that. Or generic gobbledygook that you're sending back and forth. That's kind of the stuff that I think it's really good at.

### **Neil Chilson**

Increasingly they're sort of on the other end too, where you can ask for it to summarize emails.

There is a way in which maybe this will help us.

You could see a race between [AI] generating these and then it's just LLMs reading each other's emails and summarizing them for us, which might be fine, right?

### **Neil Chilson**

Most people – and certainly policymakers – don't realize how hard it is to anticipate how these things are going to play out in the marketplace over the next five years.

People are so fired up about potential risks and not realizing that the risks that probably are going to come about probably don't look a lot like – there's probably some risks that we're not going to be able to imagine. There's certainly a ton of benefits we're not going to be able to imagine from this technology.

So it's just a very hard space to write big, comprehensive rules in.

It's a much better space to focus on – what are the harms we're worried about? Are we worried about people being lied to? We have laws against people being lied to, whether or not it's AI doing it.

Are there gaps that we need to fill? Let's think about it that way and not try to regulate how people build these things–

### **John Matze**

Like intellectual property rights.

Who owns the things that has been generated from this tool?

**Neil Chilson**

Huge, huge, huge issue – one that has, on the generative AI side, a lot of implications.

There are some really interesting open questions about how these models are trained on content, what are the IP implications of that for the people whose, you know, information was used to train it?

And then on the output end: who owns it?

Or, when you think about creating content that's false, which is then about an individual, who's liable for that content?