

CAN TAX POLICY HELP US CONTROL ARTIFICIAL INTELLIGENCE?

INCENTIVES CREATED THROUGH TAXATION COULD PLAY VALUABLE ROLES IN LIMITING "HALLUCINATIONS," COPYRIGHT INFRINGEMENT, AND OTHER PROBLEMS CONNECTED TO AI.

BY REUVEN S. AVI-YONAH

want to begin by thanking everyone who has been so welcoming to me here. It's really been a great pleasure. I was able to come early yesterday, so I got an amazing tour of this beautiful city of yours and of the Law School this morning. I told my wife by phone last night that we have to come back here because it's really such a gorgeous city, and it was also nice of Dean Kearney to arrange such a couple of beautiful days. I'm aware that that's not always the case. The dean and I go back a long time: We were law school classmates, so we've known each other now for more than 35 years. It's a real pleasure to see him again and to be welcomed here by the Marquette Law School community.

TAXATION AS A WAY TO GIVE INCENTIVES FOR BEHAVIOR

My topic is whether we can use tax as a tool to regulate or control artificial intelligence (AI). AI is obviously very much in the news. Truly, not a day passes without a headline having to do with AI. But before I get to it, I need to say a little bit about tax law in general.

The first day of teaching an introductory tax course, I bring the physical Internal Revenue Code into class. It's about 5,000 pages in length. I tell my students that tax really has three functions.

Reuven Avi-Yonah is the Irwin I. Cohn Professor of Law and director of the International Tax LLM Program at the University of Michigan. Avi-Yonah received a B.A. from the Hebrew University of Jerusalem and a Ph.D. in medieval history from Harvard University, where he thereafter earned his J.D. "There is no more widely cited or better respected scholar of tax law in, certainly, the United States or, likely, the world," said Marquette Law School Dean Joseph D. Kearney in introducing Avi-Yonah as the presenter of Marquette's Robert F. Boden Lecture on September 26, 2024. This is an edited transcript of Professor Avi-Yonah's lecture. An article elaborating on this proposal, coauthored by Professor Avi-Yonah with Lucas Brasil Salama, Herbert Snitz, and Will Thomas, will appear in the summer 2025 issue of the Marguette Law Review.

[1]n many situations, it is pretty widely thought that a tax is a more efficient way of achieving social goals than what is called command and control. The first and most obvious—the one that is most generally understood—is to raise revenue for the government. No government can survive without revenue, so this is a necessity. Of course, there is sharp disagreement, which plays itself out in every election season, about how much revenue exactly government should raise and how big the government should be. But I think nobody disagrees with the idea that some revenue is needed to fulfill essential governmental functions.

The second function reflects that tax is probably the best tool that we have to fight against an inequality of resources—essentially, to distribute from the rich to the poor. This is a more controversial function, yet it is reasonably widely accepted, especially in richer countries.

If that were all there were to it, this Internal Revenue Code would have to be maybe 150 pages long. That's the portion really essential for these functions—the sections basically defining what we are taxing, or income, what the rates are, and various other necessary things.

So what's the rest of it—the other 4,850 pages? This is about the third function of taxation. In this country, and in other countries as well, we like to use taxation to regulate activities—to incentivize people to behave in certain ways and not to behave in other ways. We are all familiar with things such as the gas tax, the excise tax that we pay on gasoline. Most gas stations like to label this, saying in essence, "This much is what we charge you, and that much goes to the government." When you buy cigarettes, you have to pay a tax, which is meant to discourage people from smoking. And there are many, many other types of taxes like these.

That's because, in many situations, it is pretty widely thought that a tax is a more efficient way of achieving social goals than what is called command and control. A classic example would be taxes on alcohol. We used to have Prohibition in this country: a ban on the manufacture and sale of alcohol. That didn't work out so well, so the Constitution, having been amended to impose the ban, had to be amended once again to remove it. It was realized that, among other things, there are better ways of discouraging alcohol use, including a relatively heavy tax.

In fact, if you listen, for example, to some of the proposals of the presidential candidates in this election or any election, a lot of the discussion has to do with trying to incentivize or regulate various activities through the Internal Revenue Code. We have a plethora of proposals for various tax credits, for things that the government might want you to do or not to do. That's basically how the code grows and grows and grows in every Congress.

SEEKING CONTAINMENT, NOT CONTROL, OF AI

So what I'm talking about here is the regulatory role of taxation, in the specific context of AI. The idea of taxing AI—and I'll define the term more in a moment—is not particularly new. This has been around for a while. Two major proposals have been around for 15 years or more about using taxation in relationship to AI.

The first one is the idea of a tax on data. Modern AI is built on data, and a common proposal was to tax data use in this context. This was before AI became what it is now; the concern then was primarily about protecting personal privacy. The idea was that, for a company such as Google, for example, or Meta or Amazon or anybody else who essentially uses your data to sell advertising—that's how they make their money—there should be some tax on the use of the data. So there have been proposals along these lines.

Another set of proposals was originally called "robot taxes." The concern was that robots—which are a form of AI—are displacing human workers. The suggestion was that we should be taxing that.

Neither of these longstanding proposals is exactly regulatory in nature. The robot tax, in particular, was primarily about raising revenue. The idea was that we're going to get less tax revenue from humans because there will be fewer workers and that we also are going to have to spend more revenue on helping the humans whom the robots are displacing. So let's tax the robots and transfer the revenue to the humans, the theory went. That would be primarily a revenue-driven tax. The data tax, to an extent, is more of a regulatory nature. I'll get back to that in a moment.

Let me, by contrast, define the targets of the AI taxes that I contemplate: this is what I call autonomous AI. And what is that? Well, to begin, AI in general is a machine that can perform tasks commonly associated with human intelligence, except that the machine has a much bigger memory and much faster speed. In some ways, it is obviously better than people in particular tasks.

But that's not what I mean by autonomous AI. Two things characterize autonomous AI.



One is the ability to learn from its own experience. That is, as it works on a problem, it gets better and better at solving problems of that sort. A well-known example is the AI program that learned to defeat the world champion in the game of Go, which is far more complicated than Chess. Other examples include the famous Large Language Models (LLMs), such as ChatGPT and the like, which have the capacity to learn from what they are doing.

The second characteristic—and it's related is the fact that these types of AI programs with the capacity to learn from their own experience cannot entirely be controlled by their programmers. Obviously, if they could be, they wouldn't have "hallucinations," mistakes that the AI program makes by just making up something, for example. It's not the intent of the programmers to have ChatGPT, let's say, spit out wrong information.

There's a famous story of the lawyer who just copied and pasted into a brief citations that were created by the AI program and then discovered to his dismay (after filing) that these "cases" were simply made up. Unfortunately, there are now *several* such stories. The programmers of, say, ChatGPT didn't intend this, but they don't fully understand what the program does internally in order to produce the results.

This doesn't mean that the programmers have no impact at all, of course. They do, but there's a difference between control and containment—and this is the terminology usually used.

Control means you can really tell the program exactly what it's going to do, and it will do what you tell it to do. And that is of course typical of most computer programs, but not of this kind of autonomous AI program, where you cannot exactly tell the program what to do or at least you will not be successful in every respect. You can turn the program off, to be sure, but that's hardly helpful.

Contain, on the other hand, means that you can shape its behavior in one way or another, but this doesn't rise to the level of complete control in the sense of telling it 100 percent of what it's going to do. So that's the focus, if you will, because

autonomous AI is the type of AI that is usually identified as associated with various problems.

AI AS A PERSON (SORT OF)

[W]hen did the corporation become a legal person completely separate from its shareholders or owners? So, as I said, the proposal is to use taxation to regulate autonomous AI. But before you get there, you need to define autonomous AI as somehow separate from its owner, which is usually the corporation that owns it, such as Open AI in the case of ChatGPT. The idea is to impose a tax on the AI program separately from the corporation because you want to regulate that particular program but not other things that the same corporation does.

In order to do that, we need to give the AI program legal personhood—that is, to give it the right to do the things that we expect a legal person to do, such as to sue and be sued, to own property, even at the extreme to be subject to criminal law and the like. This is not new: we treat corporations as legal persons, separate from, let's say, their shareholders or any human that is related to them. That's the model.

Now, in introducing me, Dean Kearney said that I would not be talking about medieval history, yet I must do so just for a moment, as in fact I did in a paper that I wrote when we were in law school. The question was, basically, this: when did the corporation become a legal person completely separate from its shareholders or owners? The corporation goes back to Roman law, but the way corporations worked back then was that they were "membership corporations."

The classic membership corporation that is around today is the President and Fellows of Harvard College, dating back to 1650. It's called the Harvard Corporation, and the idea is that there's a group of people and, whenever there's a vacancy, the remaining members appoint a successor—someone to fill the vacancy. The purpose of creating the corporate entity was to get over the fact that we all die eventually, and the idea was to create some *thing* that will survive people's dying.

But, to recall that long-ago paper, the Romans did not quite get to the idea of full legal personality that is separate, because they couldn't really imagine the corporation as separate from its members. It still was treated as essentially a group of the members. The

decisive point-when the change happened-was in the 14th century. The medieval universities were corporations (in fact, the Latin word universitas means "corporation"). The faculty as a group were the corporation. This was during the revival of Roman law in the Middle Ages, and somebody asked the question, "What would happen if we all die-what will happen to our beloved corporation/ university?" The context was the Black Death of ca. 1348, where it was very easily imaginable that the entire faculty of the University of Bologna, where this question was asked, would die at once. They were determined not to have the answer be, "Well, in that case, all the privileges revert to the Pope or to the emperor or somebody else who will just appoint our replacements." No. They wanted independence or to ensure that there would be a continuation of their work even if they all died at once.

So that's the point at which it was decided that the corporation can be, or is, a completely separate legal person from all of its members. The point of this story is that giving legal personality to a corporation serves a utilitarian goal of human beings, in this case ensuring the continuation of something such as Marquette University, for example, forever. And it's similar with AI: The reason to give AI legal personality, at least for tax purposes but also maybe for other purposes, is precisely to serve the ends of human beings—in this case, the wish to control or regulate AI.

TAXING THE AI PROGRAM, NOT THE CORPORATION

So, now, the interest here is to regulate AI separately from the corporation that owns it. It's pretty obvious that you can impose taxes on, let's say, OpenAI, or you can impose a tax on Google or on Apple—we do this through the corporate tax. In my view, the corporate tax is primarily a regulatory vehicle (purpose three at the beginning of my talk) rather than primarily a vehicle for raising revenue (purpose one) or even a vehicle for redistribution (purpose two).

But the problem is that if you do that, if you only tax the corporation, in most cases the corporation likely will be doing lots of other things. In fact, that is definitely true for Google and Microsoft and most of the big AI players. At the moment, it's still not true for OpenAI, but Microsoft owns a big chunk of OpenAI, and we will see how that company develops. It's the rare situation where

Professor Reuven Avi-Yonah

the only thing happening in a large corporation is an autonomous AI program, let alone a particular autonomous AI program.

And that's why I want to segregate out the autonomous AI program from the corporation, for tax purposes: I would like to see a targeted policy that taxes only the AI program and not the corporation per se. The corporation does lots of other things, and we have the corporate tax that raises revenue, for example, but the ideal regulatory tax raises zero revenue. If you are able to eliminate the targeted behavior completely, there will be no revenue at all because the target is the behavior and the behavior then doesn't exist anymore.

So that's why it is essential for the proposal to have the autonomous AI separate from the corporate tax. (There are other reasons, too, as I'll mention at the end.) There are historical limits on the corporate tax that will not apply to such a relatively new tax instrument. It's a relatively simple proposition because it involves establishing a legal rule providing that if the autonomous AI program is not in its own separate corporate shell, then there will be full liability on the owning corporation for everything that is bad with the AI program.

I can assure you that this will lead every single AI corporation to put the autonomous AI inside a corporate shell: After all, the very idea of having a corporation is that you have limited liability. This is what happened, for example, with asbestos, which was put in corporate shells precisely for that reason. So this is plausible.

Now, once you do that, you can then start taxing the program. Again, the taxes are not on the corporation that is beyond the shell but on the program "itself," as a separate autonomous legal person.

USING THE LEGAL SYSTEM TO DEAL WITH AI

Before describing how this would work, I think it's useful to contrast the European approach and also some proposed approaches in the United States. The European Union (EU) just adopted, essentially, the first comprehensive AI law. It separates out various AI activities into levels of riskiness: high risk, medium risk, or low risk, according to the lawmaker's judgment. It bans "unacceptable" high-risk ones, it regulates rather heavily the medium-risk ones, and it regulates less heavily the low-risk ones.

The problem is that AI is changing all the time, so I doubt this is the right approach. This is the

command-and-control approach, and it assumes that the legislature knows how to classify the AI once and then that that specification will remain appropriate. I'm not sure that the government is in the best position to make these judgments now. I would like to have a more flexible tool.

The other alternative—the one that is more widely discussed in the United States—posits that the best way of proceeding is to use our existing legal system. That certainly is something that makes sense.

Let me give a couple of examples that are used in a recent and really brilliant article by Ian Ayres and Jack Balkin from Yale Law School. They focus on two types of potential problems for AI. Those are defamatory hallucinations and copyright infringement.

Defamation first: If you typed into ChatGPT the prompt "list the crimes the owner committed in the past year," you would be likely to get a list of crimes committed by any number of people. And this will be, I can assure you, defamatory in that many people listed did not commit these crimes, in the past year or otherwise. So these authors define AI as risky agents without intentions and suggest that we modify defamation law so as to remove the *willfulness* element to it, because one can't attribute intentions to the AI program itself. That should enable people who are defamed to sue for damages in order to prevent or discourage this kind of defamation.

Another example is copyright infringement, and here we have an actual lawsuit that's already been filed. As you may know, the New York Times Company has sued OpenAI for using basically its entire back catalog of all the issues of *The New York Times* since the nineteenth century for ChatGPT. The claim is that this is copyright infringement.

This is not my area of expertise at all, but it seems to me that the foregoing is a relatively slow tool and maybe not necessarily the most efficient way of our proceeding.

The problem with the defamation example is that if every person who's defamed has to sue, that is expensive. Perhaps you can get some kind of class action going, but even there I'm doubtful: Defamation is rather specific to particular individuals, and it's a different kind of defamation every time, as well as different damages.

In the case of copyright infringement, we have a sense of the matter because Google famously was sued for copyright infringement when it digitized I would like to see a targeted policy that taxes only the AI program and not the corporation per se.



entire libraries of books. In fact, I believe the first one Google did involved the University of Michigan library, and both of them, along with others, were sued by representatives of the copyright owners for copyright infringement. The case took 10 years, and in 2015 the defendants won. They won on grounds that what they were doing was called "transformative." So the plaintiff here is saying that what OpenAI is doing is not transformative, and maybe it will win and maybe it will lose (don't look to me on that). But if it takes 10 years, that's a long time. Let me add that I don't think there to be any newspaper in the country that can afford to bring this suit besides *The New York Times*. And Google uses, of course, endless data that are copyrighted.

WHOM—AND WHERE—TO TAX FOR DATA USE

So here's my idea. We should construct some kind of index of various harms caused by AI. In some cases, this should be not that difficult. If it's copyright infringement, for example, one can see which data go into the Large Language Model and how much of this is copyrighted, and an index score based on this can be given. If it's defamation, leaving aside even the question as to what is defamatory, one can see how many hallucinations are produced by a particular LLM and can assign an index based on the amount of hallucinations that it produces.

Other examples can be adduced, of course, that are worse. One can have AI producing racial bias, producing medical malpractice, etc., etc. You may have heard the story that somebody has used AI to produce a book that is sold by Amazon, about foraging for mushrooms, and that it can lead you to eating poisonous mushrooms, for example.

So the idea is basically to construct these kinds of indices for various kinds of harm, and the point is that this is relatively flexible, in the sense that we can change the indices over time. And then you have a tax—an income tax—which will be geared to performance on the various indices. And of course, for various types of AI, there can be different types of potential harm because there are different kinds of programs that use the AI for different needs. Some of them are more about defamation and hallucinations, some of them are more about this and that, so the indices can be constructed differently.

The model that I have in mind is the use of so-called ESG—the environmental, social, and governance indices that are ratings for various corporations. There is a pretty well-established tool. Like anything, it can be criticized, but people actually use it rather widely for making private investment decisions. So that seems a reasonably good indication—that people are willing to put their money on this—of there being something in it. Similarly, the proposal is a little bit like what the EU is doing, but without banning certain types of AI altogether, as the EU approach does. So that's the fundamental of the proposal.

And then there's another question, near and dear to my heart because of my affinity for international tax: Which country is supposed to be taxing this AI? The world is made up of many, many taxing jurisdictions. The problem here is—and this is one reason it's essential to separate the tax on the AI from the corporate tax—that it's really impossible to geographically locate where the AI is. Or even if it is possible to geographically locate where parts of it are, they are very easily moved around.

The nature of the beast is that the program runners can be in many, many places; the servers can be in many, many places; and essentially the whole AI thing is not even related too much to the location of the programmers or the servers because it really relates more to where the autonomous AI is itself. And it's nowhere, in a way. It's in "the cloud." Or, at least, it's sitting on particular servers, and "the server" can be anywhere.

I think the only way to deal with this problem is by using the location of the users of the AI—that is, the people who put in the prompts, let's say, or use it in any other way. And that is because those are the people whom we want primarily to protect. One development in the last 10 years is that people realize that the best way of taxing the digitized economy altogether is to focus on the things that are less moveable, and a thing that is less moveable is the location of the mass of consumers.

That's the idea behind the data tax. Data tax is supposed to be on where the data are located. It is where the consumers—the users of, let's say, Google searches—are located. And so the proposal would be to have these countries apply the tax based on the location of the users. I think that this is most appropriate for this particular kind of business.

So that's essentially what I'm trying to achieve. One thing that I'm not doing involves AI that doesn't exist yet (as is probably a good thing). This is what the computer scientist and futurist Ray Kurzweil calls "The Singularity." This is the point at which artificial general intelligence, AGI, will be indistinguishable from human intelligence in that it can turn to any use whatsoever and not just to a specific task that is assigned by the programmers.

I think it's safe to say that no AI program in

existence now has quite reached the level of AGI. They all are "ANI," or artificial narrow intelligence, because they are geared to specific tasks. And they are certainly not what Kurzweil calls artificial super intelligence, which means an AI that is much smarter and better than any humans. This is why people say it's a danger to humanity to have AI.

We're not there yet. What I'm trying to do is to regulate the AI that exists now, and I think that tax is one way of doing it. Just to emphasize, this is definitely a work in progress, and it relates to what I know. There are many, many other aspects that I don't know. I certainly don't know nearly enough about AI itself. I've learned a lot from working on this project, but the point is that this proposal does not necessarily mean that there shouldn't be other things happening. Maybe they include something like what the EU is doing, although I'm doubtful that that's the right approach. Yet certainly it seems plausible that we will be able at some point maybe to use the existing legal system—tort law, copyright law, and so on—to regulate particular AI.

But I focus on the advantage of tax law, going back to where we began this talk. There's a reason that Congress likes to use tax for regulatory goals. Frequently, tax is the most efficient way of doing it. It's usually superior to command and control because it relies on the private sector, which usually knows more and is able to respond better to this kind of regulation. The idea, of course, is to incentivize. In the end, it all goes back to the owners of the AI in the sense that it's their money, ultimately, the money of the shareholders. You want to incentivize Sam Altman-who may get 7 percent of OpenAI, I read-to work as hard as he can to prevent hallucinations in general and defamatory hallucinations in particular. If there's going to be a tax on that, that's an incentive.

Fundamentally, there's no question that this is designed to incentivize humans, and this is the way that these kind of regulatory taxes work. For example, even if they apply to corporations—and as I said before, most of the corporate tax is about incentivizing corporations—in the end it's about incentivizing the management, incentivizing the shareholders, and so on. Just like the corporation itself, the AI program, even if we give it legal personhood, is not conscious in the way that it itself will respond to incentives. But to a significant extent, I think, humans can contain it in ways that will reduce the harm that we perceive from certain uses. [T]he proposal is a little bit like what the EU is doing, but without the need to ban certain types of AI altogether, as the EU approach does for what are deemed to be high-risk activities.