



**The Brookings Institution
The Brookings Current podcast**

“Understanding Pope Leo’s AI encyclical”

June 4, 2026

Participants:

Elham Tabassi
Director, Artificial Intelligence and Emerging Technology Initiative
Senior Fellow, Global Economy and Development
The Brookings Institution

Nicol Turner Lee
Director, Center for Technology Innovation
Senior Fellow, Governance Studies
The Brookings Institution

Valerie Wirtschafter
Fellow, Foreign Policy, Artificial Intelligence and Emerging Technology
Initiative
The Brookings Institution

Episode Summary:

Pope Leo XIV’s first encyclical is a 42,000-word Vatican treatise that elevates “the protection of the human person in the age of artificial intelligence.” Nicol Turner Lee, senior fellow and director for the Center for Technology Innovation, and Elham Tabassi, senior fellow and director for the AI and Emerging Tech Initiative, spoke to Valerie Wirtschafter, fellow in Foreign Policy, about the implications of this groundbreaking document.

TURNER LEE: People have to realize that AI is a very extractive technology. A common analogy that I use, which is something John Lewis actually coined: “When you’re not in the kitchen, you’re on the menu.”

TABASSI: It’s not the perfect that we are aiming for. It’s the struggle to be better that makes us human. And think about how we use AI to makes us more human, rather than gives us that perfect solution answer in a record time.

[music]

WIRTSCHAFTER: Hi, this is *The Brookings Current* from the Brookings Podcast Network. I’m Valerie Wirtschafter, a fellow in Foreign Policy in the AI and Emerging Technology Initiative here at Brookings. I’m joined by my Brookings colleagues, Nicol Turner Lee, a senior fellow and director for the Center for Technology Innovation, and Elham Tabassi, a senior fellow and director for the AI and Emerging Tech Initiative.

We are here today to discuss the “Magnifica Humanitas,” or “Magnificent Humanity,” Pope Leo XIV’s first encyclical. With this document, the Catholic Church has planted a flag and established a strict moral baseline for the future of artificial intelligence that centers human dignity. Since its release last month, the 42,000-word theological document has been clipped and debated across the AI community. How do the pope’s warnings against the technocratic paradigm collide with the hard realities of code, capital, and civil rights?

Nicol, welcome to *The Current*.

TURNER LEE: Oh, thank you.

WIRTSCHAFTER: Elham, welcome to you as well.

TABASSI: Thank you for having me.

[1:32]

WIRTSCHAFTER: Nicol, let’s start with you.

TURNER LEE: Oh, I thought you were gonna do that.

WIRTSCHAFTER: When “Magnifica Humanitas” dropped, what was the immediate reaction across civil society and the digital equity space?

[1:43]

TURNER LEE: So I appreciate that, so I’m gonna just say that I’m not Catholic, just so you know. But I read the document with much interest because of this idea of, you know, a power conglomerate like the Vatican really giving some type of statement on artificial intelligence in line with previous popes. So I thought that this was very timely, you know, given the fact that we’re seeing AI accelerate in such a manner that perhaps we might be losing sight of the humanity side of it.

So I thought it was interesting from this, as a sociologist in particular, that the church would have this really big moral compass on the extent to which we have to look at this tool and its potential to embolden existing civil rights concerns, as well as, you know, increase the amount of surveillance, the loss of labor, and other areas which I know my work here at Brookings is really concerned about.

So I was at first thinking maybe he read the work at Brookings and, you know, he got a head start. But most importantly, that this has served as a guide, I think some guidance for people who are following this area to really consider what does it look like to have a technology that really embraces the spiritual component in a way that values humanity and human civilization.

So I found it to be interesting. I think the reaction from civil society has been quite the same and also diverse with regards to opinion. But I do think it was a timely document to come out at a time where we're seeing so many challenges about when it comes to, you know, the safety of AI and the concerns that everyday people have around this technology.

[3:13]

WIRTSCHAFTER: And Elham, what about the reaction from the technical community? He talked a bit about kind of the Tower of Babel metaphor. How does something like that land with the technical community, and where do they see their responsibility?

[3:25]

TABASSI: I think the response from the tech community, there has been a lot of silence, so we we can talk about this. But, just like Nicole, I'm not Catholic, but I think these are important conversations and important documents for all of us to look at as they set sort of a north star.

So Pope Leo anchors the AI encyclical onto two images from scripture and probably two images in different way. One is Tower of Babel that you mentioned, and as everybody else knows, in Genesis humanity settles on a field, on a plain, and then wants to build a city and a tower, a tower so tall that it reaches heaven, so secure power. In Leo's reading, perpetuate a name.

And I think it is that ambitious and in the word of the pope, a single language, a single technology, a single direction that make it a Babel for him. And he talks about that that uniformity instead of communion is choosing homogenization.

So it's not a sort of condemning or concerning about the tower or how tall it is, uh but that the tower is being owned and directed from the top and diversity being engineered out.

And that is sort of the contrasting image where the other image is Nehemiah, that rebuilding of the Jerusalem wall after the Exile. And that wall is being built in sections by people, by communities, by different people. So you see it as a distributed, cooperative, sort of accountable to the people that they have built it.

So it provides these two images, that the act is the same, it's constructions, building, but the structures is very different. One is concentrating and homogenizing, and the other one is distributed and human scale.

And I think that the point here is that it's not about can AI build or build AI because it's happening, but it's all about this approach how this is gonna be done.

The the last thing I will say is that it was also interesting to me, and I think everybody also picked on that, that it was signed on the 135 anniversary of church's response to the first industrial revolution. So I think all of these things are things for us to look at.

[5:47]

WIRTSCHAFTER: Yeah, and I think those are both great points. And, you know, curious, we we hear a lot in in the broader tech policy space, but also specifically AI, "AI for good." Does this kind of align with that vision, the vision emanating from Vatican City is sort of aligned with this AI for good?

And what what do you you make about that in terms of the way that industry kind of brands that movement? Is there alignment or is this sort of you know, truer version or whatever kind of conversations happen around this sort of ethics of it?

[6:18]

TURNER LEE: I mean, well, I'll start. I mean, as Elham has mentioned with regards to the scriptures of the Bible, I think we all read the Bible, and we hope that it's all true, right, in terms of our forgiveness and where we land up in terms of the other side. And I think we take the Bible for its principles, right, in in many respects and its stories that allow for these type of redemptive characteristics.

Again, the Bible is one. You know, who reads the Bible is another conversation in terms of what their moral compass is. I mean, I think what the pope tried to do in this document is to provide, again, that moral compass for tech companies to start paying attention to how powerful AI has become and the extent to which its nature could actually, you know, just destroy the world in which we know it, in particular by, as Elham has said, creating these divided walls and allowing people to become more subjects of the technology.

I love the use of the words "colonization" and the way in which he constructs this power conglomerate that is essentially taking over humanity. And I think these are real concerns. I mean, I think people have to realize that AI is a very extractive technology. It's one in which, you know, it takes your data. It it looks at, areas that that data is combined and conjoined with other attributes that you have in the digital space. And it uses that in a way that I think people do not have agency.

And and when I read this document, it sort of seconded for me a common analogy that I use when I speak to audiences, which is something John Lewis actually coined: "When you're not in the kitchen, you're on the menu."

And so in many respects, you know, this technology as the pope has outlined, is one in which without the appropriate guardrails, you know, not necessarily just the

nomenclature of public good, but some real solid guardrails and some directives on how to ensure that this is going to be a peaceful, normative, inclusive, ethical technology, that it's probably not gonna be that.

And so for me, you know, the AI for good is it's a banner, it's a campaign. But unless you have the right people who are orchestrating what that good looks like, which is what I think this document tries to redirect people's aspirations towards, Valerie and Elham, I think, you know, it lands up in the same place where we are today, being very extractive and dehumanizing without any recourse of individuals to have agency over what the technology can do.

[8:31]

TABASSI: Yeah, I think, yes, intent matters and how we actually go and develop technology that reaches those goals and how do we put the safeguard, but also test them and make sure that they are doing what they're supposed to do. One point that I want to make is here that nowhere in the, at least my readout of this, is that pope is not saying that AI is the enemy. I think he uses the words such as AI is neither "a force antagonistic to humanity" nor "inherently evil," and then immediately also talk about that it's also "never neutral." And I think that tension of if we know that it's not neutral, we know that, you know, at least, you know, from from a technologist linguistic zero error, zero risk technology doesn't exist.

But at the same time, how to achieve those promising beneficial use of AI to uplift the society as it's changing the society. And that is, I think, something that a lot of people from the civil society to tech area are working on.

[9:28]

WIRTSCHAFTER: So let's move a little from the why to the how, because I actually think the how is much harder. Starting with you, Elham, you know, the document demands human-centered AI, and putting that into systems in practice is difficult. We see a lot of effort to do so. But I'm curious from your perspective, how do you code that kind of moral imperative? How do you build that into these systems?

[9:50]

TABASSI: Yeah, that's a very good question and something that I have been working on it for several years. The community has been working on this for several years. I think it's important to keep reminding ourselves and and knowing that AI is all about context of use. It's not just a technology in isolations. And AI is sociotechnical. It's not just data, compute, and algorithm. It's a complex interactions of data, compute, and algorithm with the environment that it is operating within, the human that operates it, the human that gets impacted by that.

So in a way the solutions are not just technical and technological. AI is not just a technology problem, the problems that you're talking about, making sure that it's serving the whole humanities. It's not a technology problem with a governance footnote. It's really a governance problem that's that's very much entangled with technology and technical problems.

So, the way to do this is that we need to bring communities together, the people that build the technology, the sociologists that look at the how the technology can actually serve humanity and serve the society. And if you really want the AI to serve the society, we should include the whole society in in development and advancement of the technology.

[10:57]

WIRTSCHAFTER: So in terms of kind of right now, and I know you hit on this a little bit, but I want to draw you out maybe a bit more because I know you talk about this a ton. Are we watching a new form in the current state of development, a new form of this digital colonization?

TURNER LEE: Oh, yeah. I mean, listen,

WIRTSCHAFTER: is that an easy one?

[11:13]

TURNER LEE: I know, I, you know, I I have tried to, like, work through what Elham is talking to this, like, this space, right, where you're trying to balance the trade-offs, because you do want AI. I mean, I think that it's very implicit in the document that AI will have this, like, tremendous and transformative capacity, particularly when I look at spaces like health care. Right? I mean, health care is probably one of those areas which will be primed for the type of exploitation of AI tools to be able to get to faster scientific discovery.

But I think, you know, we still have other issues. Right? I mean, one of the things that came before this that I think is worth mentioning is that the pope also apologized for slavery. Right? And what happened with colonization under a slave regime and how people, Black people and African people, were affected by that.

The challenge with the apology is that the damage is already done. Right? And so you have this situation where we're trying to correct the moral compass of AI, but you still have people who are disproportionately impacted by algorithmic decisions, you know, whether or not they get credit or they don't get credit, whether or not they qualify for a program that is being sponsored by a hospital or not.

We continue to see these examples where technology has still become a tool that, you know, applies itself to systemic discrimination and racism in ways that we have to take a look at who sits at the table. So colonization starts when you have people at the table, particularly who do not represent, as Elham has said, the context or the lived experiences of people creating these tools. Right? So you automatically start with a beta where the tool is being built for people and not necessarily with their experiences in mind. So I think, you know, for me, this is probably what many people want to bring into this government that we have here with regards centering people, centering consumers.

But it's also, as we've discussed, a moral compass, which, you know, unfortunately, is very hard to program. I mean, code encodes bias. Code can also encode good. But it, for the most part, encodes bias because the society in which that data is being

trained is already traumatized. So if you start there, you know, as we all know, you're actually starting with systems that begin with a deficit versus an asset.

And so, you know, I like to say I woke up and I looked at this document and I said, "Oh, wow, amen!" And then I turned around and realized, you know, that there are other vulnerabilities that we have to face, both on the safety and security side, as well as the consumer protection side, that we just have not resolved because the right people aren't sitting at the table creating the policies or the products.

[13:38]

WIRTSCHAFTER: So maybe quickly, and then I want to move on to kind of the policy side, because you hit a little bit on the Washington piece, and obviously there's been a lot of action in Brussels as well. But are there other kinds of technical standards that can be implemented to sort of prevent some of this dynamic? What are the limits there? What else do we need in place here?

[13:56]

TABASSI: And and to that, I will say that getting to all of the goals that Nicol was talking about, there's at least two components to that. One is building the technology that inherently trustworthy. You know, you have heard the word secure by design, privacy by design. But we should think about how to build a technology that's trustworthy by design.

And I think there is a lot of room for improvement, research, and creativity in how to build technologies that by design are secure or private or fair. And, you know, we have learned a lot about the things that we should not do.

So that's about building the technology that's trustworthy, that does the right thing.

And then also build around that the policies, the right guardrails, the the right practices for responsible use of the technology. And they go hand in hand.

You ask about the limit of the standards. I don't think we have yet even actually done enough to know the limits. I think there's a lot that we can do, and it's important for all of us to come to work on that.

[14:55]

TURNER LEE: And if I can answer that, I mean, I think it's also interesting the setting in which this was actually delivered. You know, there was Anthropic sitting next to the pope as he read that, who's actually about to go onto the private equity market. You know, I think they're IPO now, about to be worth trillions of dollars.

I mean, my point is, like, it's one of those things, and I think this is very much part of the biblical experience, where, you know, the Bible is written for those, again, who require or need some type of a transformative experience, redemptive experience, but yet we're delivering this with a company that may not have the same goals.

And the Vatican itself is not necessarily poorly resourced when it comes to pushing an agenda, which I think is appropriate, again, for the delivery of this document. But

it also suggests that they could also make a stand in terms of deplatforming some of these companies that do not stand up to some of the ethical concerns that they have.

I mean, I think at the end of the day, Elham is right. I mean, there are things that we can do. We can definitely code to compliance; we can ensure that the right actors are at the table; that we're getting feedback from consumers and validating these products and services. We can ensure that there is recourse for people who get harmed and are unable to reverse a decision that may be made by a faulty algorithm.

But at the end of the day, we have to have the will, the political will, and the will of civil society come together to talk about that. And Elham and I both know that that's a Herculean task.

And you know, Valerie, you know, just given where AI is showing up. You know, I tell people one day I woke up and I was looking at my computer, and the next day Copilot is editing my sentences. And I never loaded her, you know?

And so part of it is, like, we're now in this space where AI is becoming so entrenched in everything that we do that I think it's even harder because it's so much more fragmented to get everybody to the table to have these conversations that I think that the pope outlines in his document.

[16:40]

WIRTSCHAFTER: So let's talk governance, because we've been dancing, I think, a little bit around it. Obviously, there are different approaches. We've seen in a European context versus what's happened in Washington. Just, you know, I think 20 minutes before we taped this, the White House released their recent executive order on AI systems. So two questions, I think. Does the pope's comments, does this document have any bearing on how these governance conversations shape out? What does it do to kind of fit into a moment like this? Does it not have much of an effect on the governance conversation? Maybe Nicol, I don't know if you have thoughts there.

[17:18]

TURNER LEE: That just came out, so for people who are listening, you know, the president's second part of his executive order is basically looking at frontier models and the extent to which they should be validated by the government before they go to market. Is that about right?

[17:30]

TABASSI: Yeah, they asked for a voluntary framework for the models to be tested, but I don't know if, because it's voluntary, I don't think there is any teeth on stopping a model from release or—

[17:39]

TURNER LEE: Exactly. Yeah. So, you know, I think obviously what the initial reaction by the vice president and others in government were essentially, like, this is a great document. I think JD Vance put it, this is what the pope is supposed to do, and so, you know, expected.

I mean, I think at the end of the day, I think what this document combined with the variety of conversations Val you're talking about that are happening around the world, is really signaling that something needs to change to make sure that AI in terms of what I think Elham mentioned, there's a pragmatic framework for what responsible governance looks like.

And it was really interesting to me because, you know, this pope is also from Chicago, which I think is really interesting because he's coming with the American, like, hum in his ear. He was real clear, like, we need employment laws to work. We need, you know, these various laws on the books to work. That's actually in place for people who are listening. These things are supposed to govern these models whether explicitly or implicitly.

But with regards to what governance can look like going forward, I think we're still in a state, I think, Elham could agree, that there's so many different narratives on governance that we really don't have one central theme, even with all the information that is out there as to what could work best in a global context.

[18:46]

TABASSI: So yes, I agree. This document is aspirational. It sets principles. It sets a north star for the community to aspire to. It doesn't have the nitty-gritty things that you need for a governance or risk management or evaluations or policy.

I also agree that I don't think that, you mentioned also, you know, Europe, Europe is moving to enforcement. So come this August, the full rule for high-risk systems kicks in, so they're really moving into enforcement.

But I agree. I don't think that this is gonna be a driver for any sort of legislation or anything. And I also agree that what it can do, and and seems like it's also been doing right now, it's sort of influencing the basis of and the foundations of the conversations. A lot of conversations in the U.S. at least, is being dominated by "winning the race" and this race. I'm hoping that it brings the other angle of what is gonna do it for the people?

And and I think that can be a good a good addition to the conversations and and sort of the changing the basis of the dialogue.

[19:52]

TURNER LEE: Yeah, and if I could add on to that, I mean, I think your question about, you know, the White House advancing these principles around, you know, voluntary testing of what could potentially be very dangerous AI models. You know, when you first read that at face value, you're like, yes, we should. Right? We really

do not want technologies that can be used to weaponize itself against other countries or its own people.

But then you also think about, well, yes, that could also happen if there's complete oversight by the government only on these tools, that these tools could be used for greater surveillance. They could be used for greater identification of people that this country has made very clear that they do not think are part of the fabric of our society.

So, you know, I think as Elham has said, the aspiration of the document is real. I mean, I I personally think that this pope is the real deal. Right? I think, you know, he has really come out with some real solid apologies. Like I said, it's been a long time since there's been apology to Black Americans about slavery. And to be able to say that publicly, I still find is very powerful sentiment from the Vatican itself.

But I also think, as Elham said, aspiration is one thing, but actual action is another. And we do have to begin to think about, you know, the preservation of humanity when left to a handful of companies that feel that they can actually create and design a world in which they want to live and not everybody else.

[21:11]

WIRTSCHAFTER: Well, I think that is a great way to wrap, but I want to leave it open to one ... one, any kind of final final comments or or thoughts you wanted to add and that we haven't touched on? Anything?

[21:23]

TABASSI: You know, there is a lot of talk about AI making things more efficient.

Give us, giving us that perfect draft, you know, all of the emails has all of a sudden, you know, improved and they are all nicer and neater. But as I was reading this, and I think there was some other work that some other people has done, shout out to AEI also that has been working on this. And what does it mean to be human? and I don't think it's the perfection. It's not the perfect that we are aiming for. It's the struggle to be better that makes us human. And and think about how we use AI to makes us more human, rather than gives us that perfect solution answer in a record time.

And and I hope that can change all of the angles about, both in terms of what type of technology we develop, but also how we're gonna use it.

[22:16]

WIRTSCHAFTER: Well, Elham, Nicol, thank you both for this great conversation. And to dive deeper into the global future of AI policy, technical standards, and the ongoing work to build an inclusive digital world, please visit Brookings dot edu.

[music]

My name is Valerie, and again, thank you for joining us today.