

Transcript: Mark's Personal Challenge - Yuval Noah Harari

Mark Zuckerberg: Hey everyone. This year I'm doing a series of public discussions on the future of the internet and society and some of the big issues around that, and today I'm here with Yuval Noah Harari, a great historian and best-selling author of a number of books. His first book, "Sapiens: A Brief History of Humankind", kind of chronicled and did an analysis going from the early days of hunter-gatherer society to now how our civilization is organized, and your next two books, "Homo Deus: A Brief History of Tomorrow" and "21 Lessons for the 21st Century", actually tackle important issues of technology and the future, and that's I think a lot of what we'll talk about today. But most historians only tackle and analyze the past, but a lot of the work that you've done has had really interesting insights and raised important questions for the future. So I'm really glad to have an opportunity to talk with you today. So Yuval, thank you for joining for this conversation.

Yuval Noah Harari: I'm happy to be here. I think that if historians and philosophers cannot engage with the current questions of technology and the future of humanity, then we aren't doing our jobs. Only you're not just supposed to chronicle events centuries ago. All the people that lived in the past are dead. They don't care. The question is what happens to us and to the people in the future.

Mark Zuckerberg: So all the questions that you've outlined-- where should we start here? I think one of the big topics that we've talked about is around-- this dualism around whether, with all of the technology and progress that has been made, are people coming together, and are we becoming more unified, or is our world becoming more fragmented? So I'm curious to start off by how you're thinking about that. That's probably a big area. We could probably spend most of the time on that topic.

Yuval Noah Harari: Yeah, I mean, if you look at the long span of history, then it's obvious that humanity is becoming more and more connected. If thousands of years ago Planet Earth was actually a galaxy of a lot of isolated worlds with almost no connection between them, so gradually people came together and became more and more connected, until we reach today when the entire world for the first time is a single historical, economic, and cultural unit. But connectivity doesn't necessarily mean harmony. The people we fight most often are our own family members and neighbors and friends. So it's really a question of are we talking about connecting people, or are we talking about harmonizing people? Connecting people can lead to a lot of conflicts, and when you look at the world today, you see this duality in-- for example, in the rise of wall, which we talked a little bit about earlier when we met, which for me is something that I just can't figure out what is happening, because you have all these new connecting technology and the internet and virtual realities and social networks, and then the most-- one of the top political issues becomes building walls, and not just cyber-walls or firewalls-- building stone walls; like the most Stone Age technology is suddenly the most advanced technology. So how to make sense of this world which is more connected than ever, but at the same time is building more walls than ever before.

Mark Zuckerberg: I think one of the interesting questions is around whether there's actually so much of a conflict between these ideas of people becoming more connected and this fragmentation that you talk about. One of the things that it seems to me is that-- in the 21st century, in order to address the biggest opportunities and challenges that humanity-- I think it's both opportunities-- spreading prosperity, spreading peace, scientific progress-- as well as some of the big challenges-- addressing climate change, making sure, on the flipside, that diseases don't spread and there aren't epidemics and things like that-- we really need to be able to come together and have the world be more connected. But at the same time, that only works if we as individuals have our economic and social and spiritual needs met. So one way to think about this is in terms of fragmentation, but another way to think about it is in terms of personalization. I just think about when I was growing up-- one of the big things that I think that the internet enables is for people to connect with groups of people who share their real values and interests, and it wasn't always like this. Before the internet, you were really tied to your physical location, and I just think about how when I was growing up-- I grew up in a town of about 10 thousand people, and there were only so many different clubs or activities that you could do. So I grew up, like a lot of the other kids, playing Little League baseball. And I kind of think about this in retrospect, and it's like, "I'm not really into baseball. I'm not really an athlete. So why did I play Little League when my real passion was programming computers?" And the reality was that growing up, there was no one else really in my town who was into programming computers, so I didn't have a peer group or a club that I could do that. It wasn't until I went to boarding school and then later college where I actually was able to meet people who were into the same things as I am. And now I think with the internet, that's starting to change, and now you have the availability to not just be tethered to your physical location, but to find people who have more niche interests and different kind of subcultures and communities on the internet, which I think is a really powerful thing, but it also means that me growing up today, I probably wouldn't have played Little League, and you can think about me playing Little League as-- that could have been a unifying thing, where there weren't that many things in my town, so that was a thing that brought people together. So maybe if I was creating-- or if I was a part of a community online that might have been more meaningful to me, getting to know real people but around programming, which was my real interest, you would have said that our community growing up would have been more fragmented, and people wouldn't have had the same kind of sense of physical community. So when I think about these problems, one of the questions that I wonder is maybe-- fragmentation and personalization, or finding what you actually care about, are two sides of the same coin, but the bigger challenge that I worry about is whether-- there are a number of people who are just left behind in the transition who were people who would have played Little League but haven't now found their new community, and now just feel dislocated; and maybe their primary orientation in the world is still the physical community that they're in, or they haven't really been able to find a community of people who they're interested in, and as the world has progressed-- I think a lot of people feel lost in that way, and that probably contributes to some of the feelings. That would my hypothesis, at least. I mean, that's the social version of it. There's also the economic version around globalization, which I think is as important, but I'm curious what you think about that.

Yuval Noah Harari: About the social issue, online communities can be a wonderful thing, but they are still incapable of replacing physical communities, because there are still so many things--

Mark Zuckerberg: That's definitely true. That's true.

Yuval Noah Harari: --that you can only do with your body, and with your physical friends, and you can travel with your mind throughout the world but not with your body, and there is huge questions about the cost and benefits there, and also the ability of people to just escape things they don't like in online communities, but you can't do it in real offline communities. I mean, you can unfriend your Facebook friends, but you can't un-neighbor your neighbors. They're still there. I mean, you can take yourself and move to another country if you have the means, but most people can't. So part of the logic of traditional communities was that you must learn how to get along with people you don't like necessarily, maybe, and you must develop social mechanisms how to do that; and with online communities-- I mean, and they have done some really wonderful things for people, but also they kind of don't give us the experience of doing these difficult but important things.

Mark Zuckerberg: Yeah, and I definitely don't mean to state that online communities can replace everything that a physical community did. The most meaningful online communities that we see are ones that span online and offline, that bring people together-- maybe the original organization might be online, but people are coming together physically because that ultimately is really important for relationships and for-- because we're physical beings, right? So whether it's-- there are lots of examples around-- whether it's an interest community, where people care about running but they also care about cleaning up the environment, so a group of organize online and then they meet every week, go for a run along a beach or through a town and clean up garbage. That's a physical thing. We hear about communities where people-- if you're in a profession, in maybe the military or maybe something else, where you have to move around a lot, people form these communities of military families or families of groups that travel around, and the first thing they do when they go to a new city is they find that community and then that's how they get integrated into the local physical community too. So that's obviously a super important part of this, that I don't mean to understate.

Yuval Noah Harari: Yeah, and then the question-- the practical question for also a service provider like Facebook is: What is the goal? I mean, are we trying to connect people so ultimately they will leave the screens and go and play football or pick up garbage, or are we trying to keep them as long as possible on the screens? And there is a conflict of interest there. I mean, you could have-- one model would be, "We want people to stay as little as possible online. We just need them to stay there the shortest time necessary to form the connection, which they will then go and do something in the outside world," and that's one of the key questions I think about what the internet is doing to people, whether it's connecting them or fragmenting society.

Mark Zuckerberg: Yeah, and I think your point is right. I mean, we basically went-- we've made this big shift in our systems to make sure that they're optimized for meaningful social interactions, which of course the most meaningful interactions that you can have are physical, offline interactions, and there's always this question when you're building a service of how you measure the different thing that you're trying to optimize for. So it's a lot easier for us to measure if people are interacting or messaging online than if you're having a meaningful connection physically, but there are ways to get at that. I mean, you can ask people questions about what the most meaningful things that they did-- you can't ask all two billion people, but you can have a statistical subsample of that, and have people come in and tell you, "Okay, what are the most meaningful things that I was able to do today, and how many of them were enabled by me connecting with people online, or how much of it was me connecting with something physically, maybe around the dinner table, with content or something that I learned online or saw." So that is definitely a really important part of it. But I think one of the important and interesting questions is about the richness of the world that can be built where you have, on one level, unification or this global connection, where there's a common framework where people can connect. Maybe it's through using common internet services, or maybe it's just common social norms as you travel around. One of the things that you pointed out to me in a previous conversation is now something that's different from at any other time in history is you could travel to almost any other country and look like you-- dress like you're appropriate and that you fit in there, and 200 years ago or 300 years ago, that just wouldn't have been the case. If you went to a different country, you would have just stood out immediately. So there's this norm-- there's this level of cultural norm that is united, but then the question is: What do we build on top of that? And I think one of the things that a broader kind of set of cultural norms or shared values and framework enables is a richer set of subcultures and subcommunities and people to actually go find the things that they're interested in, and lots of different communities to be created that wouldn't have existed before. Going back to my story before, it wasn't just my town that had Little League. I think when I was growing up, basically every town had very similar things-- there's a Little League in every town-- and maybe instead of every town having Little League, there should be-- Little League should be an option, but if you wanted to do something that not that many people were interested in-- in my case, programming; in other people's case, maybe interest in some part of history or some part of art that there just may not be another person in your ten-thousand-person town who share that interest-- I think it's good if you can form those kind of communities, and now people can find connections and can find a group of people who share their interests. I think that there's a question of-- you can look at that as fragmentation, because now we're not all doing the same things, right? We're not all going to church and playing Little League and doing the exact same things. Or you can think about that as richness and depth-ness in our social lives, and I just think that that's an interesting question, is where you want the commonality across the world and the connection, and where you actually want that commonality to enable deeper richness, even if that means that people are doing different things. I'm curious if you have a view on that and where that's positive versus where that creates a lack of social cohesion.

Yuval Noah Harari: Yeah, I mean, I think almost nobody would argue with the benefits of richer social environment in which people have more options to connect around all kind of things. The key question is how do you still create enough social cohesion on the level of a country and increasing also on the level of the entire globe in order to tackle our main problems. I mean, we need global cooperation like never before because we are facing unprecedented global problems. We just had Earth Day, and to be obvious to everybody, we cannot deal with the problems of the environment, of climate change, except through global cooperation. Similarly, if you think about the potential disruption caused by new technologies like artificial intelligence, we need to find a mechanism for global cooperation around issues like how to prevent an AI arms race, how to prevent different countries racing to build autonomous weapons systems and killer robots and weaponizing the internet and weaponizing social networks. Unless we have global cooperation, we can't stop that, because every country will say, "Well, we don't want to produce killer robot-- it's a bad idea-- but we can't allow our rivals to do it before us, so we must do it first," and then you have a race to the bottom. Similarly, if you think about the potential disruptions to the job market and the economy caused by AI and automation. So it's quite obvious that there will be jobs in the future, but will they be evenly distributed between different parts of the world? One of the potential results of the AI revolution could be the concentration of immense wealth in some part of the world and the complete bankruptcy of other parts. There will be lot of new jobs for software engineers in California, but there will be maybe no jobs for textile workers and truck drivers in Honduras and Mexico. So what will they do? If we don't find a solution on the global level, like creating a global safety net to protect humans against the shocks of AI, and enabling them to use the opportunities of AI, then we will create the most unequal economic situation that ever existed. It will be much worse even than what happened in the Industrial Revolution when some countries industrialized-- most countries didn't-- and the few industrial powers went on to conquer and dominate and exploit all the others. So how do we create enough global cooperation so that the enormous benefits of AI and automation don't go only, say, to California and Eastern China while the rest of the world is being left far behind.

Mark Zuckerberg: Yeah, I think that that's important. So I would unpack that into two sets of issues-- one around AI and the future economic and geopolitical issues around that-- and let's put that aside for a second, because I actually think we should spend 15 minutes on that. I mean, that's a big set of things.

Yuval Noah Harari: Okay. Yeah, that's a big one.

Mark Zuckerberg: But then the other question is around how you create the global cooperation that's necessary to take advantage of the big opportunities that are ahead and to address the big challenges. I don't think it's just fighting crises like climate change. I think that there are massive opportunities around global--

Yuval Noah Harari: Definitely. Yeah.

Mark Zuckerberg: Spreading prosperity, spreading more human rights and freedom-- those are things that come with trade and connection as well. So you want that for the upside. But I guess my diagnosis at this point-- I'm curious to hear your view on this-- is I actually think we've spent a lot of the last 20 years with the internet, maybe even longer, working on global trade, global information flow, making it so that people can connect. I actually think the bigger challenge at this point is making it so that in addition to that global framework that we have, making it so that things work for people locally. Right? Because I think that there's this dualism here where you need both. If you just-- if you resort to just kind of local tribalism then you miss the opportunity to work on the really important global issues; but if you have a global framework but people feel like it's not working for them at home, or some set of people feel like that's not working, then they're not politically going to support the global collaboration that needs to happen. There's the social version of this, which we talked about a little bit before, where people are now able to find communities that match their interests more, but some people haven't found those communities yet and are left behind as some of the more physical communities have receded.

Yuval Noah Harari: And some of these communities are quite nasty also. So we shouldn't forget that.
<laughs>

Mark Zuckerberg: Yes. So I think they should be-- yes, although I would argue that people joining kind of extreme communities is largely a result of not having healthier communities and not having healthy economic progress for individuals. I think most people when they feel good about their lives, they don't seek out extreme communities. So there's a lot of work that I think we as an internet platform provider need to do to lock that down even further, but I actually think creating prosperity is probably one of the better ways, at a macro level, to go at that. But I guess--

Yuval Noah Harari: But I will maybe just stop there a little. People that feel good about themselves have done some of the most terrible things in human history. I mean, we shouldn't confuse people feeling good about themselves and about their lives with people being benevolent and kind and so forth. And also, they wouldn't say that their ideas are extreme, and we have so many examples throughout human history, from the Roman Empire to slave trade into modern age and colonialism, that people-- they had a very good life, they had a very good family life and social life; they were nice people-- I mean, I guess, I don't know, most Nazi voters were also nice people. If you meet them for a cup of coffee and you talk about your kids, they are nice people, and they think good things about themselves, and maybe some of them can have very happy lives, and even the ideas that we look back and say, "This was terrible. This was extreme," they didn't think so. Again, if you just think about colonialism--

Mark Zuckerberg: Well, but World War II, that came through a period of intense economic and social disruption after the Industrial Revolution and--

Yuval Noah Harari: Let's put aside the extreme example. Let's just think about European colonialism in the 19th century. So people, say, in Britain in the late 19th century, they had the best life in the world at the time, and they didn't suffer from an economic crisis or disintegration of society or anything like that, and they thought that by going all over the world and conquering and changing societies in India, in Africa, in Australia, they were bringing lots of good to world. So I'm just saying that so that we are more careful about not confusing the good feelings people have about their life-- it's not just miserable people suffering from poverty and economic crisis.

Mark Zuckerberg: Well, I think that there's a difference between the example that you're using of a wealthy society going and colonizing or doing different things that had different negative effects. That wasn't the fringe in that society. I guess what I was more reacting to before was your point about people becoming extremists. I would argue that in those societies, that wasn't those people becoming extremists; you can have a long debate about any part of history and whether the direction that a society chose to take is positive or negative and the ramifications of that. But I think today we have a specific issue, which is that more people are seeking out solutions at the extremes, and I think a lot of that is because of a feeling of dislocation, both economic and social. Now, I think that there's a lot of ways that you'd go at that, and I think part of it-- I mean, as someone who's running one of the internet platforms, I think we have a special responsibility to make sure that our systems aren't encouraging that-- but I think broadly, the more macro solution for this is to make sure that people feel like they have that grounding and that sense of purpose and community, and that their lives are-- and that they have opportunity-- and I think that statistically what we see, and sociologically, is that when people have those opportunities, they don't, on balance, as much, seek out those kind of groups. And I think that there's the social version of this; there's also the economic version. I mean, this is the basic story of globalization, is on the one hand it's been extremely positive for bringing a lot of people into the global economy. People in India and Southeast Asia and across Africa who wouldn't have previously had access to a lot of jobs in the global economy now do, and there's been probably the greatest-- at a global level, inequality is way down, because hundreds of millions of people have come out of poverty, and that's been positive. But the big issue has been that, in developed countries, there have been a large number of people who are now competing with all these other people who are joining the economy, and jobs are moving to these other places, so a lot of people have lost jobs. For some of the people who haven't lost jobs, there's now more competition for those jobs, for people internationally, so their wages-- that's one of the factors, I would-- the analyses have shown-- that's preventing more wage growth; and there are 5 to 10 percent of people, according to a lot of the analyses that I've shown, who are actually in absolute terms worse off because of globalization. Now, that doesn't necessarily mean that globalization for the whole world is negative. I think in general it's been, on balance, positive, but the story we've told about it has probably been too optimistic, in that we've only talked about the positives and how it's good as this global movement to bring people out of poverty and create more opportunities; and the reality I think has been that it's been net very positive, but if there are 5 or 10 percent of people in the world who are worse off-- there's 7 billion

people in the world, so that's many hundreds of millions of people, the majority of whom are likely in the most developed countries, in the U.S. and across Europe-- that's going to create a lot of political pressure on those in those countries. So in order to have a global system that works, it feels like-- you need it to work at the global level, but then you also need individuals in each of the member nations in that system to feel like it's working for them too, and that recurses all the way down, so even local cities and communities, people need to feel like it's working for them, both economically and socially. So I guess at this point the thing that I worry about-- and I've rotated a lot of Facebook's energy to try to focus on this-- is-- our mission used to be connecting the world. Now it's about helping people build communities and bringing people closer together, and a lot of that is because I actually think that the thing that we need to do to support more global connection at this point is making sure that things work for people locally. In a lot of ways we'd made it so the internet-- so that an emerging creator can--

Yuval Noah Harari: But then how do you balance working it locally for people in the American Midwest, and at the same time working it better for people in Mexico or South America or Africa? I mean, part of the imbalance is that when people in Middle America are angry, everybody pays attention, because they have their finger on the button. But if people in Mexico or people in Zambia feel angry, we care far less because they have far less power. I mean, the pain-- and I'm not saying the pain is not real. The pain is definitely real. But the pain of somebody in Indiana reverberates around the world far more than the pain of somebody in Honduras or in the Philippines, simply because of the imbalances of the power in the world. Earlier, what we said about fragmentation, I know that Facebook faces a lot of criticism about kind of encouraging people, some people, to move to these extremist groups, but-- that's a big problem, but I don't think it's the main problem. I think also it's something that you can solve-- if you put enough energy into that, that is something you can solve-- but this is the problem that gets most of the attention now. What I worry more-- and not just about Facebook, about the entire direction that the new internet economy and the new tech economy is going towards-- is increasing inequality between different parts of the world, which is not the result of extremist ideology, but the result of a certain economic and political model; and secondly, undermining human agency and undermining the basic philosophical ideas of democracy and the free market and individualism. These I would say are my two greatest concerns about the development of technology like AI and machine learning, and this will continue to be a major problem even if we find solutions to the issue of social extremism in particular groups.

Mark Zuckerberg: Yeah, I certainly agree that extremism isn't-- I would think about it more as a symptom and a big issue that needs to be worked on, but I think the bigger question is making sure that everyone has a sense of purpose, has a role that they feel matters and social connections, because at the end of the day, we're social animals and I think it's easy in our theoretical thinking to abstract that away, but that's such a fundamental part of who we are, so that's why I focus on that. I don't know, do you want to move over to some of the AI issues, because I think that that's a-- or do you want to stick on this topic for a second or--?

Yuval Noah Harari: No, I mean, this topic is closely connected to AI. And again, because I think that, you know, one of the disservices that science fiction, and I'm a huge fan of science fiction, but I think it has done some, also some pretty bad things, which is to focus attention on the wrong scenarios and the wrong dangers that people think, "Oh, AI is dangerous because the robots are coming to kill us." And this is extremely unlikely that we'll face a robot rebellion. I'm much more frightened about robots always obeying orders than about robots rebelling against the humans. I think the two main problems with AI, and we can explore this in greater depth, is what I just mentioned, first increasing inequality between different parts of the world because you'll have some countries which lead and dominate the new AI economy and this is such a huge advantage that it kind of trumps everything else. And we will see, I mean, if we had the Industrial Revolution creating this huge gap between a few industrial powers and everybody else and then it took 150 years to close the gap, and over the last few decades the gap has been closed or closing as more and more countries which were far behind are catching up. Now the gap may reopen and be much worse than ever before because of the rise of AI and because AI is likely to be dominated by just a small number of countries. So that's one issue, AI inequality. And the other issue is AI and human agency or even the meaning of human life, what happens when AI is mature enough and you have enough data to basically have human beings and you have an AI that knows me better than I know myself and can make decisions for me, predict my choices, manipulate my choices and authority increasingly shifts from humans to algorithms, so not only decisions about which movie to see but even decisions like which community to join, who to befriend, whom to marry will increasingly rely on the recommendations of the AI.

Mark Zuckerberg: Yeah.

Yuval Noah Harari: And what does it do to human life and human agency? So these I would say are the two most important issues of inequality and AI and human agency.

Mark Zuckerberg: Yeah. And I think both of them get down to a similar question around values, right, and who's building this and what are the values that are encoded and how does that end up playing out. I tend to think that in a lot of the conversations around AI we almost personify AI, right; your point around killer robots or something like that. But, but I actually think it's AI is very connected to the general tech sector, right. So almost every technology product and increasingly a lot of not what you call technology products have-- are made better in some way by AI. So it's not like AI is a monolithic thing that you build. It's it powers a lot of products, so it's a lot of economic progress and can get towards some of the distribution of opportunity questions that you're raising. But it also is fundamentally interconnected with these really socially important questions around data and privacy and how we want our data to be used and what are the policies around that and what are the global frameworks. And so one of the big questions that-- So, so I tend to agree with a lot of the questions that you're raising which is that a lot of the countries that have the ability to invest in future technology of which AI and data and future internet

technologies are certainly an important area are doing that because it will give, you know, their local companies an advantage in the future, right, and to be the ones that are exporting services around the world. And I tend to think that right now, you know, the United States has a major advantage that a lot of the global technology platforms are made here and, you know, certainly a lot of the values that are encoded in that are shaped largely by American values. They're not only. I mean, we, and I, speaking for Facebook, and we serve people around the world and we take that very seriously, but, you know, certainly ideas like giving everyone a voice, that's something that is probably very shaped by the American ideas around free speech and strong adherence to that. So I think culturally and economically, there's an advantage for countries to develop to kind of push forward the state of the field and have the companies that in the next generation are the strongest companies in that. So certainly you see different countries trying to do that, and this is very tied up in not just economic prosperity and inequality, but also--

Yuval Noah Harari: Do they have a real chance? I mean, does a country like Honduras, Ukraine, Yemen, has any real chance of joining the AI race? Or are they-- they are already out? I mean, they are, it's not going to happen in Yemen, it's not going to happen in Honduras? And then what happens to them in 20 years or 50 years?

Mark Zuckerberg: Well, I think that some of this gets down to the values around how it's developed, though. Right, is, you know, I think that there are certain advantages that countries with larger populations have because you can get to critical mass in terms of universities and industry and investment and things like that. But one of the values that we hear, right, both at Facebook and I think generally the academic system of trying to do research hold is that you do open research, right. So a lot of the work that's getting invested into these advances, in theory if this works well should be more open so then you can have an entrepreneur in one of these countries that you're talking about which, you know, maybe isn't a whole industry-wide thing and, you know, certainly, I think you'd bet against, you know, sitting here today that in the future all of the AI companies are going to be in a given small country. But I don't think it's far-fetched to believe that there will be an entrepreneur in some places who can use Amazon Web Services to spin up instances for Compute, who can hire people across the world in a globalized economy and can leverage research that has been done in the U.S. or across Europe or in different open academic institutions or companies that increasingly are publishing their work that are pushing the state of the art forward on that. So I think that there's this big question about what we want the future to look like. And part of the way that I think we want the future to look is we want it to be-- we want it to be open. We want the research to be open. I think we want the internet to be a platform. And this gets back to your unification point versus fragmentation. One of the big risks, I think, for the future is that the internet policy in each country ends up looking different and ends up being fragmented. And if that's the case, then I think the entrepreneur in the countries that you're talking about, in Honduras, probably doesn't have as big of a chance if they can't leverage the-- all the advances that are happening everywhere. But if the internet stays one thing and the research stays open, then I think that they have a much better shot. So when I look towards the future, one of the things that I just get very worried about is the values that I just

laid out are not values that all countries share. And when you get into some of the more authoritarian countries and their data policies, they're very different from the kind of regulatory frameworks that across Europe and across a lot of other people, people are talking about or put into place. And, you know, just to put a finer point on it, recently I've come out and I've been very vocal that I think that more countries should adopt a privacy framework like GDPR in Europe. And a lot of people I think have been confused about this. They're like, "Well, why are you arguing for more privacy regulation? You know, why now given that in the past you weren't as positive on it." And I think part of the reason why I am so focused on this now is I think at this point people around the world recognize that these questions around data and AI and technology are important so there's going to be a framework in every country. I mean, it's not like there's not going to be regulation or policy. So I actually think the bigger question is what is it going to be. And the most likely alternative to each country adopting something that encodes the freedoms and rights of something like GDPR, in my mind, the most likely alternative is the authoritarian model which is currently being spread, which says, you know, as every company needs to store everyone's data locally in data centers and you know, if I'm a government, I should be able to, you know, go send my military there and be able to get access to whatever data I want and be able to take that for surveillance or military or helping, you know, local military industrial companies. And I mean, I just think that that's a really bad future, right. And that's not-- that's not the direction that I, as, you know, someone who's building one of these internet services or just as a citizen of the world want to see the world going.

Yuval Noah Harari: To be the devil's advocate for a moment,--

Mark Zuckerberg: <laughs>

Yuval Noah Harari: I mean, if I look at it from the viewpoint, like, of India, so I listen to the American President saying, "America first and I'm a nationalist, I'm not a globalist. I care about the interests of America," and I wonder, is it safe to store the data about Indian citizens in the U.S. and not in India when they're openly saying they care only about themselves. So why should it be in America and not in India?

Mark Zuckerberg: Well, I think that there's, the motives matter and certainly, I don't think that either of us would consider India to be an authoritarian country that had-- So, so I would say that, well, it's-- <laughs>

Yuval Noah Harari: Well, it can still say--

Mark Zuckerberg: You know, it's-- <laughs>

Yuval Noah Harari: We want data and metadata on Indian users to be stored on Indian soil. We don't want it to be stored in-- on American soil or somewhere else.

Mark Zuckerberg: Yeah. And I can understand the arguments for that and I think that there's-- The intent matters, right. And I think countries can come at this with open values and still conclude that something like that could be helpful. But I think one of the things that you need to be very careful about is that if you set that precedent you're making it very easy for other countries that don't have open values and that are much more authoritarian and want the data not to-- not to protect their citizens but to be able to surveil them and find dissidents and lock them up. That-- So I think one of the-- one of the--

Yuval Noah Harari: No, I agree, I mean, but I think that it really boils down to the questions that do we trust America. And given the past two, three years, people in more and more places around the world-- I mean, previously, say if we were sitting here 10 years ago or 20 years ago or 40 years ago, then America declared itself to be the leader of the free world. We can argue a lot whether this was the case or not, but at least on the declaratory level, this was how America presented itself to the world. We are the leaders of the free world, so trust us. We care about freedom. But now we see a different America, America which doesn't want even to be-- And again, it's not a question of even what they do, but how America presents itself no longer as the leader of the free world but as a country which is interested above all in itself and in its own interests. And just this morning, for instance, I read that the U.S. is considering having a veto on the U.N. resolution against using sexual violence as a weapon of war. And the U.S. is the one that thinks of vetoing this. And as somebody who is not a citizen of the U.S., I ask myself, can I still trust America to be the leader of the free world if America itself says I don't want this role anymore.

Mark Zuckerberg: Well, I think that that's a somewhat separate question from the direction that the internet goes then, because I mean, GDPR, the framework that I'm advocating, that it would be better if more countries adopted something like this because I think that that's just significantly better than the alternatives, a lot of which are these more authoritarian models. I mean, GDPR originated in Europe, right.

Yuval Noah Harari: Yeah.

Mark Zuckerberg: And so that, because it's not an American invention. And I think in general, these values of openness in research, of cross-border flow of ideas and trade, that's not an American idea, right. I mean, that's a global philosophy for how the world should work and I think that the alternatives to that are at best fragmentation, right which breaks down the global model on this; at worst, a growth in authoritarianism for the models of how this gets adopted. And that's where I think that the precedents on some of this stuff get really tricky. I mean, you can-- You're, I think, doing a good job of playing devil's advocate in the conversation--

Yuval Noah Harari: <laughs>

Mark Zuckerberg: Because you're bringing all of the counterarguments that I think someone with good intent might bring to argue, "Hey, maybe a different set of data policies is something that we should consider." The thing that I just worry about is that what we've seen is that once a country puts that in place, that's a precedent that then a lot of other countries that might be more authoritarian use to basically be a precedent to argue that they should do the same things and, and then that spreads. And I think that that's bad, right. And that's one of the things that as the person running this company, I'm quite committed to making sure that we play our part in pushing back on that, and keeping the internet as one platform. So I mean, one of the most important decisions that I think I get to make as the person running this company is where are we going to build our data centers and store-- and store data. And we've made the decision that we're not going to put data centers in countries that we think have weak rule of law, that where people's data may be improperly accessed and that could put people in harm's way. And, you know, I mean, a lot has been-- There have been a lot of questions around the world around questions of censorship and I think that those are really serious and important. I mean, I, a lot of the reason why I build what we build is because I care about giving everyone a voice, giving people as much voice as possible, so I don't want people to be censored. At some level, these questions around data and how it's used and whether authoritarian governments get access to it I think are even more sensitive because if you can't say something that you want, that is highly problematic. That violates your human rights. I think in a lot of cases it stops progress. But if a government can get access to your data, then it can identify who you are and go lock you up and hurt you and hurt your family and cause real physical harm in ways that are just really deep. So I do think that people running these companies have an obligation to try to push back on that and fight establishing precedents which will be harmful. Even if a lot of the initial countries that are talking about some of this have good intent, I think that this can easily go off the rails. And when you talk about in the future AI and data, which are two concepts that are just really tied together, I just think the values that that comes from and whether it's part of a more global system, a more democratic process, a more open process, that's one of our best hopes for having this work out well. If it's, if it comes from repressive or authoritarian countries, then, then I just think that that's going to be highly problematic in a lot of ways.

Yuval Noah Harari: That raises the question of how do we-- how do we build AI in such a way that it's not inherently a tool of surveillance and manipulation and control? I mean, this goes back to the idea of creating something that knows you better than you know yourself, which is kind of the ultimate surveillance and control tool. And we are building it now. In different places around the world, it's been built. And what are your thoughts about how to build an AI which serves individual people and protects individual people and not an AI which can easily with a flip of a switch becomes kind of the ultimate surveillance tool?

Mark Zuckerberg: Well, I think that that is more about the values and the policy framework than the technological development. I mean, it's a lot of the research that's happening in AI are just very

fundamental mathematical methods where, you know, a researcher will create an advance and now all of the neural networks will be 3 percent more efficient. I'm just kind of throwing this out.

Yuval Noah Harari: Yeah.

Mark Zuckerberg: And that means that, all right, you know, newsfeed will be a little bit better for people. Our systems for detecting things like hate speech will be a little bit better. But it's, you know, our ability to find photos of you that you might want to review will be better. But all these systems get a little bit better. So now I think the bigger question is you have places in the world where governments are choosing to use that technology and those advances for things like widespread face recognition and surveillance. And those countries, I mean, China is doing this, they create a real feedback loop which advances the state of that technology where, you know, they say, "Okay, well, we want to do this," so now there's a set of companies that are sanctioned to go do that and they have-- are getting access to a lot of data to do it because it's allowed and encouraged. So, so that is advancing and getting better and better. It's not-- That's not a mathematical process. That's kind of a policy process that they want to go in that direction. So those are their-- the values. And it's an economic process of the feedback loop in development of those things. Compared to in countries that might say, "Hey, that kind of surveillance isn't what we want," those companies just don't exist as much, right, or don't get as much support and--

Yuval Noah Harari: I don't know. And my home country of Israel is, at least for Jews, it's a democracy.

Mark Zuckerberg: That's--

Yuval Noah Harari: And it's one of the leaders of the world in surveillance technology. And we basically have one of the biggest laboratories of surveillance technology in the world which is the occupied territories. And exactly these kinds of systems--

Mark Zuckerberg: Yeah.

Yuval Noah Harari: Are being developed there and exported all over the world. So given my personal experience back home, again, I don't necessarily trust that just because a society in its own inner workings is, say, democratic, that it will not develop and spread these kinds of technologies.

Mark Zuckerberg: Yeah, I agree. It's not clear that a democratic process alone solves it, but I do think that it is mostly a policy question, right. It's, you know, a government can quite easily make the decision that they don't want to support that kind of surveillance and then the companies that they would be working with to support that kind of surveillance would be out of business. And, and then, or at the very least, have much less economic incentive to continue that technological progress. So, so that dimension of the growth of the technology gets stunted compared to others. And that's-- and that's generally the process that I think you want to follow broadly, right. So technological advance isn't by itself good or bad. I

think it's the job of the people who are shepherding it, building it and making policies around it to have policies and make sure that their effort goes towards amplifying the good and mitigating the negative use cases. And, and that's how I think you end up bending these industries and these technologies to be things that are positive for humanity overall, and I think that that's a normal process that happens with most technologies that get built. But I think what we're seeing in some of these places is not the natural mitigation of negative uses. In some cases, the economic feedback loop is pushing those things forward, but I don't think it has to be that way. But I think that that's not as much a technological decision as it is a policy decision.

Yuval Noah Harari: I fully agree. But I mean, it's every technology can be used in different ways for good or for bad. You can use the radio to broadcast music to people and you can use the radio to broadcast Hitler giving a speech to millions of Germans. The radio doesn't care. The radio just carries whatever you put in it. So, yeah, it is a policy decision. But then it just raises the question, how do we make sure that the policies are the right policies in a world when it is becoming more and more easy to manipulate and control people on a massive scale like never before. I mean, the new technology, it's not just that we invent the technology and then we have good democratic countries and bad authoritarian countries and the question is what will they do with the technology. The technology itself could change the balance of power between democratic and totalitarian systems.

Mark Zuckerberg: Yeah.

Yuval Noah Harari: And I fear that the new technologies are inherent-- are giving an inherent advantage, not necessarily overwhelming, but they do tend to give an inherent advantage to totalitarian regimes. Because the biggest problem of totalitarian regimes in the 20th century, which eventually led to their downfall, is that they couldn't process the information efficiently enough. If you think about the Soviet Union, so you have this model, an information processing model which basically says, we take all the information from the entire country, move it to one place, to Moscow. There it gets processed. Decisions are made in one place and transmitted back as commands. This was the Soviet model of information processing. And versus the American version, which was, no, we don't have a single center. We have a lot of organizations and a lot of individuals and businesses and they can make their own decisions. In the Soviet Union, there is somebody in Moscow, if I live in some small farm or kulhose [ph?] in Ukraine, there is somebody in Moscow who tells me how many radishes to grow this year because they know. And in America, I decide for myself with, you know, I get signals from the market and I decide. And the Soviet model just didn't work well because of the difficulty of processing so much information quickly and with 1950s technology. And this is one of the main reasons why the Soviet Union lost the Cold War to the United States. But with the new technology, it's suddenly, it might become, and it's not certain, but one of my fears is that the new technology suddenly makes central information processing far more efficient than ever before and far more efficient than distributed data processing. Because the more data you have in one place, the better your algorithms and then so on and so forth. And this kind of tilts the balance

between totalitarianism and democracy in favor of totalitarianism. And I wonder what are your thoughts on this issue.

Mark Zuckerberg: Well, I'm more optimistic about--

Yuval Noah Harari: Yeah, I guess so.

<laughter>

Mark Zuckerberg: About democracy in this.

Yuval Noah Harari: Mm-hmm.

Mark Zuckerberg: I think the way that the democratic process needs to work is people start talking about these problems and then even if it seems like it starts slowly in terms of people caring about data issues and technology policy, because it's a lot harder to get everyone to care about it than it is just a small number of decision makers. So I think that the history of democracy versus more totalitarian systems is it always seems like the totalitarian systems are going to be more efficient and the democracies are just going to get left behind, but, you know, smart people, you know, people start discussing these issues and caring about them, and I do think we see that people do now care much more about their own privacy about data issues, about the technology industry. People are becoming more sophisticated about this. They realize that having a lot of your data stored can both be an asset because it can help provide a lot of benefits and services to you, but increasingly, maybe it's also a liability because there are hackers and nation states who might be able to break in and use that data against you or exploit it or reveal it. So maybe people don't want their data to be stored forever. Maybe they want it to be reduced in permanence. Maybe they want it all to be end-to-end encrypted as much as possible in their private communications. People really care about this stuff in a way that they didn't before. And that's certainly over the last several years, that's grown a lot. So I think that that conversation is the normal democratic process and I think what's going to end up happening is that by the time you get people broadly aware of the issues and on board, that is just a much more powerful approach where then you do have people in a decentralized system who are capable of making decisions who are smart, who I think will generally always do it better than too centralized of an approach. And here is again a place where I worry that personifying AI and saying, AI is a thing, right, that an institution will develop and it's almost like a sentient being, I think mischaracterizes what it actually is. Right. It's a set of methods that make everything better. Or, like, sorry. Then, sorry, let me retract that.

Yuval Noah Harari: <laughs>

Mark Zuckerberg: That's way too broad. It's a lot of technological processes more efficient. And, and I think that that's--

Yuval Noah Harari: But that's the worry.

Mark Zuckerberg: But that's--

Yuval Noah Harari: It makes also--

Mark Zuckerberg: But that's not just for-- that's not just for centralized folks, right, it's-- I mean, in our context, you know, so we build, our business is this ad platform and a lot of the way that that can be used now is we have 90 million small businesses that use our tools and now because of this access to technology, they have access to the same tools to do advertising and marketing and reach new customers and grow jobs that previously only the big companies would have had. And that's, that's a big advance and that's a massive decentralization. When people talk about our company and the internet platforms overall, they talk about how there's a small number of companies that are big. And that's true, but the flip side of it is that now there are billions of people around the world who have a voice that they can share information more broadly and that's actually a massive decentralization in power and kind of returning power to people. Similarly, people have access to more information, have access to more commerce. That's all positive. So I don't know. I'm an optimist on this. I think we have real work cut out for us and I think that the challenges that you raise are the right ones to be thinking about because if we get it wrong, that's the way in which I think it will go wrong. But I don't know. I think that the historical precedent would say that at all points, you know, where there was the competition with-- between the U.S. and Japan in the eighties and the seventies or the Cold War before that or different other times, people always thought that the democratic model, which is slow to mobilize but very strong once it does and once people get bought into a direction and understand the issue, I do think that that will continue to be the best way to spread prosperity around the world and make progress in a way that meets people's needs. And that's why, you know, when we're talking about internet policy, when you're talking about economic policy, I think spreading regulatory frameworks that encode those values I think is one of the most important things that we can do. But it starts with raising the issues that you are and having people be aware of the potential problems.

Yuval Noah Harari: Mm-hmm. Yeah, I agree and I think the last few decades it was the case that open democratic systems were better and more efficient. And this, I'm again, one of my fears is that it might have made us a bit complacent, because we assume that this is kind of a law of nature that distributed systems are always better and more efficient than centralized systems. And we lived-- we grew up in a world in which there was kind of this-- to do the good thing morally was also to do the efficient thing, economically and politically. And a lot of countries liberalized their economy, their society, their politics over the past 50 years, more because they were convinced of the efficiency argument than of the deep, moral argument. And what happens if efficiency and morality suddenly split, which has happened before in history? I mean, the last 50 years are not representative of the whole of history; we had many cases

before in human history in which repressive centralized systems were more efficient and, therefore, you got these repressive empires. And there is no law of nature, which says that “This cannot happen again.” And, again, my fear is that the new technology might tilt that balance; and, just by making central data processing far more efficient, it could give a boost to totalitarian regimes. Also, in the balance of power between, say, again, the center and the individual that for most of history the central authority could not really know you personally simply because of the inability to gather and process the information. So, there were some people who knew you very well, but usually their interests were aligned with yours. Like, my mother knows me very well, but most of the time I can trust my mother. But, now, we are reaching the point when some system far away can know me better than my mother and the interests are not necessarily aligned. Now, yes, we can use that also for good, but what I’m pointing out-- that this is a kind of power that never existed before and it could empower totalitarian and authoritarian regimes to do things that were simply, technically impossible.

Mark Zuckerberg: Mm-hm.

Yuval Noah Harari: Until today.

Mark Zuckerberg: Yeah.

Yuval Noah Harari: And, you know, if you live in an open democracy-- so, okay, you can rely on all kinds of mechanisms to protect yourself. But, thinking more globally about this issue, I think a key question is how do you protect human attention [ph?] from being hijacked by malevolent players who know you better than you know yourself, who know you better than your mother knows you? And this is a question that we never had to face before, because we never had-- usually the malevolent players just didn’t know me very well.

Mark Zuckerberg: Yeah. Okay, so, there’s a lot in what you were just talking about.

Yuval Noah Harari: Yeah.

Mark Zuckerberg: I mean, I think in general one of the things that-- do you think that there is a scale effect where one of the best things that we could do to-- if we care about these open values and having a globally connected world, I think making sure that the critical mass of the investment in new technologies encodes those values is really important. So, that’s one of the reasons why I care a lot about not supporting the spread of authoritarian policies to more countries, either inadvertently doing that or setting precedents that enable that to happen. Because the more development that happens in the way that is more open, where the research is more open, where people have the-- where the policymaking around it is more democratic, I think that that’s going to be positive. So, I think kind of maintaining that balance ends up being really important. One of the reasons why I think democratic countries over time tend to do

better on serving what people want is because there's no metric to optimize the society, right? When you talk about efficiency, a lot what people are talking about is economic efficiency, right?

Yuval Noah Harari: Yeah.

Mark Zuckerberg: Are we increasing GDP? Are we increasing jobs? Are we decreasing poverty? Those things are all good, but I think part of what the democratic process does is people get to decide on their own which of the dimensions in society matter the most to them in their lives.

Yuval Noah Harari: But if you can hijack people's attention and manipulate--

Mark Zuckerberg: See--

Yuval Noah Harari: --them, then people deciding on their own just doesn't help, because I don't realize it that somebody manipulated me to think that this is what I want. If-- and we are reaching the point when for the first time in history you can do that on a massive scale. So, again, I speak a lot about the issue of free will in this regard--

Mark Zuckerberg: Yeah.

Yuval Noah Harari: --and the people that are easiest to manipulate are the people who believe in free will and who simply identify with whatever thought or desire pops up in their mind, because they cannot even imagine--

Mark Zuckerberg: Mm-hm.

Yuval Noah Harari: --that this desire is not the result of my free will. This desire is the result of some external manipulation. Now it may sound paranoid-- and for most of history it was probably paranoid, because nobody had this kind of ability to do it on a massive scale--

Mark Zuckerberg: Yeah.

Yuval Noah Harari: --but, here, like in Silicon Valley, the tools to do that on a massive scale have been developed over the last few decades. And they may have been developed with the best intentions; some of them may have been developed with the intention of just selling stuff to people and selling products to people. But now the same tools that can be used to sell me something I don't really need can now be used to sell me a politician I really don't need or an ideology that I really don't need. It's the same tool. It's the same hacking the human animal and manipulating what's happening inside.

Mark Zuckerberg: Yeah, okay. So, there's a lot going on here. I think that there's-- when designing these systems I think that there's the intrinsic design, which you want to make sure that you get right and then there's preventing abuse--

Yuval Noah Harari: Yeah.

Mark Zuckerberg: --which I think is-- so, I think that there's two types of questions that people raise. I mean, one is we saw what the Russian government tried to do in the 2016 election. That's clear abuse. We need to build up really advanced systems for detecting that kind of interference in the democratic process and more broadly being able to identify that, identify when people are standing up networks of fake accounts that are not behaving in a way that normal people would, to be able to weed those out and work with law enforcement and election commissions and folks all around the world and the intelligence community to be able to coordinate and be able to deal with that effectively. So, stopping abuse is certainly important, but I would argue that, even more, the deeper question: Is that the intrinsic design of the systems, right?

Yuval Noah Harari: Yeah, exactly.

Mark Zuckerberg: So, not just fighting the abuse. And, there, I think that the incentives are more aligned towards a good outcome than a lot of critics might say. And here's why: I think that there's a difference between what people want first order and what they want second order over time. So, right now, you might just consume a video, because you think it's silly or fun. And, you know, you wake up-- or you kind of look up an hour later and you've watched a bunch of videos and you're like, "Well, what happened to my time?" And, okay, so, maybe in the narrow short-term period you consume some more content and maybe you saw some more ads. So, it seems like it's good for the business, but it actually really isn't over time, because people make decisions based on what they find valuable. And what we find, at least in our work, is that what people really want to do is connect with other people. Right? It's not just passively consumed content. It's-- so, we've had to find and constantly adjust our systems over time to make sure that we're rebalancing it; so, that way you're interacting with people; so, that way we make sure that we don't just measure signals in the system, like, what are you clicking on, because that can get you into a bad local optimum.

Yuval Noah Harari: Yeah.

Mark Zuckerberg: But, instead, we bring in real people to tell us what their real experience is in words, right? Not just kind of filling out scores, but also telling us what were the most meaningful experiences you had today, what content was the most important, what interaction did you have with a friend that mattered to you the most and was that connected to something that we did? And, if not, then we go and try to do

the work to try to figure out how we can facilitate that. And what we find is that, yeah, in the near-term, maybe showing some people some more viral videos might increase time, right? But, over the long term, it doesn't. It's not actually aligned with our business interests or the long-term social interest. So, kind of in strategy terms, that would be a stupid thing to do. And I think a lot of people think that businesses are just very short-term oriented and that we only care about-- people think that businesses only care about the next quarter profit, but I think that most businesses that get run well that's just not the case. And, you know, I think last year on one our earnings calls, you know, I told investors that we'd actually reduced the amount of video watching that quarter by 50 million hours a day, because we wanted to take down the amount of viral videos that people were seeing, because we thought that that was displacing more meaningful interactions that people were having with other people, which, in the near-term, might have a short-term impact on the business for that quarter, but, over the long term, would be more positive both for how people feel about the product and for the business. And, you know, one of the patterns that I think has actually been quite inspiring or a cause of optimism in running a business is that oftentimes you make decisions that you think are going to pay off long down the road, right? So, you think, "Okay, I'm doing the right thing long term, but it's going to hurt for a while." And I almost always find that the long term comes sooner than you think and that when you make these decisions that there may be taking some pain in the near term in order to get to what will be a better case down the line, that better case-- maybe you think it'll take five years, but, actually, it ends up coming in a year. Right? And I think people at some deep level know when something is good. And, like, I guess this gets back to the democratic values, because, at some level, I trust that people have a sense of what they actually care about. And it may be that, you know, if we were showing more viral videos, maybe that would be better than the alternatives that they have to do right now, right? I mean, maybe that's better than what's on TV, because at least they're personalized videos. You know, maybe it's better than YouTube, if we have better content or whatever the reason is. But I think you can still make the service better over time for actually matching what people want; and if you do that, that is going to be better for everyone. So, I do think the intrinsic design of these systems is quite aligned with serving people in a way that is pro-social and that's certainly what I care about in running this company is to get there.

Yuval Noah Harari: Yeah, and I think this is like the rock bottom, that this is the most important issue that, ultimately, what I'm hearing from you and from many other people when I have these discussions, is ultimately the customer is always right, the voter knows best, people know deep down, people know what is good for them. People make a choice: If they choose to do it, then it's good. And that has been the bedrock of, at least, Western democracies for centuries, for generations. And this is now where the big question mark is: Is it still true in a world where we have the technology to hack human beings and manipulate them like never before that the customer is always right, that the voter knows best? Or have we gone past this point? And we can know-- and the simple, ultimate answer that "Well, this is what people want," and "they know what's good for them," maybe it's no longer the case.

Mark Zuckerberg: Well, yeah, I think that the-- it's not clear to me that that has changed, but I think that that's a very deep question about democracy.

Yuval Noah Harari: Yeah, I was going to say, this is the deepest--

Mark Zuckerberg: I don't think that that's a new question. I mean, I think that people have always wondered--

Yuval Noah Harari: No, the question isn't this. The technology is new. I mean, if you lived in 19th century America and you didn't have these extremely powerful tools to decipher and influence people, then it was a different--

Mark Zuckerberg: Well, let me actually frame this a different way--

Yuval Noah Harari: Okay.

Mark Zuckerberg: --which is I actually think, you know, for all the talk around "Is democracy being hurt by the current set of tools and the media," and all this, I actually think that there's an argument the world is significantly more democratic now than it was in the past. I mean, the country was set up as-- the U.S. was set up as a republic, right? So, a lot of the foundational rules limited the power of a lot of individuals being able to vote and have a voice and checked the popular will at a lot of different stages, everything from the way that laws get written by Congress, right, and not by people, you know, so, everything-- to the Electoral College, which a lot of people think today is undemocratic, but, I mean, it was put in place because of a set of values that a democratic republic would be better. I actually think what has happened today is that increasingly more people are enfranchised and more people have a voice, more people are getting the vote, but, increasingly, people have a voice, more people have access to information and I think a lot of what people are asking is "Is that good?" It's not necessarily the question of "Okay, the democratic process has been the same, but now the technology is different." I think the technology has made it so individuals are more empowered and part of the question is "Is that the world that we want?" And, again, this is an area where it's not-- I mean, all these things are with challenges, right? And often progress causes a lot of issues and it's a really hard thing to reason through, "Wow, we're trying to make progress and help all these people join the global economy," or help people join the communities and have the social lives that they would want and be accepted in different ways, but it comes with this dislocation in the near term and that's a massive dislocation. So, that seems really painful. But I actually think that you can make a case that we're at-- and continue to be at the most democratic time and I think that overall in the history of our country at least, when we've gotten more people to have the vote and we've gotten more representation and we've made it so that people have access to more information and more people can share their experiences, I do think that that's made the country stronger and has helped

progress. And it's not that this stuff is without issues. It has massive issues. But that's, at least, the pattern that I see and why I'm optimistic about a lot of the work.

Yuval Noah Harari: I agree that more people have more voice than ever before, both in the U.S. and globally. That's-- I think you're absolutely right. My concern is to what extent we can trust the voice of people-- to what extent I can trust my voice, like I'm-- we have this picture of the world, that I have this voice inside me, which tells me what is right and what is wrong, and the more I'm able to express this voice in the outside world and influence what's happening and the more people can express their voices, it's better, it's more democratic. But what happens if, at the same time that more people can express their voices, it's also easier to manipulate your inner voice? To what extent you can really trust that the thought that just popped up in your mind is the result of some free will and not the result of an extremely powerful algorithm that understands what's happening inside you and knows how to push the buttons and press the levers and is serving some external entity and it has planted this thought or this desire that we now express? So, it's two different issues of giving people voice and trusting-- and, again, I'm not saying I know everything, but all these people that now join the conversation, we cannot trust their voices. I'm asking this about myself, to what extent I can trust my own inner voice. And, you know, I spend two hours meditating every day and I go on these long meditation retreats and my main takeaway from that is it's craziness inside there and it's so complicated. And the simple, naïve belief that the thought that pops up in my mind "This is my free will," this was never the case. But if, say, a thousand years ago the battles inside were mostly between, you know, neurons and biochemicals and childhood memories and all that; increasingly, you have external actors going under your skin and into your brain and into your mind. And how do I trust that my amygdala is not a Russian agent now? How do I know-- the more we understand about the extremely complex world inside us, the less easy it is to simply trust what this inner voice is telling, is saying.

Mark Zuckerberg: Yeah, I understand the point that you're making. As one of the people who's running a company that develops ranking systems to try to help show people content that's going to be interesting to them there's a dissonance between the way that you're explaining what you think is possible and what I see as a practitioner building this. I think you can build systems that can get good at a very specific thing, right? At helping to understand which of your friends you care the most about so you can rank their content higher in newsfeed. But the idea that there's some kind of generalized AI that's a monolithic thing that understands all dimensions of who you are in a way that's deeper than you do, I think doesn't exist and is probably quite far off from existing. So, there's certainly abuse of the systems that I think needs to be-- that I think is more of a policy and values question, which is-- you know, on Facebook, you know, you're supposed to be your real identity. So, if you have, to use your example, Russian agents or folks

from the government, the IRA, who are posing as someone else and saying something and you see that content, but you think it's coming from someone else, then that's not an algorithm issue. I mean, that's someone abusing the system and taking advantage of the fact that you trust that on this platform someone is generally going to be who they are, so you can trust that the information is coming from some place and kind of slipping in the backdoor that way and that's the thing that we certainly need to go fight. But, I don't know, as broad matter, I do think that there's this question of, you know, to what degree are the systems-- this kind of brings it full circle to where we started on "Is it fragmentation or is it personalization?" You know, is the content that you see-- if it resonates, is that because it actually just more matches your interests or is it because you're being incepted and convinced of something that you don't actually believe and doesn't-- and is dissonant with your interests and your beliefs. And, certainly, all the psychological research that I've seen and the experience that we've had, is that when people see things that don't match what they believe, they just ignore it.

Yuval Noah Harari: Mm-hm.

Mark Zuckerberg: Right? So, certainly, there is a-- there can be an evolution that happens where a system shows information that you're going to be interested in; and if that's not managed well, that has the risk of pushing you down a path towards adopting a more extreme position or evolving the way you think about it over time. But I think most of the content, it resonates with people because it resonates with their lived experience. And, to the extent that people are abusing that and either trying to represent that they're someone who they're not or are trying to take advantage of a bug in human psychology where we might be more prone to an extremist idea, that's our job in either policing the platform, working with governments and different agencies, and making sure that we design our systems and our recommendation systems to not be promoting things that people might engage with in the near term, but over the long term will regret and resent us for having done that. And I think it's in our interests to get that right. And, for a while, I think we didn't understand the depth of some of the problems and challenges that we faced there and there's certainly still a lot more to do. And when you're up against nation-states, I mean, they're very sophisticated, so they're going to keep on evolving their tactics. But the thing that I would-- that I think is really important is that the fundamental design of the systems I do think-- and our incentives are aligned with helping people connect with the people they want, have meaningful interactions, not just getting people to watch a bunch of content that they're going to resent later that they did that and certainly not making people have more extreme or negative viewpoints than what they actually believe. So.

Yuval Noah Harari: Mm-hm. Maybe I can try and summarize my view in that we have two distinct dangers coming out of the same technological tools. We have the easier danger to grasp, which is of extreme totalitarian regimes of the kind we haven't seen before, and this could happen in different-- maybe not in the U.S., but in other countries, that these tools, you say that-- I mean, that these are abuses. But in some countries, this could become the norm. That you're living from the moment you are

born in this system that constantly monitors and surveils you and constantly kind of manipulates you from a very early age to adopt particular ideas, views, habits, so forth, in a way which was never possible before.

Mark Zuckerberg: Mm-hm.

Yuval Noah Harari: And this is like the full-fledged totalitarian dystopia, which could be so effective that people would not even resent it, because they will be completely aligned with the values or the ideals of the sys-- it's not "1984" where you need to torture people all the time. No! If you have agents inside their brain, you don't need the external secret police. So, that's one danger. It's like the full-fledged totalitarianism. Then, in places like the U.S., the more immediate danger or problem to think about is what, increasingly, people refer to as surveillance capitalism; that you have these systems that constantly interact with you and come to know you and it's all supposedly in your best interests to give you better recommendations and better advice. So, it starts with recommendation for which movie to watch and where to go on vacation. But, as the system becomes better, it gives you recommendation on what to study at college and where to work, ultimately, whom to marry who to vote for, which religion to join-- like, join a community. Like, "You have all these religious communities. This is the best religion for you for your type of personality, Judaism, nah, it won't work for you. Go with Zen Buddhism. It's a much better fit for your personality. You would thank us. In five years, you would look back and say, 'This was an amazing recommendation. Thank you. I so much enjoy Zen Buddhism.'" And, again, people will-- it will feel that this is aligned with their own best interests and the system improves over time. Yeah, there will be glitches. Not everybody will be happy all the time. But what does it mean that all the most important decisions in my life are being taken by an external algorithm? What does it mean in terms of human agency, in terms of the meaning of life?

Mark Zuckerberg: Mm-hm.

Yuval Noah Harari: You know, for thousands of years, humans tended to view life as a drama of decision-making. Like, life is-- it's a journey, you reach an intersection after intersection and you need to choose. Some decisions are small, like what to eat for breakfast, and some decisions are really big like whom to marry. And almost all of art and all of religion is about that. Like, almost every-- whether it's a Shakespeare tragedy or a Hollywood comedy, it's about the hero or heroine needing to make a big decision, "To be or not to be," to marry X or to marry Y. And what does it mean to live in a world in which, increasingly, we rely on the recommendations of algorithms to make these decisions until we reach a point when we simply follow them all the time or most of the time. And they make good recommendations. I'm not saying that this is some abuse, something sinister-- no! They are good recommendations, but I'm just-- we don't have a model for understanding what is the meaning of human life in such a situation?

Mark Zuckerberg: Well, I think the biggest objection that I'd have to what-- to both of the ideas that you just raised is that we have access to a lot of different sources of information, a lot of people to talk to about different things. And it's not just like there's one set of recommendations or a single recommendation that gets to dominate what we do and that that gets to be overwhelming either in the totalitarian or the capitalist model of what you were saying. To the contrary, I think people really don't like and are very distrustful when they feel like they're being told what to do or just have a single option. One of the big questions that we've studied is how to address when there's a hoax or clear misinformation. And the most obvious thing that it would seem like you'd do intuitively is tell people, "Hey, this seems like it's wrong. Here is the other point of view that is right," or, at least, if it's a polarized thing, even if it's not clear what's wrong and what's right, "here's the other point of view," on any given issue. And that really doesn't work, right? So, what ends up happening is if you tell people that something is false, but they believe it, then they just end up not trusting you.

Yuval Noah Harari: Yeah.

Mark Zuckerberg: Right? So, that ends up not working. And if you frame two things as opposites-- right? So, if you say, "Okay, well, you're a person who doesn't believe in-- you're seeing content about not believing in climate change, I'm going to show you the other perspective, right? Here's someone that argues that climate change is a thing," that actually just entrenches you further, because it's, "Okay, someone's trying to kind of control--"

Yuval Noah Harari: Yeah, it's a-- mm-hm.

Mark Zuckerberg: Okay, so what ends up working, right-- sociologically and psychologically, the thing that ends up actually being effective is giving people a range of choices. So, if you show not "Here's the other opinion," and with a judgement on the piece of content that a person engaged with, but instead you show a series of related articles or content, then people can kind of work out for themselves, "Hey, here's the range of different opinions," or things that exist on this topic. And maybe I lean in one direction or the other, but I'm kind of going to work out for myself where I want to be. Most people don't choose the most extreme thing and people end up feeling like they're informed and can make a good decision. So, at the end of the day, I think that that's the architecture and the responsibility that we have is to make sure that the work that we're doing gives people more choices, that it's not a given-- a single opinion that can kind of dominate anyone's thinking but where you can, you know, connect to hundreds of different friends. And even if most of your friends share your religion or your political ideology, you're probably going to have five or 10 percent of friends who come from a different background, who have different ideas and, at least that's getting in as well. So, you're getting a broader range of views. So, I think that these are really important questions and it's not like there's an answer that is going to fully solve it one way or another.

Yuval Noah Harari: That's-- definitely not. [ph?]

Mark Zuckerberg: But I feel these are the right things to talk through. You know, we've been going for 90 minutes. So, we probably should wrap up. But I think we have a lot of material to cover in the next one of these--

Yuval Noah Harari: Yeah.

Mark Zuckerberg: --that, hopefully, we'll get to do at some point in the future. And thank you so much for coming and joining and doing this. This has been a really interesting series of important topics to discuss.

Yuval Noah Harari: Yeah, so, thank you for hosting me and for being open about these very difficult questions, which I know that you, being the head of a global corpora-- I can just sit here and speak whatever I want--

<laughter>

Yuval Noah Harari: --but you have many more responsibilities on your head. So, I appreciate that kind of you putting yourself on the firing line and dealing with these questions.

Mark Zuckerberg: Thanks. All right.

Yuval Noah Harari: Thank you.

Mark Zuckerberg: Yeah.

End of MZ_Yuval_low_res.mp3