

TRANSCRIPT - MARK ZUCKERBERG AT HARVARD / FIRST PERSONAL CHALLENGE 2019

Jonathan Zittrain: Very good. So, thank you, Mark, for coming to talk to me and to our students from the Techtopia program and from my “Internet and Society” course at Harvard Law School. We’re really pleased to have a chance to talk about any number of issues and we should just dive right in. So, privacy, autonomy, and information fiduciaries.

Mark Zuckerberg: All right!

Jonathan Zittrain: Love to talk about that.

Mark Zuckerberg: Yeah! I read your piece in The New York Times.

Jonathan Zittrain: The one with the headline that said, “Mark Zuckerberg can fix this mess”? **Mark Zuckerberg:** Yeah.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: Although that was last year.

<laughter>

Jonathan Zittrain: That’s true! Are you suggesting it’s all fixed?

<laughter>

Mark Zuckerberg: No. No.

<laughter>

Jonathan Zittrain: Okay, good. So--

Mark Zuckerberg: I’m suggesting that I’m curious whether you still think that we can fix this mess?

Jonathan Zittrain: Ah! <laughter>

Jonathan Zittrain: I hope-- <laughter>

Jonathan Zittrain: “Hope springs eternal”--

Mark Zuckerberg: Yeah, there you go.

Jonathan Zittrain: --is my motto. So, all right, let me give a quick characterization of this idea that the coinage and the scaffolding for it is from my colleague, Jack Balkin, at Yale. And the two of us have been developing it out further. There are a standard number of privacy questions with which you might have some familiarity, having to do with people conveying information that they know they’re conveying or they’re not so sure they are, but “mouse droppings” as we used to call them when they run in the rafters of the Internet and leave traces. And then the standard way of talking about that is you want to make sure that that stuff doesn’t go where you don’t want it to go. And we call that “informational privacy”. We don’t want people to know stuff that we want maybe our friends only to know. And on a place like Facebook, you’re supposed to be able to tweak your settings and say, “Give them to this and not to that.” But there’s also ways in which stuff that we share with consent could still sort of be used against us and it feels like, “Well, you consented,” may not end the discussion. And the analogy that my colleague Jack brought to bear was one of a doctor and a patient or a lawyer and a client or-- sometimes in America, but not always-- a financial advisor and a client that says that those professionals have certain expertise, they get trusted with all sorts of sensitive information from their clients and patients and, so, they have an extra duty to act in the interests of those clients even if their own interests conflict. And, so, maybe just one quick hypo to get us started. I wrote a piece in 2014, that maybe you read, that was a hypothetical about elections in which it said, “Just hypothetically, imagine that Facebook had a view about which candidate should win and they reminded people likely to vote for the favored candidate that it was Election Day,” and to others they simply sent a cat photo. Would that be wrong? And I find-- I have no idea if it’s illegal; it does seem wrong to me and it might be that the fiduciary approach captures what makes it wrong.

Mark Zuckerberg: All right. So, I think we could probably spend the whole next hour just talking about that! <laughter>

Mark Zuckerberg: So, I read your op-ed and I also read Balkin's blogpost on information fiduciaries. And I've had a conversation with him, too.

Jonathan Zittrain: Great.

Mark Zuckerberg: And the-- at first blush, kind of reading through this, my reaction is there's a lot here that makes sense. Right? The idea of us having a fiduciary relationship with the people who use our services is kind of intuitively-- it's how we think about how we're building what we're building. So, reading through this, it's like, all right, you know, a lot of people seem to have this mistaken notion that when we're putting together news feed and doing ranking that we have a team of people who are focused on maximizing the time that people spend, but that's not the goal that we give them. We tell people on the team, "Produce the service--" that we think is going to be the highest quality that-- we try to ground it in kind of getting people to come in and tell us, right, of the content that we could potentially show what is going to be-- they tell us what they want to see, then we build models that kind of-- that can predict that, and build that service.

Jonathan Zittrain: And, by the way, was that always the case or--

Mark Zuckerberg: No.

Jonathan Zittrain: --was that a place you got to through some course adjustments?

Mark Zuckerberg: Through course adjustments. I mean, you start off using simpler signals like what people are clicking on in feed, but then you pretty quickly learn, "Hey, that gets you to local optimum," right? Where if you're focusing on what people click on and predicting what people click on, then you select for click bait. Right? So, pretty quickly you realize from real feedback, from real people, that's not actually what people want. You're not going to build the best service by doing that. So, you bring in people and actually have these panels of-- we call it "getting to ground truth"-- of you show people all the candidates for what can be shown to them and you have people say, "What's the most meaningful thing that I wish that this system were showing us? So, all this is kind of a way of saying that our own self image of ourselves and what we're doing is that we're acting as fiduciaries and trying to build the best services for people. Where I think that this ends up getting interesting is then the question of who gets to decide in the legal sense or the policy sense of what's in people's best interest? Right? So, we come in every day and think, "Hey, we're building a service where we're ranking newsfeed trying to show people the most relevant content with an assumption that's backed by data; that, in general, people want us to show them the most relevant content. But, at some level, you could ask the question which is "Who gets to decide that ranking newsfeed or showing relevant ads?" or any of the other things that we choose to work on are actually in people's interest. And we're doing the best that we can to try to build the services [ph?] that we think are the best. At the end of the day, a lot of this is grounded in "People choose to use it." Right? Because, clearly, they're getting some value from it. But then there are all these questions like you say about, you have-- about where people can effectively give consent and not.

Jonathan Zittrain: Yes.

Mark Zuckerberg: So, I think that there's a lot of interesting questions in this to unpack about how you'd implement a model like that. But, at a high level I think, you know, one of the things that I think about in terms of we're running this big company; it's important in society that people trust the institutions of

society. Clearly, I think we're in a position now where people rightly have a lot of questions about big internet companies, Facebook in particular, and I do think getting to a point where there's the right regulation and rules in place just provides a kind of societal guardrail framework where people can have confidence that, okay, these companies are operating within a framework that we've all agreed. That's better than them just doing whatever they want. And I think that that would give people confidence. So, figuring out what that framework is, I think, is a really important thing. And I'm sure we'll talk about that as it relates--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --to a lot of the content areas today. But getting to that question of how do you-- "Who determines what's in people's best interest, if not people themselves?"--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --is a really interesting question.

Jonathan Zittrain: Yes, so, we should surely talk about that. So, on our agenda is the "Who decides?" question.

Mark Zuckerberg: All right.

Jonathan Zittrain: Other agenda items include-- just as you say, the fiduciary framework sounds nice to you-- doctors, patients, Facebook users. And I hear you saying that's pretty much where you're wanting to end up anyway. There are some interesting questions about what people want, versus what they want to want.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: People will say "On January 1st, what I want--" New Year's resolution-- "is a gym membership." And then on January 2nd, they don't want to go to the gym. They want to want to go to the gym, but they never quite make it. And then, of course, a business model of pay for the whole year ahead of time and they know you'll never turn up develops around that. And I guess a specific area to delve into for a moment on that might be on the advertising side of things, maybe the dichotomy between personalization and does it ever going into exploitation? Now, there might be stuff-- I know Facebook, for example, bans payday loans as best it can.

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: That's just a substantive area that it's like, "All right, we don't want to do that." **Mark**

Zuckerberg: Mm-hm.

Jonathan Zittrain: But when we think about good personalization so that Facebook knows I have a dog and not a cat, and a targeter can then offer me dog food and not cat food. How about, if not now, a future day in which an advertising platform can offer to an ad targeter some sense of "I just lost my pet, I'm really upset, I'm ready to make some snap decisions that I might regret later, but when I make them--"

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: "--I'm going to make them." So, this is the perfect time to tee up-- **Mark Zuckerberg:** Yeah.

Jonathan Zittrain: --a Cubic Zirconia or whatever the thing is that-- .

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: That seems to me a fiduciary approach would say, ideally-- how we get there I don't know, but ideally we wouldn't permit that kind of approach to somebody using the information we've gleaned from them to know they're in a tough spot--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: --and then to exploit them. But I don't know. I don't know how you would think about something like that. Could you write an algorithm to detect something like that?

Mark Zuckerberg: Well, I think one of the key principles is that we're trying to run this company for the long term. And I think that people think that a lot of things that-- if you were just trying to optimize the profits for next quarter or something like that, you might want to do things that people might like in the near term, but over the long term will come to resent. But if you actually care about building a community and achieving this mission and building the company for the long term, I think you're just much more aligned than people often think companies are. And it gets back to the idea before, where I think our self image is largely acting as-- in this kind of fiduciary relationship as you're saying-- and across-- we could probably go through a lot of different examples. I mean, we don't want to show people content that they're going to click on and engage with, but then feel like they wasted their time afterwards. Where we don't want to show them things that they're going to make a decision based off of that and then regret later. I mean, there's a hard balance here which is-- I mean if you're talking about what people want to want versus what they want-- you know, often people's revealed preferences of what they actually do shows a deeper sense of what they want than what they think they want to want. So, I think there's a question between when something is exploitative versus when something is real, but isn't what you would say that you want.

Jonathan Zittrain: Yes.

Mark Zuckerberg: And that's a really hard thing to get at. **Jonathan Zittrain:** Yes.

Mark Zuckerberg: But on a lot of these cases my experience of running the company is that you start off building a system, you have relatively unsophisticated signals to start, and you build up increasingly complex models over time that try to take into account more of what people care about. And there are all these examples that we can go through. I think probably newsfeed and ads are probably the two most complex ranking examples--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --that we have. But it's-- like we were talking about a second ago, when we started off with the systems, I mean, just start with newsfeeds-- but you could do this on ads, too-- you know, the most naïve signals, right, are what people click on or what people "Like". But then you just very quickly realize that that doesn't-- it approximates something, but it's a very crude approximation of the ground truth of what people actually care about. So, what you really want to get to is as much as possible getting real people to look at the real candidates for content and tell you in a multi-dimensional way what matters to them and try to build systems that model that. And then you want to be kind of conservative on preventing downside. So, your example of the payday loans-- and when we've talked about this in the past, your-- you've put the question to me of "How do you know when a payday loan is going to be exploitative?" right? "If you're targeting someone who is in a bad situation?" And our answer is, "Well, we

don't really know when it's going to be exploitative, but we think that the whole category potentially has a massive risk of that, so we just ban it--

Jonathan Zittrain: Right. Which makes it an easy case.

Mark Zuckerberg: Yes. And I think that the harder cases are when there's significant upside and significant downside and you want to weigh both of them. So, I mean, for example, once we started putting together a really big effort on preventing election interference, one of the initial ideas that came up was "Why don't we just ban all ads that relate to anything that is political?" And they you pretty quickly get into, all right, well, what's a political ad? The classic legal definition is things that are around elections and candidates, but that's not actually what Russia and other folks were primarily doing. Right? It's-- you know, a lot of the issues that we've seen are around issue ads, right, and basically sewing division on what are social issues. So, all right, I don't think you're going to get in the way of people's speech and ability to promote and do advocacy on issues that they care about. So, then the question is "All right, well, so, then what's the right balance?" of how do you make sure that you're providing the right level of controls, that people who aren't supposed to be participating in these debates aren't or that at least you're providing the right transparency. But I think we've veered a little bit from the original question--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --but the-- but, yeah. So, let's get back to where you were--

Jonathan Zittrain: Well, here's-- and this is a way of maybe moving it forward, which is: A platform as complete as Facebook is these days offers lots of opportunities to shape what people see and possibly to help them with those nudges, that it's time to go to the gym or to avoid them from falling into the depredations of the payday loan. And it is a question of so long as the platform to do it, does it now have an ethical obligation to do it, to help people achieve the good life?

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: And I worry that it is too great a burden for any company to bear to have to figure out, say, if not the perfect, the most reasonable newsfeed for every one of the-- how many? Two and a half billion active users? Something like that.

Mark Zuckerberg: Yeah. On that order.

Jonathan Zittrain: All the time and there might be some ways that start a little bit to get into the engineering of the thing that would say, "Okay, with all hindsight, are there ways to architect this so that the stakes aren't as high, aren't as focused on just, "Gosh, is Facebook doing this right?" It's as if there was only one newspaper in the whole world or one or two, and it's like, "Well, then what The New York Times chooses to put on it's home page, if it were the only newspaper, would have outsize importance."

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: So, just as a technical matter, a number of the students in this room had a chance to hear from Tim Berners-Lee, inventor of the World Wide Web, and he has a new idea for something called "Solid". I don't know if you've heard of Solid. It's a protocol more than it is a product. So, there's no car to move off the lot today. But its idea is allowing people to have the data that they generate as they motor around the web end up in their own kind of data locker. Now, for somebody like Tim, it might mean literally in a locker under his desk and he could wake up in the middle of the night and see where his data

is. For others, it might mean Iraq somewhere, guarded perhaps by a fiduciary who's looking out for them, the way that we put money in a bank and then we can sleep at night knowing the bankers are-- this is maybe not the best analogy in 2019, but watching.

<laughter>

Mark Zuckerberg: We'll get there.

Jonathan Zittrain: We'll get there. But Solid says if you did that, people would then-- or their helpful proxies-- be able to say, "All right, Facebook is coming along. It wants the following data from me and including that data that it has generated about me as I use it, but stored back in my locker and it kind of has to come back to my well to draw water each time. And that way if I want to switch to Schmacebook or something, it's still in my well and I can just immediately grant permission to Schmacebook to see it and I don't have to do a kind of data slurp and then re-upload it. It's a fully distributed way of thinking about data. And I'm curious from an engineering perspective does this seem doable with something of the size and the number of spinning wheels that Facebook has and does it seem like a--

Mark Zuckerberg: Yeah--

Jonathan Zittrain: --and I'm curious your reaction to an idea like that.

Mark Zuckerberg: So, I think it's quite interesting. Certainly, the level of computation that Facebook is doing and all the services that we're building is really intense to do in a distributed way. I mean, I think as a basic model I think we're building out the data center capacity over the next five years and our plan for what we think we need to do that we think is on the order of all of what AWS and Google Cloud are doing for supporting all of their customers. So, okay, so, this is like a relatively computationally intense thing. Over time you assume you'll get more compute. So, decentralized things which are less efficient computationally will be harder-- sorry, they're harder to do computation on, but eventually maybe you have the compute resources to do that. I think the more interesting questions there are not feasibility in the near term, but are the philosophical questions of the goodness of a system like that.

So, one question if you want to-- so, we can get into decentralization, one of the things that I've been thinking about a lot is a use of block chain that I am potentially interesting in-- although I haven't figured out a way to make this work out, is around authentication and bringing-- and basically granting access to your information and to different services. So, basically, replacing the notion of what we have with Facebook Connect with something that's fully distributed.

Jonathan Zittrain: "Do you want to login with your Facebook account?" is the status quo--

Mark Zuckerberg: Basically, you take your information, you store it on some decentralized system and you have the choice of whether to login to different places and you're not going through an intermediary, which is kind of like what you're suggesting here--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --in a sense. Okay, now, there's a lot of things that I think would be quite attractive about that. You know, for developers one of the things that is really troubling about working with our system, or Google's system for that matter, or having your services through Apple's app store, is that you don't want to have an intermediary between serving your-- the people who are using your service and you, right, where someone can just say, "Hey, we as a developer have to follow your policy and if we

don't, then you can cut off access to the people we're serving." That's kind of a difficult and troubling position to be in. I think developers--

<overlapping conversation>

Jonathan Zittrain: --you're referring to a recent incident. **Mark Zuckerberg:** No, well, I was-- well, sure-- <laughter>

Mark Zuckerberg: But I think it underscores the-- I think every developer probably feels this: People are using any app store but also login with Facebook, with Google; any of these services, you want a direct relationship with the people you serve.

Jonathan Zittrain: Yes.

Mark Zuckerberg: Now, okay, but let's look at the flip side. So, what we saw in the last couple of years with Cambridge Analytica, was basically an example where people chose to take data that they-- some of it was their data, some of it was data that they had seen from their friends, right? Because if you want to do things like making it so alternative services can build a competing newsfeed, then you need to be able to make it so that people can bring the data that they see you [ph?] within the system. Okay, they-- basically, people chose to give their data to a developer who's affiliated with Cambridge University, which is a really respected institution, and then that developer turned around and sold the data to the firm Cambridge Analytica, which is in violation of our policies. So, we cut off the developers' access. And, of course, in a fully distributed system there would be no one who could cut off the developers' access. So, the question is if you have a fully distributed system, it dramatically empowers individuals on the one hand, but it really raises the stakes and it gets to your questions around, well, what are the boundaries on consent and how people can really actually effectively know that they're giving consent to an institution? In some ways it's a lot easier to regulate and hold accountable large companies like Facebook or Google, because they're more visible, they're more transparent than the long tail of services that people would chose to then go interact with directly. So, I think that this is a really interesting social question. To some degree I think this idea of going in the direction of block chain authentication is less gated on the technology and capacity to do that. I think if you were doing fully decentralized Facebook, that would take massive computation, but I'm sure we could do fully decentralized authentication if we wanted to. I think the real question is do you really want that?

Jonathan Zittrain: Yes.

Mark Zuckerberg: Right? And I think you'd have more cases where, yes, people would be able to not have an intermediary, but you'd also have more cases of abuse and the recourse would be much harder.

Jonathan Zittrain: Yes. What I hear you saying is people as they go about their business online are generating data about themselves that's quite valuable, if not to themselves, to others who might interact with them. And the more they are empowered, possibly through a distributed system, to decide where that data goes, with whom they want to share it, the more they could be exposed to exploitation. this is a genuine dilemma--

Mark Zuckerberg: Yeah, yeah.

Jonathan Zittrain: --because I'm a huge fan of decentralization. **Mark Zuckerberg:** Yeah, yeah.

Jonathan Zittrain: But I also see the problem. And maybe one answer is there's some data that's just so toxic there's no vessel we should put it in; it might eat a whole through it or something, metaphorically

speaking. But, then again, innocuous data can so quickly be assembled into something scary. So, I don't know if the next election--

<overlapping conversation>

Mark Zuckerberg: Yeah. [ph?] I mean, I think in general we're talking about the large-scale of data being assembled into meaning something different from what the individual data points mean.

Jonathan Zittrain: Yes.

Mark Zuckerberg: And I think that's the whole challenge here. But I philosophically agree with you that-- I mean, I want to think about the-- like, I do think about the work that we're doing as a decentralizing force in the world, right? A lot of the reason why I think people of my generation got into technology is because we believe that technology gives individuals power and isn't massively centralizing. Now you've built a bunch of big companies in the process, but I think what has largely happened is that individuals today have more voice, more ability to affiliate with who they want, and stay connected with people, ability to form communities in ways that they couldn't before, and I think that's massively empowering to individuals and that's philosophically kind of the side that I tend to be on. So, that's why I'm thinking about going back to decentralized or block chain authentication. That's why I'm kind of bouncing around how could you potentially make this work, because from my orientation is to try to go in that direction.

Jonathan Zittrain: Yes.

Mark Zuckerberg: An example where I think we're generally a lot closer to going in that direction is encryption. I mean, this is, like, one of the really big debates today is basically what are the boundaries on where you would want a messaging service to be encrypted. And there are all these benefits from a privacy and security perspective, but, on the other hand, if what we're trying to do-- one of the big issues that we're grappling with content governance and where is the line between free expression and, I suppose, privacy on one side, but safety on the other as people do really bad things, right, some of the time. And I think people rightfully have an expectation of us that we're going to do everything we can to stop terrorists from recruiting people or people from exploiting children or doing different things. And moving in the direction of making these systems more encrypted certainly reduces some of the signals that we would have access to be able to do some of that really important work.

But here we are, right, we're sitting in this position where we're running WhatsApp, which is the largest end-to-end encrypting service in the world; we're running messenger, which is another one of the largest messaging systems in the world where encryption is an option, but it isn't the default. I don't think long term it really makes sense to be running different systems with very different policies on this. I think this is sort of a philosophical question where you want to figure out where you want to be on it. And, so, my question for you-- now,

I'll talk about how I'm thinking about this-- is all right, if you were in my position and you got to flip a switch is probably too glib, because there's a lot of work that goes into this, and go in one direction for both of those services, who would you think about that?

Jonathan Zittrain: Well, the question you're putting on the table, which is a hard one is "Is it okay," and let's just take the simple case, "for two people to communicate with each other in a way that makes it difficult for any third party to casually listen in?" Is that okay? And I think that the way we normally answer

that question is kind of a form of what you might call status quo-ism, which is not satisfying. It's whatever has been the case is—

Mark Zuckerberg: Yeah, yeah.

Jonathan Zittrain: --whatever has been the case is what should stay the case.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: And, so, for WhatsApp, it's like right now WhatsApp, as I understand it, you could correct me if I'm wrong, is pretty hard to get into if--

Mark Zuckerberg: It's fully end-to-end encrypted.

Jonathan Zittrain: Right. So, if Facebook gets handed a subpoena or a warrant or something from name-your-favorite-country--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: --and you're just like, "Thank you for playing. We have nothing to--" <overlapping conversation>

Mark Zuckerberg: Oh, yeah, we've had employees thrown in jail because we have gotten court orders that we have to turnover data that we wouldn't probably anyway, but we can't because it's encrypted.

Jonathan Zittrain: Yes. And then, on the other hand, and this is not as clean as it could be in theory, but Messenger is sometimes encrypted, sometimes not. If it doesn't happen to have been encrypted by the users, then that subpoena could work and, *more than that*, there could start to be some automated systems either on Facebook's own initiative or under pressure from governments in the general case, not a specific warrant, to say, "Hey, if the following phrases appear, if there's some telltale that says, "This is somebody going after a kid for exploitation," it should be forwarded up. If that's already happening and we can produce x-number of people who have been identified and a number of crimes averted that way, who wants to be the person to be like, "Lock it down!" Like, "We don't want any more of that!" But I guess, to put myself now to your question, when I look out over years rather than just weeks or months, the ability to casually peek at any conversation going on between two people or among a small group of people or even to have a machine do it for you, so, you can just set your alert list, you know, crudely speaking, and get stuff back, that-- it's always trite to call something Orwellian, but it makes Orwell look like a piker. I mean, it seems like a classic case where you-- the next sentence would be "What could possible go wrong?"

Jonathan Zittrain: And we can fill that in! And it does mean, though, I think that we have to confront the fact that if we choose to allow that kind of communication, then there's going to be crimes unsolved that could've been solved. There's going to be crimes not prevented that could have been prevented. And the only thing that kind of blunts it a little is it is not really all or nothing. The modern surveillance states of note in the world, have a lot of arrows in their quivers. And just being able to darken you door and demand surveillance of a certain kind, that might be a first thing they would go to, but they've got a Plan B, and Plan C, and a Plan D. And I guess it really gets to what's your threat model? If you think everybody is kind of a threat, think about the battles of copyright 15 years ago. Everybody is a potential infringer. All they have to do is fire up Napster, then you're wanting some massive technical infrastructure to prevent the bad thing. If what you're thinking is instead, they are a few really bad apples and they tend to-- when they congregate online or otherwise with one another-- tend to identify themselves and then we might have to send somebody near their house to listen with a cup at the window, metaphorically speaking. That's a different threat model and [sic] might not need it.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: Is that getting to an answer to your question?

Mark Zuckerberg: Yeah, and I think I generally agree. I mean, I've already said publically that my inclination is to move these services in the direction of being all encrypted, at least the private communication version. I basically think if you want to kind of talk in metaphors, messaging is like people's living room, right? And I think we-- you know, we definitely don't want a society where there's a camera in everyone's living room watching the content of those conversations.

Jonathan Zittrain: Even as we're now-- I mean, it is 2019, people are happily are putting cameras in their living rooms.

Mark Zuckerberg: That's their choice, but I guess they're putting cameras in their living rooms, well, for a number of reasons, but--

Jonathan Zittrain: And Facebook has a camera that you can go into your living room- <laughter>

Mark Zuckerberg: That is, I guess--

Jonathan Zittrain: I just want to be clear.

<laughter>

Mark Zuckerberg: Yeah, although that would be encrypted in this world. **Jonathan Zittrain:** Encrypted between you and Facebook!

Mark Zuckerberg: No, no, no. I think-- but it also--

<overlapping conversation>

Jonathan Zittrain: Doesn't it have like a little Alexa functionality, too?

Mark Zuckerberg: Well, Portal works over Messenger. So, if we go towards encryption on Messenger, then that'll be fully encrypted, which I think, frankly, is probably what people want.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: The other model, beside the living room is the town square and that, I think, just has different social norms and different policies and norms that should be at play around that. But I do think that these things are very different. Right? You're not going to-- you may end up in a world where the town square is a fully decentralized or fully encrypted thing, but it's not clear what value there is in encrypting something that's public content anyway, or very broad.

Jonathan Zittrain: But, now, you were put to it pretty hard in that as I understand it there's now a change to how WhatsApp works, that there's only five forwards permitted.

Mark Zuckerberg: Yeah, so, this is a really interesting point, right? So, when people talk about how encryption will darken some of the signals that we'll be able to use, you know, both for potentially providing better services and for preventing harm. One of the-- I guess, somewhat surprising to me, findings of the last couple of years of working on content governance and enforcement is that it often is much more effective to identify fake accounts and bad actors upstream of them doing something bad by patterns of activity rather than looking at the content.

Jonathan Zittrain: So-called meta data.

Mark Zuckerberg: Sure.

Jonathan Zittrain: "I don't know what they're saying, but here's who's they're calling" kind of thing.

Mark Zuckerberg: Yeah, or just like they-- this account doesn't seem to really act like a person, right? And I guess as AI gets more advanced and you build these adversarial networks or generalized adversarial networks, you'll get to a place where you have AI that can probably more effectively--

Jonathan Zittrain: Go under mimic [ph?] cover. Mimic act like another person-- <overlapping conversation>

Mark Zuckerberg: --for a while.

Mark Zuckerberg: Yeah. But, at the same time, you'll be building up contrary AI on the other side, but is better at identifying AIs that are doing that. But this has certainly been the most effective tactic across a lot of the areas where we've needed to focus to preventing harm. You know, the ability to identify fake accounts, which, like, a huge amount of the-- under any category of issue that you're talking about, a lot of the issues downstream come from fake accounts or people who are clearly acting in some malicious or not normal way. You can identify a lot of that without necessarily even looking at the content itself. And if you have to look at a piece of content, then in some cases, you're already late, because the content exists and the activity has already happened. So, that's one of the things that makes me feel like encryption for these messaging services is really the right direction to go, because you're-- it's a very pro-privacy and per security move to give people that control and assurance and I'm relatively confident that even though you are losing some tools to-- on the finding harmful content side of the ledger, I don't think at the end of the day that those are going to end up being the most important tools--

Jonathan Zittrain: Yes.

Mark Zuckerberg: --for finding the most of the--

<overlapping conversation>

Jonathan Zittrain: But now connect it up quickly to the five forwards thing.

Mark Zuckerberg: Oh, yeah, sure. So, that gets down to if you're not operating on a piece of content directly, you need to operate on patterns of behavior in the network. And what we, basically found was there weren't that many good uses for people forwarding things more than five times except to basically spam or blast stuff off. It was being disproportionately abused. So, you end up thinking about different tactics when you're not operating on content specifically; you end up thinking about patterns of usage more.

Jonathan Zittrain: Well, spam, I get and that-- I'm always in favor of things that reduce spam. However, you could also say the second category was just to spread content. You could have the classic, I don't know, like Les Mis, or Paul Revere's ride, or Arab Spring-esque in the romanticized vision of it: "Gosh, this is a way for people to do a tree," and pass along a message that "you can't stop the signal," to use a Joss Whedon reference. You really want to get the word out. This would obviously stop that, too.

Mark Zuckerberg: Yeah, and then I think the question is you're just weighing whether you want this private communication tool where the vast majority of the use and the reason why it was designed was the vast majority of just one-on-one; there's a large amount of groups that people communicate into, but it's a pretty small edge case of people operating this with, like-- you have a lot of different groups and

you're trying to organize something and almost hack public content-type or public sharing- type utility into an encrypted space and, again, there I think you start getting into "Is this the living room or is this the town square?" And when people start trying to use tools that are designed for one thing to get around what I think the social norms are for the town square, that's when I think you probably start to have some issues. This is not-- we're not done addressing these issues. There's a lot more to think through on this--

Jonathan Zittrain: Yeah.

Mark Zuckerberg: --but that's the general shape of the problem that at least I perceive from the work that we're doing.

Jonathan Zittrain: Well, without any particular segue, let's talk about fake news.

<laughter>

Jonathan Zittrain: So, insert your favorite segue here. There's some choice or at least some decision that gets made to figure out what's going to be next in my newsfeed when I scroll up a little more.

Mark Zuckerberg: Mm-hm.

Jonathan Zittrain: And in the last conversation bit, we were talking about how much we're looking at content versus telltales and metadata, things that surround the content.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: For knowing about what that next thing in the newsfeed should be, is it a valid desirable material consideration, do you think, for a platform like Facebook to say is the thing we are about to present true, whatever true means?

Mark Zuckerberg: Well, yes, because, again, getting at trying to serve people, people tell us that they don't want fake content. Right. I mean, I don't know anyone who wants fake content. I think the whole issue is, again, who gets to decide. Right. So broadly speaking, I don't know any individual who would sit there and say, "Yes, please show me things that you know are false and that are fake." People want good quality content and information. That said, I don't really think that people want us to be deciding what is true for them and people disagree on what is true. And, like, truth is, I mean, there are different levels of when someone is telling a story, maybe the meta arc is talking about something that is true but the facts that were used in it are wrong in some nuanced way but, like, it speaks to some deeper experience. Well, was that true or not? And do people want that disqualified from to them? I think different people are going to come to different places on this.

Now, so I've been very sensitive, which, on, like, we really want to make sure that we're showing people high quality content and information. We know that people don't want false information. So we're building quite advanced systems to be able to-- to make sure that we're emphasizing and showing stuff that is going to be high quality. But the big question is where do you get the signal on what the quality is? So the kind of initial v.1 of this was working with third party fact checkers.

Right, I believe very strongly that people do not want Facebook and that we should not be the arbiters of truth in deciding what is correct for everyone in the society. I think people already generally think that we

have too much power in deciding what content is good. I tend to also be concerned about that and we should talk about some of the governance stuff that we're working on separately to try to make it so that we can bring more independent oversight into that.

Jonathan Zittrain: Yes.

Mark Zuckerberg: But let's put that in a box for now and just say that with those concerns in mind, I'm definitely not looking to try to take on a lot more in terms of also deciding in addition to enforcing all the content policies, also deciding what is true for everyone in the world. Okay, so v.1 of that is we're going to work with--

Jonathan Zittrain: Truth experts.

Mark Zuckerberg: We're working with fact checkers. **Jonathan Zittrain:** Yeah.

Mark Zuckerberg: And, and they're experts and basically, there's like a whole field of how you go and assess certain content. They're accredited. People can disagree with the leaning of some of these organizations.

Jonathan Zittrain:

Mark Zuckerberg: Jonathan Zittrain: Mark Zuckerberg: Jonathan Zittrain: Mark Zuckerberg: <laughter> Who does accredited fact checkers? <laughs> The Poynter Institute for Journalism.

Okay, good.

And then-- And--

I should apply for my certification. You may.

You'd probably get it, but you have to-- You'd have to go through the process. The issue there is there aren't enough of them, right. So there's a large content. There's obviously a lot of information is shared every day and there just aren't a lot of fact checkers. So then the question is okay, that is probably--

Mark Zuckerberg:

Jonathan Zittrain: But the portion-- You're saying the food is good, it's just the portions are small. But the food is good.

Mark Zuckerberg: I think in general, but so you build system, which is what we've done especially leading up to elections where I think are some of the most fraught times around this where people really are aggressively trying to spread misinformation.

Jonathan Zittrain: Yes.

Mark Zuckerberg: You build systems that prioritize content that seems like it's going viral because you want to reduce the prevalence of how widespread the stuff gets, so that way the fact checkers have tools to be able to, like, prioritize what they need to go-- what they need to go look at. But it's still getting to a relatively small percent of the content. So I think the real thing that we want to try to get to over time is more of a crowd sourced model where people, it's not that people are trusting some sort, some basic set of experts who are accredited but are in some kid of lofty institution somewhere else. It's like do you trust, yeah, like, if you get enough data points from within the community of people reasonably looking at

something and assessing it over time, then the question is can you compound that together into something that is a strong enough signal that we can then use that?

Jonathan Zittrain: Kind of in the old school like a slash-dot moderating system--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: With only the worry that if the stakes get high enough, somebody wants to Astro turf that.

Mark Zuckerberg: Yes.

Jonathan Zittrain: I'd be--

Mark Zuckerberg: There are a lot of questions here, which is why I'm not sitting here and announcing a new program.

<laughter>

Mark Zuckerberg: But what I'm saying is this is, like,-- **Jonathan Zittrain:** Yeah,--

Mark Zuckerberg: This is the general direction that I think we should be thinking about when we have-- and I think that there's a lot of questions and--

Jonathan Zittrain: Yes.

Mark Zuckerberg: And we'd like to run some tests in this area to see whether this can help out. Which

would be upholding the principles which are that we want to stop-- **Jonathan Zittrain:** Yes.

Mark Zuckerberg: The spread of misinformation.

Jonathan Zittrain: Yes.

Mark Zuckerberg: Knowing that no one wants misinformation. And the other principle, which is that we do not want to be arbiters of truth.

Jonathan Zittrain: Want to be the decider, yes.

Mark Zuckerberg: And I think that that's the basic-- those are the basic contours I think of that, of that problem.

Jonathan Zittrain: So let me run an idea by you that you can process in real time and tell me the eight reasons I have not thought of why this is a terrible idea. And that would be people see something in their Facebook feed. They're about to share it out because it's got a kind of outrage factor to it. I think of the classic story from two years ago in *The Denver Guardian* about "FBI agent suspected in Hilary Clinton email leak implicated in murder-suicide." I have just uttered fake news.

None of that was true if you clicked through *The Denver Guardian*. There was just that article. There is *Denver Guardian*. If you live in Denver, you cannot subscribe. Like, it is unambiguously fake. And it was shared more times than the most shared story during the election season of *The Boston Globe*. And so--

Mark Zuckerberg: So, and this is actually an example, by the way, of where trying to figure out fake accounts is a much simpler solution.

Jonathan Zittrain: Yes.

Mark Zuckerberg: Than trying to down--

Jonathan Zittrain: So if newspaper has one article--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: Wait for ten more before you decide they're a newspaper.

Mark Zuckerberg: Yeah. Or, you know, I mean, it's there are any number of systems that you could build to basically detect, "Hey, this is--"

Jonathan Zittrain: A Potemkin.

Mark Zuckerberg: This is a fraudulent thing.

Jonathan Zittrain: Yes.

Mark Zuckerberg: And then you can take that down. And again, that ends up being a much less controversial decision because you're doing it upstream based on the basis of inauthenticity.

Jonathan Zittrain: Yes.

Mark Zuckerberg: In a system where people are supposed to be their real and represent that they're their real selves than downstream, trying to say, "Hey, is this true or false?"

Jonathan Zittrain: I made a mistake in giving you the easy case. **Mark Zuckerberg:** Okay.

<laughter>

Jonathan Zittrain: So I should have not used that example. **Mark Zuckerberg:** Too simple.

Jonathan Zittrain: You're right and you knocked that one out of the park and, like, *Denver Guardian*, come up with more articles and be real and then come back and talk to us.

<laughter>

Jonathan Zittrain: So, here's the harder case which is something that might be in an outlet that is, you know, viewed as legitimate, has a number of users, et cetera. So you can't use the metadata as easily. Imagine if somebody as they shared it out could say, "By the way, I want to follow this. I want to learn a little bit more about this." They click a button that says that. And I also realized when I talked earlier to somebody at Facebook on this that adding a new button to the homepage is, like, everybody's first idea

Mark Zuckerberg: Oh, yeah.

Jonathan Zittrain: And it's--

Mark Zuckerberg: But it's a reasonable thought experiment, even though it would lead to a very bad UI.

Jonathan Zittrain: Fair enough. I understand this is already--

Mark Zuckerberg: Yeah. <laughs>

Jonathan Zittrain: In the land of fantasy. So they add the button. They say, "I want to follow up on this." If enough people are clicking comparatively on the same thing to say, "I want to learn more about this. If anything else develops, let me know, Facebook," that, then, if I have my pneumatic tube, it then goes to a convened virtually panel of three librarians. We go to the librarians of the nation and the world at public and private libraries across the land who agree to participate in this program. Maybe we set up a little foundation for it that's endowed permanently and no long connected to whoever endowed it. And those librarians together discuss the piece and they come back with what they would tell a patron if somebody came up to them and said, "I'm about to cite this in my social studies paper. What do you think?" And librarians, like, live for questions like that.

Mark Zuckerberg: Mm-hmm, yeah.

Jonathan Zittrain: They're like, "Wow. Let us tell you." And they have a huge fiduciary notion of patron duty that says, "I may disapprove of you even studying this, whatever, but I'm here to serve you, the user."

Mark Zuckerberg: Yeah.

Jonathan Zittrain: "And I just think you should know, this is why maybe it's not such a good source." And when they come up with that they can send it back and it gets pushed out to everybody who asks for follow-up--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: And they can do with it as they will. And last piece of the puzzle, we have high school students who apprentice as librarian number three for credit.

<laughter>

Jonathan Zittrain: And then they can get graded on how well they participated in this exercise which helps generate a new generation of librarian-themed people who are better off at reading things, so.

Mark Zuckerberg: All right, well, I think you have a side goal here which I haven't been thinking about on the librarian thing.

<laughter>

Mark Zuckerberg: Which is the evil goal of promoting libraries. **Jonathan Zittrain:** Well, it's--

Mark Zuckerberg: No, but I mean, look, I think solving-- preventing misinformation or spreading misinformation is hard enough without also trying to develop high school students in a direction.

Jonathan Zittrain: Ah. My colleague Charlies Foote--

Mark Zuckerberg: So, that's solving a problem with a problem. **Jonathan Zittrain:** All right. Well, anyway, yes.

Mark Zuckerberg: So I actually think I agree with most of what you have in there. It doesn't need to be a button on the home page, it can be-- I mean, it turns out that there's so many people using these services that even if you get-- even if you put something that looks like it's not super prominent, like, behind the three dots on a given newsfeed story, you have the options, yeah, you're not-- not everyone is going to-- is going to like something.

Jonathan Zittrain: If 1 out of 1000 do it, you still get 10,000 or 100,000 people, yeah.

Mark Zuckerberg: You get pretty good signal. But I actually think you could do even better, which is, it's not even clear that you need that signal. I think that that's super helpful. I think really what matters is looking at stuff that's getting a lot of distribution. So, you know, I think that there's kind of this notion, and I'm going back to the encryption conversation, which is all right, if I say something that's wrong to you in a one-on-one conversation, I mean, does that need to be fact checked? I mean, it's, yeah, it would be good if you got the most accurate information.

Jonathan Zittrain: I do have a personal librarian to accompany me for most conversations, yes. There you go.

<laughter>

Mark Zuckerberg: Well, you are--

Jonathan Zittrain: Unusual.

Mark Zuckerberg: Yeah, yeah. Yes.

<laughter>

Mark Zuckerberg: That's the word I was looking for. <laughter>

Jonathan Zittrain: I'm not sure I believe you, but yes.

Mark Zuckerberg: It's-- <laughs> But I think that there's limited-- I don't think anyone would say that every message that goes back and forth in especially an encrypted messaging service should be--

Jonathan Zittrain: Fact checked.

Mark Zuckerberg: Should be fact checked. **Jonathan Zittrain:** Correct.

Mark Zuckerberg: So I think the real question is all right, when something starts going viral or getting a lot of distribution, that's when it becomes most socially important for it to be-- have some level of validation or at least that we know where that the community in general thinks that this is a reasonable thing. So it's actually, while it's helpful to have the signal of whether people are flagging this as something that we should look at, I actually think increasingly you want to be designing systems that just prevent like alarming or sensational content from going viral in the first place. And making sure that that, that the stuff that is getting wide distribution is doing so because it's high quality on whatever front you care about. So then, okay--

Jonathan Zittrain: And that quality is still generally from Poynter or some external party that-- **Mark**

Zuckerberg: Well, well quality has many dimensions.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: But certainly accuracy is one dimension of it. You also, I mean, you pointed out I think in one of your questions, is this piece of content prone to incite outrage. If you don't mind, I'll get to your panel of three things in a second, but as a slight detour on this.

Jonathan Zittrain: Yes.

Mark Zuckerberg: One of the findings that has been quite interesting is, you know, there's this question about whether social media in general increases, basically makes it so that sensationalist content gets the most distribution. And what we've found is that, all right, so we're going to have rules, right, about what content is allowed. And what we found is that generally within whatever rules you set up, as content approaches the line of what is allowed, it often gets more distribution. So if you'll have some rule on, you know, what-- And take a completely different example and our nudity policies. Right. It's like, okay, you have to define what is unacceptable nudity in some way. As you get as close to that as possible it's like, all right. Like, this is maybe a photo of someone--

Jonathan Zittrain: The skin to share ratio goes up until it gets banned at which point it goes to zero.

Mark Zuckerberg: Yes. Okay. So that is a bad property of a system, right, that I think you want to generally address. Or you don't want to design a community where or systems for helping to build a community where things that get as close to the line as what is bad get the most distribution.

Jonathan Zittrain: So long as we have the premise, which in many cases is true, but I could probably try to think of some where it wouldn't be true, that as you near the line, you are getting worse.

Mark Zuckerberg: That's a good point. That's a good point. There's--

Jonathan Zittrain: You know, there might be humor that's really edgy.

Mark Zuckerberg: That's true.

Jonathan Zittrain: And that conveys a message that would be impossible to convey without the edginess, while not still--

Mark Zuckerberg: That is-- **Jonathan Zittrain:** But, I-- **Mark Zuckerberg:** That's true. **Jonathan Zittrain:** Yeah.

Mark Zuckerberg: So but then you get the question of what's the cost benefit of allowing that. And obviously, where you can accurately separate what's good and bad which you, like in the case of misinformation I'm not sure you could do it fully accurately, but you can try to build systems that approximate that, there's certainly the issue, which is that, I mean, there is misinformation which leads to massive public harm, right. So if it's misinformation that is also spreading hate and leading to genocide or public attacks or, it's like, okay, we're not going to allow that. Right. That's coming down. But then generally if you say something that's wrong, we're not going to try to block that.

Jonathan Zittrain: Yes.

Mark Zuckerberg: We're just going to try to not show it to people widely because people don't want content that is wrong. So then the question is as something is approaching the line, how do you assess that? This is a general theme in a lot of the content governance and enforcement work that we're doing, which is there's one piece of this which is just making sure that we can as effectively as possible enforce the policies that exist. Then there's a whole other stream of work, which I called borderline content, which is basically this issue of as content

approaches the line of being against the policies, how do you make sure that that isn't the content that is somehow getting the most distribution? And a lot of the things that we've done in the last year were focused on that problem and it really improves the quality of the service and people appreciate that.

Jonathan Zittrain: So this idea would be stuff that you're kind of letting down easy without banning and letting down easy as it's going to somehow have a coefficient of friction for sharing that goes up. It's going to be harder--

Mark Zuckerberg: Yeah. **Jonathan Zittrain:** For it to go viral. **Mark Zuckerberg:** Yeah.

Jonathan Zittrain: And--

Mark Zuckerberg: So it's fascinating because it's just against-- Like, you can take almost any category of policy that we have, so I used nudity a second ago. You know, gore and violent imagery.

Jonathan Zittrain: Yes.

Mark Zuckerberg: Hate speech. **Jonathan Zittrain:** Yeah.

Mark Zuckerberg: Any of these things. I mean, there's, like, hate speech, there's content that you would just say is mean or toxic, but that did not violate-- But that you would not want to have a society that banned being able to say that thing. But it's, but you don't necessarily want that to be the content that is getting the most distribution.

Jonathan Zittrain: So here's a classic transparency question around exactly that system you described. And when you described this, I think you did a post around this a few months ago. This was fascinating. You had graphs in the post depicting this, which was great. How would you feel about sharing back to the person who posted or possibly to everybody who encounters it its coefficient of friction? Would that freak people out? Would it be, like, all right, I-- And in fact, they would then probably start conforming their posts, for better or worse,--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: To try to maximize the sharability. But that rating is already somewhere in there by design. Would it be okay to surface it?

Mark Zuckerberg: So, as a principle, I think that that would be good, but I don't-- The way that the systems are designed isn't that you get a score of how inflammatory or sensationalist a piece of content is. The way that it basically works is you can build classifiers that identify specific types of things. Right. So we're going down the list of, like, all right, there's 20 categories of harmful content that you're trying to identify.

You know, everything from terrorist propaganda on the one hand to self-harm issues to hate speech and election interference. And basically, each of these things while it uses a lot of the same underlying machine learning infrastructure, you're doing specific work for each of them. So if you go back to the example on Nudity for a second, you know, what you-- you're not necessarily scoring everything on a scale of not at all nude to nude. You're basically enforcing specific policies. So, you know, you're saying, "Okay, if--"

Jonathan Zittrain: So by machine learning it would just be give me an estimate of the odds by which if a human looked at it who was employed to enforce policy--

Mark Zuckerberg: Well, basically--

Jonathan Zittrain: Whether it violates the policy.

Mark Zuckerberg: And you have a sense of, okay, this is-- So what are the things that are adjacent to the policy, right? So you might say, okay, well, if the person is completely naked, that is something that you can definitely build a classifier to be able to identify with relatively high accuracy. But even if they're not, you know, then the question is you kind of need to be able to qualitatively describe what are the things that are adjacent to that. So maybe the person is wearing a bathing suit and is in a sexually suggestive position. Right. It's not like any piece of content you're going to score from not at all nude to nude. But you kind of have the cases for what you think are adjacent to the issues and, and again, you ground this and qualitatively, people, like, people might click on it, they might engage with it, but at the end, they don't necessarily feel good about it. And you want to get at when you're designing these systems not just what people do, but also you want to make sure we factor in, too, like is this the content that people say that they really want to be seeing? Do they--?

Jonathan Zittrain: The constitutional law, there's a formal kind of definition that's emerged for the word "prurient." If something appeals to the prurient interest--

Mark Zuckerberg: Okay.

Jonathan Zittrain: As part of a definition of obscenity, the famous Miller test, which was not a beer-oriented test. And part of a prurient interest is basically it excites me and yet it completely disgusts me. And it sounds like you're actually converging to the Supreme Court's vision of prurience with this.

Mark Zuckerberg: Maybe.

<laughter>

Jonathan Zittrain: And it might be-- Don't worry, I'm not trying to nail you down on that. But it's very interesting that machine learning, which you invoked, is both really good, I gather, at something like this. It's the kind of thing that's like just have some people tell me with their expertise, does this come near to violating the policy or not and I'll just through a Spidey sense start to tell you whether it would.

Mark Zuckerberg: Mm-hmm.

Jonathan Zittrain: Rather than being able to throw out exactly what the factors are. I know the person's fully clothed, but it still is going to--

Mark Zuckerberg:

Jonathan Zittrain:

drawbacks where--

Yeah.

Invoke that quality. So all of the benefits of machine learning and all of, of course, the

Yeah.

Mark Zuckerberg: Jonathan Zittrain:

It classifies something and somebody's like, "Wait a minute. That was me doing a parody of blah, blah, blah." That all comes to the fore.

Mark Zuckerberg: Yeah and I mean, when you ask people what they want to see in addition to looking at what they actually engage with, you do get a completely different sense of what people value and you can build systems that approximate that. But going back to your question, I think rather than giving people a score of the friction--

Jonathan Zittrain: Yes.

Mark Zuckerberg: I think you can probably give people feedback of, "Hey, this might make people uncomfortable in this way, in this specific way." And this fits your--

Jonathan Zittrain: It might affect how much it gets-- how much it gets shared.

Mark Zuckerberg: Yeah. And this gets down to a different-- There's a different AI ethics question which I think is really important here, which is designing AI systems to be understandable by people

Jonathan Zittrain: Right.

Mark Zuckerberg: Right and to some degree, you don't just want it to spit out a score of how offensive or, like, where it scores on any given policy. You want it to be able to map to specific things that might be problematic.

Jonathan Zittrain: Yes.

Mark Zuckerberg: And that's the way that we're trying to design the systems overall.

Jonathan Zittrain: Yes. Now we have something parked in the box we should take out, which is the external review stuff. But before we do, one other just transparency thing maybe to broach. It basically just occurred to me, I imagine it might be possible to issue me a score of how much I've earned for Facebook this year. It could simply say, "This is how much we collected on the basis of you in particular being exposed to an ad." <laughs> And I know sometimes people, I guess, might compete to get their numbers up. <laughs> But I'm just curious, would that be a figure? I'd kind of be curious to know, in part because it might even lay the groundwork of being like, "Look, Mark, I'll double it. You can have double the money and then don't show me any ads." Can we get a car off of that lot today?

<laughter>

Mark Zuckerberg: <laughs> Okay, well, there's a lot-- <laughter>

Mark Zuckerberg: There's a lot in there.

Jonathan Zittrain: It was a quick question.

Mark Zuckerberg: <laughs> So there's a question in what you're saying which is so we built an ad-supported system. Should we have an option for people to pay to not see ads.

Jonathan Zittrain: Right.

Mark Zuckerberg: I think is kind of what you're saying. I mean, just as the basic primer from first principles on this. You know, we're building this service. We want to give everyone a voice. We want everyone to be able to connect with who they care about. If you're trying to build a service for everyone,--

Jonathan Zittrain: Got to be free. That's just <inaudible 01:04:45> +++ 12:08:17 +++

Mark Zuckerberg: If you want them to use it, that's just going to be the argument. Yes, yes. **Jonathan Zittrain:** Okay. All right.

Mark Zuckerberg: So then, so this is a kind of a tried and true thing. There are a lot of companies over time that have been ad supported. In general what we find is that if people are going to see ads, they want them to be relevant. They don't want them to be junk. Right. So then within that you give people control over how their data is used to show them ads. But the vast majority of people say, like, show me the most relevant ads that you can because I get that I have to see ads. This is a free service. So now the question is, all right, so there's a whole set of questions around that that we could get into, but but then--

Jonathan Zittrain: For which we did talk about enough to reopen it, the personalization exploitation.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: Or even just philosophical question. Right now, Uber or Lyft are not funded that way. We could apply this ad model to Uber or Lyft, "Free rides. Totally free. It's just every fifth ride takes you to Wendy's and idles outside the drive through window."

<laughter>

Jonathan Zittrain: "Totally up to you what you want to do, but you're going to sit here for a while," and then you go on your way. I don't know how-- and status quo-ism would probably say people would have a problem with that, but it would give people rides that otherwise wouldn't get rides.

Mark Zuckerberg: I have not thought about that case in their--

<laughter>

Mark Zuckerberg: In their business, so, so--

Jonathan Zittrain: Well, that's my patent, damn it, so don't you steal it.

Mark Zuckerberg: But certainly some services, I think tend themselves better towards being ad supported than others.

Jonathan Zittrain: Okay.

Mark Zuckerberg: Okay and I think generally information-based ones tend to-- **Jonathan Zittrain:** Than my false imprisonment hypo, I'd-- Okay, fair enough. **Mark Zuckerberg:** I mean, that seems--

Jonathan Zittrain: Yeah.

Mark Zuckerberg: There might be, you know, more-- <laughs> more issues there. But okay, but go to the subscription thing.

Jonathan Zittrain: Yes.

Mark Zuckerberg: When people have questions about the ad model on Facebook, I don't think the questions are just about the ad model, I think they're about both seeing ads and data use around ads. And the thing that I think, so when I think about this it's, I don't just think you want to let people pay to not see ads because I actually think then the question is the questions are around ads and data use and I don't think people are going to be that psyched about not seeing ads but then not having different controls over how their data is used. Okay, but now you start getting into a principle question which is are we going to let people pay to have different controls on data use than other people. And my answer to that is a hard no, right. So the prerequisite--

Jonathan Zittrain: What's an example of data use that isn't ad-based, just so we know what we're talking about?

Mark Zuckerberg: That isn't ad-based? **Jonathan Zittrain:** Yeah.

Mark Zuckerberg: Like what do you mean?

Jonathan Zittrain: You were saying, I don't want to see ads. But you're saying that's kind of just the wax on the car. What's underneath is how the data gets used.

Mark Zuckerberg: So, well, look-- Maybe-- let me keep going with this explanation and then I think this'll be clear.

Jonathan Zittrain: Yeah, sure.

Mark Zuckerberg: So one of the things that we've been working on is this tool that we call clear history. And the basic idea is it is you can kind of analogize it to a web browser where you can clear your cookies. That's kind of a normal thing. You know that when you clear your cookies you're going to get logged out of a bunch of stuff. A bunch of stuff might get more annoying.

Jonathan Zittrain: Which is why my guess is, am I right, probably nobody clears their cookies.

Mark Zuckerberg: I don't know.

Jonathan Zittrain: They might use incognito mode or something, but.

Mark Zuckerberg: I think-- I don't know. How many of you guys clear your cookies every once in a while, right?

Jonathan Zittrain: This is not a representative group, damn it.

<laughter>

Mark Zuckerberg: Okay. Like, maybe once a year or something I'll clear my cookies.

Jonathan Zittrain: <laughs>

Mark Zuckerberg: But no, it's, I think--

Jonathan Zittrain: Happy New Year.

<laughter>

Mark Zuckerberg: No, over some period of time, all right, but--

Jonathan Zittrain: Yeah, okay.

Mark Zuckerberg: But not necessarily every day. But it's important that people have that tool even though it might in a local sense make their experience worse.

Jonathan Zittrain: Yes.

Mark Zuckerberg: Okay. So that kind of content of what different services, websites and apps send to Facebook that, you know, we use to help measure the ads in effectiveness there, right, so things like, you know, if you're an app developer and you're trying to pay for ads to help grow your app, we want to only charge you when we actually, when something that we show leads to an install, not just whether someone sees the ad or clicks on it, but if they add--

Jonathan Zittrain: That requires a whole infrastructure to, yeah.

Mark Zuckerberg: Okay, so then, yeah, so you build that out. It helps us show people more relevant ads. It can help show more relevant content. Often a lot of these signals are super useful also on the security side for some of the other things that we've talked about, so that ends up being important. But fundamentally, you know, looking at the model today, it seems like you should have something like this ability to clear history. It turns out that it's a much more complex technical project. I'd talked about this at our developer conference last year, about how I'd hoped that we'd roll it out by the end of 2018 and just, the plumbing goes so deep into all the different systems that it's, that-- But we're still working on it and we're going to do it. It's just it's taking a little bit longer.

Jonathan Zittrain: So clear history basically means I am as if a newb, I just show-- **Mark Zuckerberg:** Yes.

Jonathan Zittrain: Even though I've been using Facebook for a while, it's as if it knows nothing about me and it starts accreting again.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: And I'm just trying to think just as a plain old citizen, how would I make an informed

judgment about how often to do that or when I should do it? What--?

Mark Zuckerberg: Well, hold on. Let's go to that in a second.

Jonathan Zittrain: Okay.

Mark Zuckerberg: But one thing, just to connect the dots on the last conversation.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: Clear history is a prerequisite, I think, for being able to do anything like subscriptions. Right. Because, like, partially what someone would want to do if they were going to really actually pay for a not ad supported version where their data wasn't being used in a system like that, you would want to have a control so that Facebook didn't have access or wasn't using that data or associating it with your account. And as a principled matter, we are not going to just offer a control like that to people who pay. Right. That's going to, if we're going to give controls over data use, we're going to do that for everyone in the community. So that's the first thing that I think we need to go do.

Mark Zuckerberg: So that's, so that's kind of-- This is sort of the how we're thinking about the projects and this is a really deep and big technical project but we're committed to doing it because I think it's that's what it's there for. [ph?]

Jonathan Zittrain: And I guess like an ad block or somebody could then write a little script for your browser that would just clear your history every time you visit or something.

Mark Zuckerberg: Oh, yeah, no, but the plan would also be to offer something that's an ongoing thing.

Jonathan Zittrain: I see.

Mark Zuckerberg: In your browser, but I think the analogy here is you kind of have, in your browser you have the ability to clear your cookies. And then, like, in some other place you have under your, like, nuclear settings, like, don't ever accept any cookies in my browser. And it's like, all right, your browser's not really going to work that well.

Jonathan Zittrain: Yeah. <laughs>

Mark Zuckerberg: But, but you can do that if you want because you should have that control. I think that these are part and parcel, right. It's I think a lot of people might go and clear their history on a periodic basis because they-- Or, or actually in the research that we've done on this as we've been developing it, the real thing that people have told us that they want is similar to cookie management, not necessarily wiping everything, because that ends in inconvenience of getting logged out of a bunch of things, but there are just certain services or apps that you don't want that data to be connected to your Facebook account. So having the ability on an ad hoc basis to go through and say, "Hey, stop associating this thing," is going to end up being a quite important thing that I think we want to try to deliver. So that's, this is partially as we're getting into this, it's a more complex thing but I think it's very valuable. And I think if any conversation around the-- around subscriptions, I think you would want to start with giving people these, make sure that everyone has these kind of controls. So that's, we're kind of in the early phases of doing that. The philosophical downstream question of whether you also let people pay to not have ads, I don't know. There were a bunch of questions around whether that's actually a good thing, but I personally don't believe that very many people would like to pay to not have ads. That all of the research that we have, it's it may still end up being the right thing to offer that as a choice down the line, but all of the data that I've seen suggests that the vast, vast, vast majority of people want a free service and that the ads, in a lot of places are not even that different from the organic content in terms of the quality of what people are being able to see.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: People like being able to get information from local businesses and things like that too, so. So there's a lot of good there.

Jonathan Zittrain: Yeah. Forty years ago it would have been the question of ABC versus HBO and the answer turned out to be yes.

<laughter>

Jonathan Zittrain: So you're right. And people might have different things. **Mark Zuckerberg:** Yeah.

Jonathan Zittrain: There's a little paradox lingering in there about if something's so important and vital that we wouldn't want to deprive anybody of access to it but therefore nobody gets it until we figured out how to remove it for everybody.

Mark Zuckerberg: What we-- [ph?]

Jonathan Zittrain: In other words, if I could buy my way out of ads and data collection it wouldn't be fair to those who can't and therefore we all subsist with it until the advances you're talking about.

Mark Zuckerberg: Yeah, but I guess what I'm saying is on the data use, I don't believe that that's something that people should buy. I think the data principles that we have need to be uniformly available to everyone. That to me is a really important principle. It's, like, maybe you could have a conversation about whether you should be able to pay and not see ads. That doesn't feel like a moral question to me.

Jonathan Zittrain: Yes.

Mark Zuckerberg: But the question of whether you can pay to have different privacy controls feels wrong. So that to me is something that in any conversation about whether we'd evolve towards having a subscription service, I think you have to have these controls first and it's a very deep thing. A technical problem to go do, but we're-- that's why we're working through that.

Jonathan Zittrain: Yes. So long as the privacy controls that we're not able to buy our way into aren't controls that people ought to have. You know, it's just the kind of underlying question of is the system as it is that we can't opt out of a fair system. And that's of course, you know, you have to go into the details to figure out what you mean by it. But let's in the remaining time we have left--

Mark Zuckerberg: How are we doing on time?

Jonathan Zittrain: We're good. We're 76 minutes in.

Mark Zuckerberg: All right, into--

<laughter>

Mark Zuckerberg: We're going to get through maybe half the topics. **Jonathan Zittrain:** Yeah, yeah, yeah.

Mark Zuckerberg: And I'll come back and do another one later.

Jonathan Zittrain: I'm going to bring this in for a landing soon. On my agenda left includes such things as taking out of the box the independent review stuff, chat a little bit about that. I'd be curious, and this might be a nice thing, really, as we wrap up, which would be a sense of any vision you have for what would Facebook look like in 10 or 15 years and how different would it be than the Facebook of 10 years ago is compared to today. So that's something I'd want to talk about. Is there anything big on your list that you want to make sure we talk about?

Mark Zuckerberg: Those are good. Those are good topics.

Jonathan Zittrain: <laughs> Fair enough.

Mark Zuckerberg: <laughs>

Jonathan Zittrain: So all right, the external review board.

Mark Zuckerberg: Yeah. So one of the big questions that I have just been thinking about is, you know, we make a lot of decisions around content enforcement and what stays up and what comes down. And having gone through this process over the last few years of working on the systems, one of the themes that I feel really strongly about is that we shouldn't be making so many of these decisions ourselves. You know, one of the ways that I try to reason about this stuff is take myself out of the position of being CEO of the company, almost like a Rawlsian perspective. If I was a different person, what would I want the CEO of the company to be able to do? And I would not want so many decisions about content to be concentrated with any individual. So--

Jonathan Zittrain: It is weird to see big impactful, to use a terrible word, decisions about what a huge swath of humanity does or doesn't see inevitably handled as, like, a customer service issue. It does feel like a mismatch, which is what I hear you saying.

Mark Zuckerberg: So let's, yeah, so I actually think the customer service analogy is a really interesting one. Right. So when you email Amazon, because they don't, they make a mistake with your package, that's customer support. Right. I mean, they are trying to provide a service and generally, they can invest more in customer support and make people happier. We're doing something completely different, right. When someone emails us with an issue or flags some content, they're basically complaining about something that someone else in the community did. So it's more like it's almost more like a court system in that sense. Doing more of that does not make people happy because in every one of those transactions one person ends up the winner and one is the loser. Either you said that that content, that the content was fine, in which case the person complaining is upset, or you the someone's content down, in which case the person is really upset because you're now telling them that they don't have the ability to express something that they feel is a valid thing that they should be able to express.

So in some deep sense while some amount of what we do is customer support, people get locked out of their account, et cetera, you know, we now have, like, more than 30,000 people working on content review and safety review, doing the kind of judgments that, you know, it's basically a lot of the stuff, we have machine learning systems that flag things that could be problematic in addition to people in the community flagging things, but making these assessments of whether the stuff is right or not. So one of the questions that I just think about, it's like, okay, well, you have many people doing this.

Regardless of how much training they have, we're going to make mistakes, right. So you want to start building in principles around, you know, what you would kind of think of as due process, right. So we're building in an ability to have an appeal, right, which already is quite good in that we are able to overturn a bunch of mistakes that the first line people make in making these assessments. But at some level I think you also

want a level of kind of independent appeal, right, where if, okay, let's say, so the appeals go to maybe a higher level of Facebook employee who is a little more trained in the nuances of the policies; but at some point, I think you also need an appeal to an independent groups, which is, like, is this policy fair? Was this--? Like is this piece of content really getting on the wrong side of the balance of free expression and safety? And I just don't think at the end of the day that that's something that you want centralized in a

single company. So now the question is how do you design that system and that's a real question, right, so that we don't pretend to have the answers on this. What we're basically working through is we have a draft proposal and we're working with a lot of experts around the world to run a few pilots in the first half of this year that can hopefully we can codify into something that's a longer term thing. But I just, I believe that this is just an incredibly important thing. As a person and if I take aside the role that I have as CEO of the company, I do not want the company being able to make all of those final decisions without a check and balance and accountability, so I want to use the position that I'm in to help build that kind of an institution.

Jonathan Zittrain: Yes. And when we talk about an appeal, then, it sounds like you could appeal two distinct things. One is this was the rule but it was applied wrong to me. This, in fact, was parody [ph?] so it shouldn't be seen as near the line.

Mark Zuckerberg: Yeah.

Jonathan Zittrain: And I want the independent body to look at that. The other would be the rule is wrong. The rule should change because--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: And you're thinking the independent body could weigh in on both of those?

Mark Zuckerberg: Yeah. Over time, I would like the role of the independent oversight board to be able to expand to do additional things as well. I think the question is it's hard enough to even set something up that's going to codify the values that we have around expression and safety on a relatively defined topic.

Jonathan Zittrain: Yeah.

Mark Zuckerberg: So I think the question is if you kind of view this as an experiment in institution building where we're trying to build this thing that is going to have real power to--

Jonathan Zittrain: Yes.

Mark Zuckerberg: I mean, like, I will not be able to make a decision that overturns what they say. Which I think is good. I think also just it raises the stakes. You need to make sure we get this right, so.

Jonathan Zittrain: It's fascinating. It's huge. I think the way you're describing it, I wouldn't want to understate--

Mark Zuckerberg: Yeah.

Jonathan Zittrain: That this is not a usual way of doing business.

Mark Zuckerberg: Yeah, but I think it-- I think this is-- I really care about getting this right. **Jonathan Zittrain:** Yeah.

Mark Zuckerberg: But I think you want to start with something that's relatively well-defined and then hopefully expand it to be able to cover more things over time. So in the beginning I think one question that could come up is my understanding, I mean, it's always dangerous talking about legal precedence when I'm, this might be one of my first times at Harvard Law School. I did not spend a lot of time here--

<laughter>

Mark Zuckerberg: When I was an undergrad. But, you know what I mean, the, if the Supreme Court overturns something, they don't tell Congress what the law should be, they just say there's an issue here, right. And then basically there's a process. All right. So if I'm getting that wrong--

<laughter>

Mark Zuckerberg: All right. I shouldn't have done that.

Jonathan Zittrain: No, no. That's quite honest. [ph?]

Mark Zuckerberg: I knew that was dangerous.

<laughter>

Mark Zuckerberg: And that that was a mistake.

Jonathan Zittrain: There are people who do agree with you.

<laughter>

Mark Zuckerberg: Okay. Oh, so that's an open question that that's how it works. **Jonathan Zittrain:** It's a highly debated question, yes.

Mark Zuckerberg: All right.

Jonathan Zittrain: There's the I'm just the umpire calling balls and strikes and in fact, the first type of question we brought up, which was, "Hey, we get this is the standard. Does it apply here?" lends itself a little more to, you know, you get three swings and if you miss them all, like, you can't keep playing. The umpire can usher you away from the home plate. This is, I'm really digging deep into my knowledge now of baseball. There's another thing about, like,--

Mark Zuckerberg: That's okay. I'm not the person who's going to call you out on getting something wrong there.

Jonathan Zittrain: I appreciate that.

Mark Zuckerberg: That's why I also need to have a librarian next to me at all times. <laughter>

Jonathan Zittrain: Very good. I wonder how much librarians tend to know about baseball. <laughter>

Mark Zuckerberg: Aww.

Jonathan Zittrain: But we digress. Ah, we're going to get letters.

Mark Zuckerberg: Yeah.

<laughter>

Jonathan Zittrain: mentions. But whether or not the game is actually any good with a three strikes rule, maybe there should be two or four or whatever, starts to ask of the umpire more than just, you know, your best sense of how that play just went. Both may be something. Both are surely beyond standard customer service issues, so both could maybe be usefully externalized. What you'd ask the board to do in the category one kind of stuff maybe it's true that, like, professional umpirage [ph?] could help us and there are people who are jurists who can do that worldwide. For the other, whether it's the Supreme--

Jonathan Zittrain: --court, or the so-called common law and state courts where often a state supreme court will be like, "Henceforth, 50 feet needs to be the height of a baseball net," and like, "If you don't agree, Legislature, we'll hear from you, but until then it's 50 feet." They really do kind of get into the weeds. They derive maybe some legitimacy for decisions like that from being close to their communities, and it really regresses them to a question of: Is Facebook a global community, a community of 2.X billion people worldwide, transcending any national boundaries, and for which I think so far on these issues, it's meant to be, "The rule is the rule," it doesn't really change in terms of service from one place to another--

versus how much do we think of it as somehow localized-- whether or not localized through government-- but where different local communities make their own judgments?

Mark Zuckerberg: That is one of the big questions. I mean, right now we have community standards that are global. We follow local laws, as you say. But I think the idea is-- I don't think we want to end up in a place where we have very different norms in different places, but you want to have some sense of representation and making sure that the body that can deliberate on this has a good diversity of views. So these are a lot of the things that we're trying to figure out, is like: Well, how big is the body? When decisions are made, are they made by the whole body, or do you have panels of people that are smaller sets? If there are panels, how do you make sure that you're not just getting a random sample that kind of skews in the values perspective towards one thing? So then there a bunch of mechanisms like, okay, maybe one panel that's randomly constituted decides on whether the board will take up a question or one of the issues, but then a separate random panel of the group actually does the decisions, so that way you eliminate some risk that any given panel is going to be too ideologically skewed. So there's a bunch of things that I think we need to think through and work through, but the goal on this is to, over time, have it grow into something that can provide greater accountability and oversight to potentially more of the hard questions that we face, but I think it's so high-stakes that starting with something that's relatively defined is going to be the right way to go in the beginning. So regardless of the fact that I was unaware of the controversy around the legal point that I made a second ago, I do think in our case it makes sense to start with not having this group say what the policies are going to be, but just have there be-- have it be able to say, "Hey, we think that you guys are on the wrong side on this, and maybe you should rethink where the policy is because we think you're on the wrong side." There's one other thing that I think is worth calling out, which is in a typical kind of judicial analog, or at least here in the U.S., my understanding, is there's the kind of appeal route to the independent board considering an issue, but I also think that we want to have an avenue where we as the company can also just raise hard issues that come up to the board without having-- which I don't actually know if there's any mechanism for that.

Jonathan Zittrain: It's called an advisory opinion.

<audience laughter>

Jonathan Zittrain: But under U.S. federal law, it's not allowed because of Article III Case or Controversy requirement, but state courts do it all the time. You'll have a federal court sometimes say-- because it's a federal court but it's deciding something under state law. It'll be like, "I don't know, ask Florida." And they'll be like, "Hey Florida," and then Florida is just Florida.

<audience laughter>

Mark Zuckerberg: Sure. So I think that--

Jonathan Zittrain: So you can do an advisory opinion.

Mark Zuckerberg: --that'll end up being an important part of this too. We're never going to be able to get out of the business of making frontline judgments. We'll have the AI systems flag content that they think is against policies or could be, and then we'll have people-- this set of 30 thousand people, which is growing-- that is trained to basically understand what the policies are. We have to make the frontline decisions, because a lot of this stuff needs to get handled in a timely way, and a more deliberative process that's thinking about the fairness and the policies overall should happen over a different timeframe than what is often relevant, which is the enforcement of the initial policy. But I do think overall for a lot of the biggest questions, I just want to build a more independent process.

Jonathan Zittrain: Well, as you say, it's an area with fractal complexity in the best of ways, and it really is

terra incognita, and it'd be exciting to see how it might be built out. I imagine there's a number of law professors around the world, including some who come from civil rather than common law jurisdictions, who are like, "This is how it works over here," from which you could draw. Another lingering question would be-- lawyers often have a bad reputation. I have no idea why. But they often are the glue for a system like this so that a judge does not have to be oracular or omniscient. There's a process where the lawyer for one side does a ton of work and looks at prior decisions of this board and says, "Well, this is what would be consistent," and the other lawyer comes back, and then the judge just gets to decide between the two, rather than having to just know everything. There's a huge tradeoff here for every appealed content decision, how much do we want to build it into a case, and you need experts to help the parties, versus they each just sort of come before Solomon and say, "This kind of happened," and-- or Judge Judy maybe is a more contemporary reference.

Mark Zuckerberg: Somewhere between the two, yeah.

<audience laughter>

Jonathan Zittrain: Yeah. So it's a lot of stuff-- and for me, I both find myself-- I don't know if this is the definition of prurient-- both excited by it and somewhat terrified by it, but very much saying that it's better than a status quo, which is where I think you and I are completely agreeing, and maybe a model for other firms out there. So that's the last question in this area that pops to my mind, which is: What part of what you're developing at Facebook-- a lot of which is really resource-intensive-- is best thought of as a public good to be shared, including among basically competitors, versus, "That's part of our comparative advantage and our secret sauce"? If you develop a particularly good algorithm that can really well detect fake news or spammers or bad actors-- you've got the PhDs, you've got the processors-- is that like, "In your face, Schmitter [ph?]," or is like, "We should have somebody that-- some body-- that can help democratize that advance"? And it could be the same to be said for these content decisions. How do you think about that?

Mark Zuckerberg: Yeah, so certainly the threat-sharing and security work that you just referenced is a good area where there's much better collaboration now than there was historically. I think that that's just because everyone recognizes that it's such a more important issue. And by the way, there's much better collaboration with governments now too on this, and not just our own here in the U.S., and law enforcement, but around the world with election commissions and law enforcement, because there's just a broad awareness that these are issues and that--

Jonathan Zittrain: Especially if you have state actors in the mix as the adversary.

Mark Zuckerberg: Yes. So that's certainly an area where there's much better collaboration now, and that's good. There's still issues. For example, if you're law enforcement or intelligence and you have developed a-- "source" is not the right word-- but basically if you've identified someone as a source of signals that you can watch and learn about, then you may not want to come to us and tell us, "Hey, we've identified that this state actor is doing this bad thing," because then the natural thing that we're going to want to do is make sure that they're not on our system doing bad things, or that they're not-- either they're not in the system at all or that we're interfering with the bad things that they're trying to do. So there's some mismatch of incentives, but as you build up the relationships and trust, you can get to that kind of a relationship where they can also flag for you, "Hey, this is what we're at." So I just think having that kind of baseline where you build that up over time is helpful. And I think on security and safety is probably the biggest area of that kind of collaboration now, across all the different types of threats; not just election and democratic process type stuff, but any kind of safety issue. The other area where I tend to think about what we're doing is-- it should be open-- is just technical infrastructure overall. I mean, that is probably a less controversial piece, but we open-source a lot of the basic

stuff that runs our systems, and I think that that is a-- that's a contribution that I'm quite proud of that we do. We have sort of pioneered this way of thinking about how people connect, and the data model around that is more of a graph, and the idea of graph database and a lot of the infrastructure for being able to efficiently access that kind of content I think is broadly applicable beyond the context of a social network. When I was here as an undergrad, even though I wasn't here for very long, I studied psychology and computer science, and to me-- I mean, my grounding philosophy on this stuff is that basically people should be at the center of more of the technology that we build. I mean, one of the early things that I kind of recognized when I was a student was like-- at the time, there were internet sites for finding almost anything you cared about, whether it's books or music or news or information or businesses-- but as people, we think about the world primarily in terms of other people, not in terms of other objects, not cutting things up in terms of content or commerce or politics or different things, but it's like-- the stuff should be organized around the connections that people have, where people are at the centerpiece of that, and one of the missions that I care about is over time just pushing more technology development in the tech industry overall to develop things with that mindset. I think-- and this is a little bit of a tangent-- but the way that our phones work today, and all computing systems, organized around apps and tasks is fundamentally not how people-- how our brains work and how we approach the world. It's not-- so that's one of the reasons why I'm just very excited longer-term about especially things like augmented reality, because it'll give us a platform that I think actually is how we think about stuff. We'll be able to bring the computational objects into the world but fundamentally we'll be interacting as people around them. The whole thing won't be organized around an app or a task; it'll be organized around people, and that I think is a much more natural and human system for how our technology should be organized. So open-sourcing all of that infrastructure-- to do that, and enabling not just us but other companies to kind of get that mindset into more of their thinking and the technical underpinning of that, is just something that I care really deeply about.

Jonathan Zittrain: Well, this is nice, and this is bringing us in for our landing, because we're talking about 10, 20, 30 years ahead. As a term of art, I understand augmented reality to mean, "I've got a visor"-- version 0.1 was Google Glass-- something where I'm kind of out in the world but I'm literally online at the same time because there's data coming at me in some-- that's what you're talking about, correct?

Mark Zuckerberg: Yeah, although it really should be glasses like what you have. I think we'll probably-- maybe they'll have to be a little bigger, but not too much bigger or else it would start to get weird.

<audience laughter>

Mark Zuckerberg: So I don't think a visor is going to catch. I don't think anyone is psyched about that feature.

Jonathan Zittrain: And anything involving surgery starts to sound a little bad too. **Mark Zuckerberg:** No, no, we're definitely focused on--

<audience laughter>

Mark Zuckerberg: --on external things. Although--

Jonathan Zittrain: Like, "Don't make news, don't make news, don't make news."

<audience laughter>

Mark Zuckerberg: No, no, no. Although we have showed this demo of basically can someone type by thinking, and of course when you're talking about brain-computer interfaces, there's two dimensions of that work. There's the external stuff, and there's the internal stuff, and invasive, and yes, of course if

you're actually trying to build things that everyone is going to use, you're going to want to focus on the noninvasive things.

Jonathan Zittrain: Yes. Can you type by thinking? **Mark Zuckerberg:** You can.

Jonathan Zittrain: It's called a Ouija Board. No. But you're subvocalizing enough or there's enough of a read of--

Mark Zuckerberg: No, no, no. So there's actually a bunch of the research here-- there's a question of throughput and how quickly can you type and how many bits can you express efficiently, but the basic foundation for the research is someone-- a bunch of folks who are doing this research showed a bunch of people images-- I think it was animals-- so, "Here's an elephant, here's a giraffe"-- while having kind of a net on their head, noninvasive, but shining light and therefore looking at the level of blood activity and-- just blood flow and activity in the brain-- trained a machine learning basically on what the pattern of that imagery looked like when the person was looking at different animals, then told the person to think about an animal, right? So think about-- just pick one of the animals to think about, and can predict what the person was thinking about in broad strokes just based on matching the neural activity. So the question is, so you can use that to type.

Jonathan Zittrain: Fifth amendment implications are staggering. <audience laughter>

Jonathan Zittrain: Sorry.

Mark Zuckerberg: Well, yes. I mean, presumably this would be something that someone would choose to use a product. I'm not-- yeah, yeah. I mean, yes, there's of course all the other implications, but yeah, I think that this is going to be-- that's going to be an interesting thing down the line.

Jonathan Zittrain: But basically your vision then for a future--

Mark Zuckerberg: I don't know how we got onto that.

<audience laughter>

Jonathan Zittrain: You can't blame me. I think you brought this up.

Mark Zuckerberg: I did, but of all the things that-- I mean, this is exciting, but we haven't even covered yet how we should talk about-- tech regulation and all this stuff I figured we'd get into. I mean, we'll be here for like six or seven hours. I don't know how many days you want to spend here to talking about this, but--

Jonathan Zittrain: "We're here at the Zuckerberg Center and hostage crisis."

<audience laughter>

Jonathan Zittrain: "The building is surrounded."

Mark Zuckerberg: Yeah. But I think a little bit on future tech and research is interesting too, so.

Jonathan Zittrain: Please.

Mark Zuckerberg: Yeah, we're good.

Jonathan Zittrain: Oh, we did cover it, is what you're saying.

<audience laughter>

Mark Zuckerberg: I mean, but going back to your question about what-- if this is the last topic-- what I'm excited about for the next 10 or 20 years-- I do think over the long term, reshaping our computing platforms to be fundamentally more about people and how we process the world is a really fundamental thing. Over the nearer term-- so call it five years-- I think the clear trend is towards more private communication. If you look at all of the different ways that people want to share and communicate across the internet-- but we have a good sense of the cross-strength, everything from one-on-one messages to kind of broadcasting publicly-- the thing that is growing the fastest is private communication. Right?

So between WhatsApp and Messenger, and Instagram now, just the number of private messages-- it's about 100 billion a day through those systems alone, growing very quickly, growing much faster than the amount that people want to share or broadcast into a feed-type system. Of the type of broadcast content that people are doing, the thing that is growing by far the fastest is stories. Right?

So ephemeral sharing of, "I'm going to put this out, but I want to have a timeframe after which the data goes away." So I think that that just gives you a sense of where the hub of social activity is going. It also is how we think about the strategy of the company. I mean, people-- when we talk about privacy, I think a lot of the questions are often about privacy policies and legal or policy-type things, and privacy as a thing not to be breached, and making sure that you're within the balance of what is good. But I actually think that there's a much more-- there's another element of this that's really fundamental, which is that people want tools that give them new contexts to communicate, and that's also fundamentally about giving people power through privacy, not just not violating privacy, right? So not violating privacy is a backstop, but actually-- you can kind of think about all the success that Facebook has had-- this is kind of a counterintuitive thing-- has been because we've given people new private or semi-private ways to communicate things that they wouldn't have had before.

So thinking about Facebook as an innovator in privacy is certainly not the mainstream view, but going back to the very first thing that we did, making it so Harvard students could communicate in a way that they had some confidence that their content and information would be shared with only people within that community, there was no way that people had to communicate stuff at that scale, but not have it either be completely public or with just a small set of people before. And people's desire to be understood and express themselves and be able to communicate with all different kinds of groups is, in the experience that I've had, nearly unbounded, and if you can give people new ways to be able to communicate safely and express themselves, then that is something that people just have a deep thirst and desire for.

So encryption is really important, because I mean, we take for granted in the U.S. that there's good rule of law, and that the government isn't too much in our business, but in a lot of places around the world, especially where WhatsApp is the biggest, people can't take that for granted. So having it so that you really have confidence that you're sharing something one-on-one and it's not-- and it really is one-on-one, it's not one-on-one and the government there-- actually makes it so people can share things that they wouldn't be comfortable otherwise doing it. That's power that you're giving people through building privacy innovations.

Stories I just think is another example of this, where there are a lot of things that people don't want as part of the permanent record but want to express, and it's not an accident that that is becoming the primary way that people want to share with all of their friends, not putting something in a feed that goes on their permanent record. There will always be a use for that too-- people want to have a record and there's a lot of value that you can build around that-- you can have longer-term discussions-- it's harder to do that around stories. There's different value for these things. But over the next five years, I think we're going to see all of social networking kind of be reconstituted around this base of private communication, and that's something that I'm just very excited about. I think that that's-- it's going to unlock a lot of people's ability to express themselves and communicate things that they haven't had the tools to do before, and it's going to be the foundation for building a lot of really important tools on top of that too.

Jonathan Zittrain: That's so interesting to me. I would not have predicted that direction for the next five years. I would have figured, "Gosh, if you already know with whom you want to speak, there are so many tools to speak with them," some of which are end-to-end, some of which aren't, some of which are roll-your-own and open-source, and there's always a way to try to make that easier and better, but that feels a little bit to me like a kind of crowded space, not yet knowing of the innovations that might lie ahead and means of communicating with the people you already know you want to talk to. And for that, as you say, if that's where it's at, you're right that encryption is going to be a big question, and otherwise technical design so that if the law comes knocking on the door, what would the company be in a position to say. This is the Apple iPhone Cupertino-- sorry, San Bernardino case-- and it also calls to mind will there be peer-to-peer implementations of the things you're thinking about that might not even need the server at all, and it's basically just an app that people use, and if it's going to deliver an ad, it can still do that app-side, and how much governments will abide it. They have not, for the most part, demanded technology mandates to reshape how the technology works. They're just saying, "If you've got it"-- in part you've got it because you want to serve ads-- "we want it." But if you don't even have it, it's been rare for the governments to say, "Well, you've got to build your system to do it." It did happen with the telephone system back in the day. CALEA, the Communications Assistance to Law Enforcement Act, did have federal law in the United States saying, "If you're in the business of building a phone network, AT&T, you've got to make it so we can plug in as you go digital," and we haven't yet seen those mandates in the internet software side so much. So we can see that coming up again. But it's so funny, because if you'd asked me, I would have figured it's encountering people you haven't met before and interacting with them, for which all of the stuff about air traffic control of what goes into your feed and how much your stuff gets shared-- all of those issues start to rise to the fore, and it gets me thinking about, "I ought to be able to make a feed recipe that's my recipe, and fills it according to Facebook variables, but I get to say what the variables are." But I could see that if you're just thinking about people communicating with the people they already know and like, that is a very different realm.

Mark Zuckerberg: It's not necessarily-- it's not just the people that you already know. I do think-- we've really focused on friends and family for the last 10 or 15 years, and I think a big part of what we're going to focus on now is around building communities in different ways and all the utility that you can build on top of, once you have a network like this in place. So everything from how people can do commerce better to things like dating, which is-- a lot of dating happens on our services, but we haven't built any tools specifically for that.

Jonathan Zittrain: I do remember the Facebook joint experiment-- "experiment" is such a terrible word-- study, by which one could predict when two Facebook members are going to declare themselves in a relationship, months ahead of the actual declaration. I was thinking some of the ancillary products were in-laws.

Mark Zuckerberg: That was very early. Yeah. So you're right that a lot of this is going to be about utility that you can build on top of it, but a lot of these things are fundamentally private, right? So if you're thinking about commerce, that people have a higher expectation for privacy, and the question is: Is the right context for that going to be around an app like Facebook, which is broad, or an app like Instagram? I think part of it is-- the discovery part of it, I think we'll be very well served there-- but then we'll also transition to something that people want to be more private and secure. Anyhow, we could probably go on for many hours on this, but maybe we should save this for the Round 2 of this that we'll do in the future.

Jonathan Zittrain: Indeed. So thanks so much for coming out, for talking at such length, for covering such a kaleidoscopic range of topics, and we look forward to the next time we see you.

Mark Zuckerberg: Yeah. Thanks.

Jonathan Zittrain: Thanks.

<audience applause>