



TECHNOLOGY
POLICY
INSTITUTE

2023 TPI Aspen Forum Academic Keynote with Professor Matthew Gentzkow

Scott Wallsten:

All right. Good morning. Good morning, everyone. We're going to keep playing that music all day because it's so inspiring. I hope you all slept well and you're ready for an intense day of tech policy discussions, tech hub fun, and then exploring Aspen later today. And remember, like I said, we have a packed schedule, so we're going to stay on time. Speakers, remember the taser I mentioned last night. Jane's around here somewhere. But it's going to be a great day.

We've got panels, keynotes, firesides, nice break in the afternoon for some outdoor exploration. We'll have the working dinners, the tech hub of course, a dessert bar and a happy hour at the end of it all, sponsored by Net Choice at Hooch, which will be great. Now, the only one substantive announcement I want to make here, which is about Sli.do. This is how we take questions. We will ask for questions, and you can raise your hand and people will call on you, the old-fashioned way. We love that, someone will come around with a microphone. But there's no reason to wait until the Q&A period. If you go to this platform, you can, through either of those, TinyURL/tpass@2023, or go to SLI.DO and type in TPI Aspen 20... No, just TPI Aspen, or the QR code.

It will take you to the platform, where you can type in questions. And they will show up and we can work them into the conversation too. So you don't have to wait until the question time. And just to get it going, we'll do a little poll on here, so you can sort of see how it works. Nathaniel, do you have the poll ready? Yeah, can you put it up there? Oh, they've put it up there? Somebody put it up there.

Matthew Gentzkow:

Yeah. Okay, good. Oh, so right, that's not it. Oh, here it is. No, that's not it either. I'll just assume you can see it. The poll is, should AI be regulated? Now we know that's kind of a silly question, a little bit broad. But anyways, it's just to show how it works. So you can enter your votes and hopefully it'll show up there. But mostly we're not using it for polls, we're using it for questions. But this is to make sure everybody gets on the screen and that you can ask all your questions there. I don't know what's happening with that. It doesn't matter. It'll show up. The questions will show up. Nathaniel, do you know why it's not showing results? Okay. We won't worry about it. But put your questions in there because the questions, we will see. And you can also see other people's questions as they come in. And I think you can upvote questions too, if you think that they're better than others. So they're more likely to be answered. Okay. Now let's get onto the program.

Or to society more broadly. It's not hurting me, it's just hurting society. And the top of that list is this whole discussion about the political impacts of social media, impacts on things like polarization and misinformation, and people's perception of the integrity of elections, and so on. As well as other things like hate crimes and ethnic violence, and so on. Those are all external effects of my use of social media on others. And then the other part of the conversation is about things that we would sometimes call internalities, which are behavioral biases or mistakes that people make, where not only might my using social media be hurting all kinds of other people, but it might actually be hurting me in some ways.

Or in other words, people might be making decisions about how many hours they spend watching videos on TikTok or whatever, that are not even necessarily maximizing their own wellbeing and utility correctly. And so, that includes conversations about addiction or their addictive qualities to social media, as well as discussions of impacts on mental health and wellbeing, particularly pertaining to youth and kids, loneliness, addiction, depression, suicide and so on. So we have this long list of harms.

There has been a ton of work on this, and I would say one thing... Discussions of this topic are often kind of depressing, in general. One good thing I think, as a note of optimism here, is our ability to learn about the answers to these questions and to understand what is happening, is way, way better than it was in any previous generation of media technologies. Many of these are kinds of issues that have been discussed in the context of different technologies for a long time, and we really have powerful tools now, experiments that we can run, data that we can gather, ability to monitor and see what's going on. So I think the potential for research to answer these questions is really big. So as I said, I want to give selective, subjective spin through, what do I see as the big picture here?

And given the short time, I'm going to zoom in on the externality side, on the specific question of, how has, and is social media impacting polarization? And on the more internality side, on what are the impacts on mental health and wellbeing, including on kids? And to give you just the kind of takeaway at the beginning, my overall view is that the evidence we have has consistently moved me in the direction of being less worried about the political impacts of social media. Less worried about the impacts of social media on polarization. And the evidence is, if anything, made me more worried about the impacts in this other category, on wellbeing and mental health. Okay. So I'll talk through evidence in each of these categories. So starting with polarization. There is now a tremendous amount of evidence that polarization, measured in a variety of different ways, has in fact, increased dramatically in the US, as well as in a number of other places around the world.

So there was a big debate about this. And there was... In political science, circa 2005, 2008, 2010, there was a view that the whole perception that polarization was increasing might be some kind of myth, that it's not really something we can see in the data. Since then, we can see it in the data very, very clearly. And this trend is very real. And when I'm going to say, I may be a little bit less worried about the impact of social media on polarization, do not mistake that to mean I am not worried about polarization. Because I think that is an incredibly important issue right now, and obviously the potential social harms of that are very big. Just to give you a flavor of the kinds of data points, these are just facts from Pew surveys. One way that you can look at this is in terms of people's views on political issues and political values.

That is something that in general, has been quite stable in the US, over time. But this is showing distributions of an index of a bunch of different ideological views, by party. And you can see, going from 2004 to 2017, this quite dramatic separation in those views. Incidentally, a lot of that is about correlation among the views, more so than any individual view becoming more polarized. So what's happened is people have become a lot more sorted in terms of their views. You're much less likely to have people who have a mix of views, some conservative, some liberal, than you did before. A very different way to look at things is what's called in the literature, affective polarization. How do people feel about people on the other side? Those feelings, as measured in different ways, have become dramatically more negative over time. And so, you can look at this data in lots of different ways.

There are things that... There are ways of looking at polarization that it has not increased, but overall, these increases are quite large and quite clear, and quite worrying. So the question is, what does all this have to do with social media? One thing that you can see in the data, first of all, that I don't really have on the slide... I guess you can sort of see it from this picture even. Is,

this is a trend that started before social media existed. So there's some part of it that we know for sure had nothing to do with social media because it happened before the early 2000s, before the introduction of these things. And if you look at an overall trend, you might see some inflection point around the early 2000s, but that doesn't stand out. It looks like something that has been going on for quite a long time.

Now I want to show you two facts that... This comes from some of my earlier work with colleagues. These are not causal experimental estimates, but just descriptive facts that I find quite diagnostic, or at least provocative and interesting to think about with this question. So the first is, suppose we take measures, these kinds of measures of polarization, and break them down across demographic groups, and ask, who are the groups, say within the US, for whom polarization has increased the most? And you might think about, what would you expect that to look like, if social media was a key driver? If social media is the thing making people more polarized, who would you expect to be getting more polarized, or most polarized? So you could think about that on lots of different dimensions. A really obvious one you might think about is age. So social media use, particularly until quite recently, is highly correlated with age.

Until quite recently, there were very few people say, over 60, over 70, who were using social media to a significant degree, certainly before 2016. And so, one thing you might think is, if social media is a key driver, we might see the biggest increases in polarization among younger people who are more likely to use it. And so, we compiled data on a bunch of different measures of polarization, put them together into a big index. I'm not going to tell you how all that happened, but the overall picture is, if you just break down those trends by age, you see that if anything, it is the oldest people, older people who have seen the biggest increase in polarization.

So that is not dispositive, it doesn't... You're all smart people. You can tell me a long list of reasons why social media could nevertheless, be the key driver here, even though this picture is true. But it's not the picture that you would've most immediately imagined in a world where this is all about social media. Something must be going on with these 75 plus year olds, who I can tell you were not using social media to any meaningful degree over most of this time.

Second fact that I think is kind of diagnostic is, what is going on in other countries? And something we know about social media is, it is available and widely used, and has grown kind of at similar rates in most countries, and certainly in most OECD countries. And so again, just loosely, if you thought social media is key driver of polarization, you might expect to see similar increasing trends in polarization everywhere. So in this work, we went and tried to gather... It's a little bit tricky because there's not harmonized cross-country data on measures of polarization, but many countries do various versions of election studies, surveys and things that include measures like, how do you feel about people from the other party? How do you feel about the other party? From which you can construct measures, at least of this effective kind of polarization, and see how that has trended over countries.

So this is that trend for the US. So it's kind of similar to the Pew data that I showed you, but now it's built up from some underlying survey measures. It's basically... I'm not going to give you all the details of how this is built, but it's basically the difference between how positively do people feel toward people in their own party, relative to people in the opposite party. And most of this increase is because people are feeling more and more negative about people in the opposite party. So this is the trend for the US, from 1980 to... Here we can only go up in these data to about 2016. We see this big increase. So question, do we see that same increase across lots of countries? And the answer is really no. So here we were able to do this for 12 countries, and here I'm going to sort them by how big the increase in polarization is.

So the US stands out for having the largest increase in polarization among any country in this sample, in this effective polarization measure. Some other countries have seen increases.

There's an increase for Switzerland, although the data there are a little bit sparse. France. Small increases in Denmark, Canada, New Zealand. Basically flat overall, in Japan, Australia, Britain. And decreasing in Scandinavia, and decreasing quite dramatically in Germany. A few things to keep in mind here. One, note the timeframe here. We're talking over a long period of time, not what has happened in the last few years. So if you look at the UK plot, you might wonder, "Well, what about Brexit and things like that?" If you just look at the most recent years, maybe you can see a little uptick, but over this long horizon, it's been pretty flat. Another thing to say is, these are all constructed to be measures of, how does the average person feel about the average person who is not part of their party?

So in these multi-party democracies, that gets a little bit complicated. And so, this downward trend for Germany, you might say, "How is that consistent with the rise of AFD and right-wing parties, and all this stuff we've heard about?" What's happened in Germany is, the two large mainstream parties, both in terms of their ideologies, in terms of how their legislators vote, and also in terms of how people feel, have gotten closer and closer to each other. So an average person in Germany, feels a lot less polarized vis-a-vis another average person in Germany. However, you have a small minority in Germany, who are part of say, an extreme right-wing party, who people really don't like. And those people really don't like everybody else. This is not to say that you don't have that kind of very different phenomenon of the rise of extreme parties. But in terms of overall polarization, you don't see the same kind of increases you do in the US.

Okay. So that's second descriptive fact that again, is not dispositive, but I think is different than what you would expect in a world where this is mainly about social media. Now, a little more recent evidence. All of that is descriptive. What do we know about causal impacts? One of the things, as I said, I think that's exciting, is the ability in this domain, to really do randomized experiments, isolate causal effects, really see, at least in a short run sense where we can do that, how does being on social media, not being on social media, or the nature of what you're seeing on social media, impact polarization?

And so, there's been a bunch of recent work on this, which I think is neat. Let me tell you about... There's this project that the first batch of papers from were just published last month, in science. So this is something I'm part of, but it's a big group of academics. It's a collaboration with Meta, that started with a goal of analyzing the impact of Facebook and Instagram on the 2020 election. And so, as part of this project, we were able to do a bunch of randomized experiments in that run up to the election, to see what impact it had. There are a bunch of interesting things to talk about, about this project, in terms of how that collaboration was set up. How did we try to guarantee the independence of this research? I'd love to talk about that more at the Q&A, if people are interested.

Let me not go into the details now, but just say, this big group of academics who are working on this project and had pretty free rein to look at data from inside Meta, to design experiments the way that we wanted to design them, to look at these impacts. Now there are a number of different studies which are part of this project, but I want to focus on the ones that are randomized experiments. And those... Two main things I want to tell you about. So the first is, what I would call a deactivation experiment, which is similar to some prior work that I had done with co-authors. Where basically, we're going to randomize people into a group where they're not on Facebook or not on Instagram, for the six months prior to the 2020 election.

So think about all this stuff that was going on in the run-up to 2020, and all this stuff about mail-in ballots, and Trump saying things about the integrity of the election, all of the content that was swirling around. What happens if somebody is not on Facebook during that time? Not exposed to any of that? Do they feel differently about issues? Do they feel differently about the integrity of the election? About whether the election is conducted fairly? How is that different? Here we're lucky to have a pretty big sample. So this is 35,000 people that we could randomize

into these treatment and control groups, looking at Facebook. And also, in a kind of separate experiment, looking at Instagram.

The other set of experiments I want to talk about are what I would group together as platform experiments. Which were basically a set of experiments that varied things about people's newsfeeds. So if you enrolled in this study, you agreed to complete some surveys and then you knew that various aspects of your Facebook or Instagram experience might change, but you don't really know how or in what ways it's going to change. In fact, there were four main treatments and a control group. So those treatments were, first, a group that was randomized into having just a purely chronological newsfeed on Facebook and Instagram. So basically no algorithm at all. You just see things in the order in which they're posted. Could think of that as some kind of omnibus measure of what is the impact of the algorithm overall, on political outcomes and polarization.

Two, a group that saw no re-shared content. We pulled all re-shared content out of their feed. So if you think of re-shares as a key vector for viral content and the kind of things that are promoted through that virality channel, this experiment is going to shut that down. Three, there's a kind of breaking your echo chambers treatment, where we dramatically down ranked like-minded content. That is content that comes from users or pages, or groups, that look like they have the same ideology that you do. So down rank that content and therefore, replace it with content from users who either show up as somewhere in the middle or show up on the opposite side of the spectrum. And finally, there's a treatment where we just remove political ads altogether, and a subset of that where we just remove micro targeted political ads.

So that's a set of... There are also some descriptive studies and other things that are part of this project, but these are the key ones. I can only talk... For reasons that are related to the ground rules of this project, I can only talk about the results of the ones that are published, which are the chrono feed, re-shares and like-minded experiments, not the political ad experiment. But we will have all of these results hopefully published soon. Each of those is a paper. Each of those has a whole bunch of different interesting results and nuances, and things we should talk about. We don't have time for that. So let me just give you a really broad overview. Which is, these changes had big impacts on what people experienced. So the deactivation experiment obviously had a big impact on what people experienced, because they weren't on Facebook or Instagram at all.

But the platform experiments also changed quite dramatically, both what people saw in their newsfeed, the like-minded treatment reduced by about a third, the amount of content people saw from like-minded sources. Also, those had big impacts on the time that users spent. So you think about something like the chrono feed treatment, get rid of the algorithm, just show people essentially random content. You would hope that that would reduce engagement. If not, there's a bunch of people at Meta who probably should lose their jobs, because they just spent 10 years trying to optimize an algorithm for this target, and if it doesn't reduce engagement, they might be in trouble. It does reduce engagement quite dramatically, dramatically reduce time people spent on the platform. Despite that, despite the fact that people's experience changed, all of the published results show very precise zero effects on polarization, on beliefs about election integrity, on favorability toward candidates, on whether people voted, participation.

These are big samples. So that means the index of polarization in the control group and the index of polarization in the treatment group, are very, very close to each other. So this does not settle the issue. This is one of many streams of research. There are a whole bunch of caveats that, again, I'm sure you all can give me. These experiments are about short-term impacts. We can only do, say the deactivation, we can only do these things over a duration of six weeks. Maybe if they ran for twelve weeks or two years, or ten years, they could be different. This certainly doesn't tell you what would happen if Facebook never existed, but I think they're

informative nevertheless. All of these things are measuring average effects. That doesn't rule out the possibility that there are small groups of users who are profoundly affected in ways that we should care about.

This is a big election. People are saturated with information in the run-up to a presidential election. That could make these effects small. It could be that in less saturated environments, where people don't know as much, the effects could be bigger. But notwithstanding all those caveats, I think the body of evidence as a whole, has pushed me toward being less and less... Believing it less and less likely that social media has been the primary, or a key central driver of the increase in polarization over time, as well as a key driver in the short run sense of polarization right now.

So we can talk about all that more. What about internalities? Impacts on wellbeing and mental health. One thing to say here is, this is a place where I think the evidence we have is much thinner. There's still... Particularly about impacts on kids, there's just still very little causal evidence, and we need more. Anybody who's interested in supporting, assisting with, helping us do more large scale RCTs to really understand those impacts on kids, please talk to me. Because we need to do this, we're working on it. I think what we have, to begin, is a large body of what I would call, circumstantial evidence, that something important might be going on. You can see that with adults, just in the sense that there's an overwhelming sense, in terms of surveys and self-reports, that people are not very happy with the way they're using social media. Not everybody, there's about 40% of people who say, "It's great. I wouldn't change a thing." But about 60% of people say they wish they used it less.

If you look at standard kind of survey measures that are used in medical or psychology literatures to assess addiction, you can see many people look kind of addicted by those measures. And maybe more importantly... Or I think, definitely to me, more importantly, on the side of youth, adolescents, kids, there is really, genuinely a crisis of mental health, which happens to coincide very precisely with the timing of the rise of use of social media. That's not causal, that's not dispositive, it's circumstantial. But it's the kind of circumstantial evidence that I think should make all of us pay attention. I put Jean Twenge's book, *iGen*, up here, which documents this evidence in great detail. I think we could debate the kind of inferences in that book and how far she's willing to go, and what she draws from that, but I think, the data at least tell you there's something to be worried about.

So what do we know from randomized experiments? I think basically, we have large scale, randomized experiments on mental health impacts, on 18+ adult populations. Let us see what's happening with adults. And in some cases, are big enough to let us zoom in on 18 to 24, college age kids, but not look at kids under 18. And then there's a small number of very small RCTs with younger populations, that I think are maybe harder to interpret, but at least suggestive. One of those is a deactivation study that we did, that was published in 2020. That again, was just like, take a treatment group of people, pay them to deactivate their Facebook accounts, not be on Facebook for a month. See what does a month long detox from Facebook do to your wellbeing. Those people who were not on Facebook saw substantial improvements in their self-reported wellbeing, were less likely to self-report that they're depressed, that they're anxious and so on.

Not going to tell you exactly what that figure is, but just to give you the flavor. Those are treatment effects across a set of mental health outcomes from that experiment. Other studies have shown similar things. One way to gauge the magnitude is to compare this to, what do we know about the RCT effect of putting somebody in therapy? Like CBT or something, for three weeks or a month. And these effects are reasonably large relative to that. You also saw that people, if they spend a month off Facebook, themselves said they wanted to use it less in the future, and in fact, used it about 20% less in the future. There's then been a number of other experiments since then. Still small set. There's a study that we did focused on addiction, which I

don't have time to go into the details of, but I think really try to zero in on what are the actual behavioral mechanisms associated with addiction, like habit formation and also self-control problems, and temptation.

And I think shows pretty convincing evidence that there is an addictive aspect to those technologies. In the 2020 election study, it's about the election, but as part of that, we do have one study that looks at emotional state and wellbeing outcomes. So I can't tell you those results yet, but they are going to allow us to look at impacts on mental health, particularly for that kind of 18 to 24 group. Importantly here, we have Instagram as well as Facebook. So there's been no large scale RCT study of the impact of Instagram. We're going to be able to do that. We're going to be able to look at gender, how different is it for boys and girls? Which I think has been really important. And then as I said, there's a big literature here, including a number of smaller scale, randomized experiments, that are at least consistent with the hypothesis that there are negative effects on kids.

There's also, last thing I'll talk about in terms of evidence here, a really nice study by Luca Braghieri, Ro'ee Levy and Alexey Makarin, which falls into the kind of natural experiment category. So not randomizing things, but going and finding variation in the data. They look at the original rollout of Facebook across colleges, and in ways that I don't have time to tell you about, they build a very convincing case that there's quasi random variation. If you look at basically somebody who was in a cohort in college, where they had Facebook in their last year of college, compared to somebody who graduated just before and just missed having Facebook in college... There are surveys and data that actually have tracked mental health outcomes for graduating seniors for a very long time. So they can look at those effects and find quite substantial, negative effects on mental health of having Facebook in college.

Again, just to give you the impression, this is the set of outcomes that they look at, which includes actually, because this is a really high quality survey, stuff like, not just self-reports, but also contact with the health systems at their universities, diagnoses, prescriptions for medication, stuff like that. And these effects I think, if we talked about the magnitude, are maybe small to medium. But they also show that the effects are quite a lot larger, if you focus on people who at baseline, are at risk. So take kids who were... Or college students, I shouldn't probably call them kids, who have baseline mental health profiles that suggest they might be at risk of negative effects. If they then have Facebook, have much more substantial negative effects. Okay. So let me stop there. To conclude, I want to talk, just for a second, about policy, and maybe we can talk about this more.

So as I said, overall, kind of big picture takeaway for me on polarization, is that it makes me less worried about the impacts of social media and not about polarization in general. Polarization is also a domain where the effectiveness of the policy levers that we have, has frankly not been very impressive. And so, this leaves me... I think more work in this domain, people thinking hard about how to use policy to improve things like misinformation, is great. We should be doing that, but I'm less optimistic about the return, or the possibility of really major impact through those policies. In contrast, on the mental health side, I think we really need aggressive policy action here, particularly vis-a-vis youth and kids. That is not because I think the evidence is definitive. I think the evidence is still very limited. We don't really know. But it's a case where we need to make decisions under uncertainty. And there is enough circumstantial and causal evidence to suggest the probability that there are large harms here is reasonably high.

And on the flip side, I have seen no evidence that has convinced me that the harms of say, not letting 12-year olds have Instagram accounts, in the scenario where we were making a mistake and actually Instagram was fine, and there were no... But we screwed up and didn't let 12-year olds use Instagram, that there's any probability of corresponding harms that would be of similar magnitude. So in terms of the decision uncertainty, we have big risks on one side, no substantial

risks that I'm all that worried about on the other side. And so, I think that there's a real case here for being bold in policy responses, and we can talk more about what those might be. Okay.

Matthew Gentzkow:

I think we have time for a few questions.

Laura Martin:

Thank you. Laura Martin with Needham. So really impressive, interesting work, and congratulations on your Nobel Prize early. I want to be the first to congratulate. My question to you is this, MySpace started in 2003, Facebook started in 2004. If you were sitting as a policymaker in 2005, knowing it does a lot of good, my question is, what legislative remedies would you have put in place to mitigate some of the harms, while keeping the good? Is there a regulatory framework that actually would've prevented some of these harms from happening, back in 2005, before they ran away with what they did?

Matthew Gentzkow:

Yeah. It's a great question and a long conversation. And I think obviously, the answer is going to be different across these different categories of harms. What would we have done if we had known ahead of time? Foresaw the political impacts and thinking about-

Laura Martin:

That's the question, what's your answer?

Matthew Gentzkow:

Content and misinformation. I'm going to answer the question. I think there are many different answers, depending on which one we're talking about. Let me talk about the wellbeing side of this. Which is, we have a long history of regulating both access and the type of content people can see, by age, in this country and lots of other countries. That is, in some sense, a solved problem. There are of course, lots of details about, how do you do age verification? How do you make that work in these technological contexts? But I think, legislation that applied a similar framework to the regulation of under 18 children and youth access to, and use of these technologies, to what existed for television, to what existed for a number of other things, I think broadly speaking, would've been something I would've supported and would've potentially been effective.

Peter Davidson:

Yeah. Peter Davidson from Mental Side. Thank you so much, first of all, for the research that you've done, and also your kind of bold call for action on the internal side. I'm a pre-market person, so I don't know how to operationalize that. The last question was in that direction. So it's a great topic for the next panel and for this discussion. So thank you for the work there. First question is really more... My question is more on a methodology on the externality side. So when you isolate people for 30 days, or whatever it is, and you take them off Facebook, and you find that their polarization's not that different or whatever, I'm just wondering, to what extent do you take into account the echo chamber that's around them?

So 99%... Well, all of their friends, other than that one person, are not taken off. They're all saying, "Hey, did you see what Donald Trump said today? Or did you see what this person responded to what Donald Trump said today?" How do you level that out? And once you do that, are you still more optimistic that that's not that big a problem on the externality side?

Matthew Gentzkow:

Yeah. Yeah. Great question. And I think the answer is, we don't level it out. So absolutely, when people are not on Facebook, they are in an environment where number one, they have access to lots of other media, including cable television. Which I think we didn't talk about here. Have a ton of evidence it actually has very large impacts on polarization. Two, they have social networks that are full of people that they're hearing these things from. Three, we are in a situation where the behavior and rhetoric of leaders in this country, I think has a profound impact on polarization, and people are exposed to all of that still. They're hearing what they're hearing from presidents, from presidential candidates, from congresspeople and senators, and so on, and so on, and so on. So all of those forces are super important. We're not leveling them out. We're saying, "Holding that background constant, does whether you're on social media or not matter?"

And I think a fair read of that is, in that environment, the marginal impact of social media is not that large. And those other things are probably playing a really important role. Now, one really important caveat that's sort of on my list is, social media could have indirect effects through all of those things. Cable TV could be different in a world with social media than in a world without. The behavior of Marjorie Taylor Green and all kinds of politicians could be different in a world with social media than without. So there are other channels there, and I think that's an important caveat for what we're doing.

Scott Wallsten:

This will be the last question.

Jamie Susskind:

Hey. Hi. Hi. Jamie Susskind, I work for Senator Blackburn. I don't know if you can hear me. Thank you, first of all, for just talking about this for wellbeing. I put it on the Sli.do. I was curious. So arguably kids don't use Facebook that much, or at least that's sort of what the findings have shown, and sort of their own findings show that internally. So I was curious why you all focused on Facebook instead of TikTok, and as you alluded to, Instagram and some of the other sites that kids are frequenting more.

Matthew Gentzkow:

Great. So to be clear, in the 2020 study that I was talking about here, we focus on Facebook and Instagram. We focus on Facebook in that study because it's mainly about political impacts, it's not mainly about kids. And if we think about political content and channels for political impact in the US, in the adult population, Facebook remains I think, the most important platform. For kids, it is clearly not. And so, as I said, I think one of the exciting things there is that we actually do have this kind of randomized variation for Instagram, which is not the only thing. I think doing that for TikTok as well would be something we'd love to do, and really important and interesting. But I think Instagram, for that 18 to 24 population, is really important. And at least that's a first step, if we can get that number nailed down. But we would love... As I said, we're doing ongoing work in this space. And I think... I mean, I have a 13-year-old and a 10-year-old, so it is vivid to me, what's important and what's not. And so, Instagram, TikTok, Snapchat for sure.

Audience Question:

I just wanted to ask, have there been studies on the persistence of political viewpoints? Because another issue is, if I am at the extreme, you could put me in a closet for two months and I'd still come out extreme.

Matthew Gentzkow:

Yeah. We can consider that on the list of policy remedies. Yeah, so absolutely. So there is a big body of literature on whose views can move. Broadly speaking, there are two margins here. For the more extreme people, the really important margin is like, do they turn out and vote, or not turn out and vote? Do they come out of their closet on election day? But there are also a set of people whose views are more malleable potentially. And in these studies, we can zoom in on moderates, swing voters, see is there any effect there. And the zeros are basically similar in that case.

Scott Wallsten:

Okay. Great. We do have to move on. So thank you very much. Really great, really great talk.