FTC PrivacyCon 2018
February 28, 2018
Segment 3
Transcript

YAN LAU: Hi. Welcome to section number three, Economics, Markets, and Experiments. My name is Yan Lau, I'm one of the economists in the Bureau of Economics here at the FTC. So I guess we will just kind of start with the presentations, and afterwards, we'll have a panel for questions.

And as a reminder, remember, there are those cards that are being handed around, so if you have a question, feel free to jot it down and then pass it back to one of the FTC people and we'll get it to me and I'll ask them.

All right. So without further ado, this is Ying Lei Toh from Toulouse.

YING LEI TOH: Toulouse.

YAN LAU: Toulouse, yes.

YING LEI TOH: All right, so thank you for being here. And today I'll be presenting a paper of mine which is about the firm's incentives to invest in data security, and particularly looking at the role of reputation.

So we know with rapid digitization, the number of data breaches have also risen rapidly. Just in 2016 alone, there have been 1,800 data breaches, which resulted in the compromise of approximately 1.4 billion records. And this number of-- this prevalence of data breaches seems to be indicative that firms are not investing sufficiently in data security, and there could be two market failures which contribute to this lack of investment incentive.

First, we have input that consumer information-- so consumers that typically aren't aware of a firm's data security practices, as well as the occurrences of the data breaches. And secondly, we have investment externalities where firms, when we make their investment decision, they don't take into account the losses that are incurred by consumers and other third parties in the event of the data breach.

And so given that data breaches can actually result in adverse consequences with victims, including things like identity theft and payment fraud, one important question that we would ask here is whether or not regulatory intervention is necessary. Or is it possible for us to rely on the market to provide firms with incentives to invest in security?

Now in theory, there is a mechanism through which the market could provide firms with incentive to invest, and that is by reputation consent. So if you think about data breaches as the one at Target, when this data breach becomes publicly known or consumers become aware of these data breaches, the firm will suffer damage to its reputation. And as a result of this reputation damage, some consumers may choose not to purchase from the firm in a future period.

And it's precisely this potential loss of future business that could provide incentives for firms to invest in order to lower the probability that a data breach occurs.

Now in practice, the effectiveness of this mechanism is going to depend on two things. So first, whether consumers are willing to punish a firm for the data breach by voting with their feet, and that is going to depend on how high the losses they expect to incur as compared to how much the value the product or the services offered by the firm.

And second is going to depend on whether the consumer is able to punish the firm, and this is related to whether the consumer actually is aware of the data breach, because the consumer cannot punish the firm because the consumer does not know that his or her identity has been breached.

And in practice, there are several factors that could affect the consumer's willingness or ability to punish a firm for data breaches. And in the context of payment card detail breaches, what I focus on in my paper, factors such as high liability protection, we know with many credit card companies, if offer consumers your liability policy, where consumers are not liable at all for any fraudulent charges made on their card; and also, your bank do actively engage in fraud mitigation that also lowers the likelihood that a data breach would actually result in losses, both of these factors are going to limit the amount of losses that consumers face in the event of a data breach and make them less willing to punish the firm for data breaches.

And on top of that, it's also very difficult for consumers to actually become aware that their data has been breached. This low likelihood of breach detection make them unable to punish the firm, even if they were willing to do so.

So what this implies is that although in theory, reputation concerns could provide like a market-based solution to addressing this under-investment in data security, in practice, it is unlikely to be very effective, and therefore, we may actually need to have a more direct regulatory intervention.

So I considered various forms of regulation interventions in my paper. So I think about those interventions as, you know, thinking two possible-- in two possible categories. So first, the regulators could try to intervene indirectly-- that is, by relying on this reputation mechanism. So they could impose policies that increase the consumer's willingness or ability to punish the firms for data breaches. Or they could directly correct for the underlying market failures of imperfect information and externalities.

And I consider a few of these indirect approaches in my paper, including expelling a breached merchant from a payment card network, having the banks stay active through the monitoring of fraudulent activities, and as well as mandatory breach notification. And all of these policies are seeking to improve either the consumer's ability or willingness to punish the firm.

And I show that basically, any policy that increases the consumer's willingness or ability to punish the firm, as you see in this table above, it's going to lead to more investment in data security.

However, it may not result in higher consumer surplus. In particular, this is the case when you have a policy that tries to increase the-- or when the policy is trying to increase the consumer's willingness to punish the firm, because we know that the consumer is going to be more willing to punish the firm if they incur high or more losses in the event of a data breach.

So while we are able to increase the amount of investment made by the firm, we are doing so by putting a high cost on the consumer. So overall, it's not clear that the consumer is going to be better off under such a policy.

One of the positives that I would like to spend a bit more time talking about here is actually mandatory breach notification. For those of you who are familiar with this policy, it's often said that mandatory breach notification brings about two benefits to the consumer.

Now first, we have improved transparency, because, you know, with breach notification, consumers are more aware of data breaches, and as such, they'll be more able to punish the firm for a data breach. And the second benefit is loss mitigation. When consumers are aware that a breach has occurred, they could take actions in order to reduce the losses that they may incur, such as applying for fraud monitoring or closing their own user account and so on.

And although this is actually a benefit to the consumer, it may have an adverse impact on investment incentive, because if you recall from what I said earlier, the consumer is more willing to punish the firm when they have more to lose in the event of a data breach. Well, loss mitigation is basically doing the reverse, it's going to reduce the consumers' willingness to punish.

So we have this policy that on the one hand, you increase the consumer's ability to punish, but on the other hand, reduces their willingness to do so. And as a result, we may have an ambiguous impact-- to possibly have an ambiguous impact on investment level and consumer surplus.

And in terms of direct intervention, I consider two different forms of regulation. So in order to address the problem of imperfect information. I look at a regime that mandates the firm to review either their money is invested in data security or its state of security, and this could be implemented in practice by mandatory certification.

And to address the problem of externalities, I consider a liability move that would shift the losses that are incurred by consumers and third parties to the firm.

And both of these policies do not affect the consumer's willingness or ability to punish, but directly address the underlying market failures, and therefore without an increase in investment level as well as consumer surplus.

So to wrap up, what I show is that from a consumer protection viewpoint, a direct type s intervention which addresses the underlying market failures are actually preferable to indirect interventions which rely on consumers to punish the firm for data breaches. And in terms of policies, it is always good to have policy to improve information. So for example, either about the security level of the firm or all about data breaches.

And finally, I think like one of the most important points I would like to make with my paper is that there may be this potential trade-off between protecting consumers exposed from the losses from data breaches and increasing or fostering investment incentives [INAUDIBLE]. So for a regulator who wants to achieve both objectives, it may be necessary for us to have actually two different policy instruments.

So this is all that I have. Thank you.

YAN LAU: Thank you.

[APPLAUSE]

Next we have Sasha Romanosky from RAND.

SASHA ROMANOSKY: Well this is fun. I'm Sasha, I'm trying to figure out what I can say to you in 10 minutes. Maybe I can try and answer all of your questions about cyber insurance. I probably won't get to all of it, but I'll try. This is very weird.

I want to acknowledge my co-authors, Lily, Therese, and Andreas, you did a great job of helping collect and code these policies. What I'll walk you through is the exercise of-- well, or the insights that we collected from acquiring these policies and mining them to really understand how insurance companies price cyber risk and privacy risk.

And the reason that's important-- I think-- is because people tend to turn to insurance companies when they want to understand how to price risk generally. And in this crazy world of trying to understand privacy harms and trying to quantify privacy harms, and even try to understand the taxonomy of privacy harms, we think-- many people think that insurance companies have the answers. And so I wanted to explore that-- really, what do they know, how do they price these risks, how do they categorize them, and what is the process they use?

I mean, this is the motivations slide. You probably don't need to hear all the background of this, but the short story is that there are lots of interventions that we can think of, we can imagine, that we can try and implement if we were a policy advisor to incentivize firms to invest in appropriate security and privacy controls.

Now, maybe they are already investing in their cost minimizing level. Maybe also we want them, some people want them to invest in higher levels because we feel that they're underinvesting. But the point is, how can we do that? What are different kinds of interventions we can apply.

Cyber insurance is one of these that come up. Ex-post litigation, or at least a privacy right of harm-- sorry, a cause of action has been argued by the Department of Commerce, and in other work, we talk about how those lawsuits end up. But cyber insurance, this market-driven incentive created by firms and insurers to help regulate these practices is certainly on people's minds.

The challenge, of course, with cyber insurance is familiar moral hazard, the adverse selection. There have been solutions to this in different kinds of markets. I'm not here to really talk about those, but only to highlight that there are pros and cons of this insurance market overall.

So really, what we try to understand, what is the current state of cyber insurance policies? And to be clear, for those who aren't super familiar with them, these are corporate policies that companies take out in order to recover costs from data breaches and privacy breaches. So of all the costs you've heard about today-- the cost of notification, the costs of forensic investigations, PR and marketing, in some cases, business interruption, ransomware, all of those kinds of losses. And then to understand how carriers price these different risks.

To give you just a sense of the market, $2 to $3 billion in US premiums is where we're at now, expected to be quite a bit higher in just a few years. To be clear, though, this represents a very small percentage of overall corporate insurance. If we're talking directors and officers, error and omissions, professional liability insurance of corporations, they represent-- all of this cyber insurance-sy stuff represents a very small percentage of that-- but still in the billions of dollars, which is not nothing.

You see some estimates here of typical kinds of premiums, typical kinds of limits, and towers. So a tower of insurance, you typically don't buy a policy with a limit of $500 million, you might go to AIG first to get $50 million or $75 million, you might go to Lloyd's to get another 50 and you build this tower, literally a tower of insurance-- it's priced differently based on that.

To give you some context, Target had a tower of $100 million, Sony had a tower about $100 million, Equifax I think had a tower of $125 million, which everyone expects them to blow through pretty quickly.

So the fun thing for us is that insurance is regulated at the state level, and because of that, we are able to go to the state insurance commissioners-- to the websites-- and acquire these policies and download the policies. We're not doing anything tricky about these policies, anyone can go in to look at them. And whether these are home or auto or health or whatever policies, we collected cyber insurance policies.

We didn't have infinite resources, as much as we would like to have, so we collected policies-- we focused on some of the largest states, because that's where we figured we would have the most variation in policies, and those who would target, say, the largest firms. And this gives you some sense of the numbers of policies and the different components of the policy.

So if we talk about-- if we think about a holistic insurance policy-- at least in this context-- made up of three components. The coverages, like that which is covered, that which is excluded; the security questionnaires, the questionnaires that are provided to the firms ostensibly to help the insurance carrier understand and assess risk and differentiate the risk across firms; and the rate schedules, and the rate schedules are the algorithms essentially that determine the pricing.

So as I mentioned, what is being covered? Operational loss, so if you have an attack against your website that causes you to lose business for three days, some portion of that lost business can be

covered. Forensic analysis, you need to understand who is it who attacked you, you need to restore your systems, get everything up and running. Breach notification, all the PR stuff that I mentioned.

In some cases-- and certainly there are exclusions. So anyone who has purchased insurance policy, it's very important to understand what is being covered, but it's also equally important to understand what is being excluded. Very often these acts of war are not being covered. If Russia cyber bombs a particular firm, you know, the insurance company may not cover all those losses. Sometimes they will, but not always. Other exceptions relate to ransomware. As I mentioned before, it could be covered, but it could also be excluded.

The questionnaires, as we looked at it, generally relate to these four categories and they ask questions, all sorts of questions related to what industry are you in, how many public-facing devices do you have, more technical stuff-- what of security controls do you have? Firewalls and intrusion detection, two-factor authentication.

What sort of information policies do you have? Do you have a website policy? Do you have a user acceptance policy or a user awareness policy or security training for your employees? Third party policies if you're doing business with other parties-- all the stuff that you would expect you would want to know if you were an auditor, if you were an assessor of a company and you want to understand what their security posture was, what their security maturity was.

So, the really fun thing is that in these policies we collected, sometimes there is justification by the firm, by the carrier as to how they derived these policies-- how did they come up with the pricing? And so of course, they need to provide this to the auditor so the auditors can-- the insurance auditors can make sure that the prices aren't unfair or disadvantaging anyone, and so we see these explanations.

And this helps us understand how carriers price cyber risk. Privacy risk, cyber risk, they're all part of these policies. And so some might argue suboptimally, so they would say, limitations of available data constrain the traditional actuarial methods used to support the rates. So I kind of interpret that as, we don't really know. We as a carrier, we don't know how to assess this risk, we really have no idea, we're going to make it up.

The base retentions, these limits in deductables are what we believe to be an appropriate level. Well, we're kind of guessing. Again, we don't really know. We sort of have some sense because we're in this business, presumably, but we don't have proper information.

The rates for the above-mentioned coverage have been developed by analyzing rates of other competitors. So this is kind of curious. So we don't know, we're using someone else's guess. We don't have the information, we kind of hope they do, so we're going to take their rates. Now actually, and maybe to their fairness, this is not uncommon inside the insurance business-- often competitors use each other's prices, that's probably fine.

It is true other circumstances, other firms have a little bit more sophistication, a little more justification. Loss trend is determined by examining 10 years, fiduciary frequencies vary. So

we're looking at other books of business, right? We don't really understand the cyber, but the distribution of loss kind of looks like this other field, so we're going to incorporate that. Taken from miscellaneous product liability. Look, we're looking at errors and omissions, so there's some evolution here, at least there's some variation.

So the answer to the question of how they price cyber risk. What we found was generally three categories, first is flat rate. So a single premium offered to policyholders generally targeted towards these small companies. So the premiums here are a couple of hundred dollars a year. Obviously there's no variation in risk across different kinds of companies.

The second approach, we'll cost it the base rate. So generally, they use a lookup table of, say, asset size. If your assets fall between $2.5 billion and $5 billion, your base rate will be $16,500-- meaning your base premium will be $16,500-- oh my goodness. And then there are a series of factors. It's a linear product of different kinds of factors that the insurance company will use to price this premium.

If you want higher limits, your premium will go up. If you want lower retention, your price will go down. Depending on what kind of industry you're on, they've developed these hazard tables, so they price a premium appropriately and you can imagine.

And then the other kind of evolution, the third evolution is to look at-- and here's now we're incorporating answers from the security questionnaires-- do you have this kind of policy, do you have that kind of policy, do you use this technology, do you have that kind of technology? And you can see it going on and on and on.

But the-- you know, what we're doing with this is to expose in a sense to add some transparency to the pricing. And I guess I'm here to say that the equations, the premiums are really nothing more complicated than this. They take this base factor of the firm-- base properties of the firm and incorporate these different properties, these different characteristics of the firm often in terms of, say, a high, medium, low, so there's still a lot of guessing that's involved, but really amounts to nothing more than that.

So I'm over my time, thank you very much.

[APPLAUSE]

YAN LAU: Next we have Jaspreet Bhatia from Carnegie Mellon.

JASPREET BHATIA: Hello, everyone. I'm just Jaspreet. I'm a PhD student in the School of Computer Science at Carnegie Mellon. Today I'll talk about a framework that we have developed to measure privacy risk. This is a joint work with my advisor, Travis Breaux.

So let's first start with a summary of some of the interesting results we've had so far. So the first one was-- so we've developed a risk likelihood scale that I will show in a few slides. And what we found from that scale was that more likely a harm or a privacy violation is, the less willing people to share their information.

And we use the scale-- willingness of a person to share their information as a dependent variable-- that is, we use it to measure the privacy risk. And the idea is that if a person is more willing to share their information, that means that they're more willing to accept the privacy risk or any other kind of risk that they see in the situation.

So the second big observation we had was that as the societal benefit of doing something increases, so does the user's willingness to share their information for that benefit, and this was also talked about in one of the previous sessions.

And in an interesting finding, what I was not sure was that when being asked to share their information with the federal government, the participants were more willing to share information about who they are-- that is, things like their IP address, their MAC address, their device information, and so on; ask them types of things about what they actually do online. So things like their browsing history, the videos they've watched, or some of, you know, their email content, and so on.

You know, we did a lot of studies over the last one-and-a-half years to measure the effect of privacy risk on different factors, such as benefits, the type of information, with whom the information would be shared, and so on. And what we concluded was that measuring privacy risk is a very complex task. We have a lot of factors, then that each factor has different levels. So for example, sharing your SSN is very different from sharing your age range or maybe the country you live in.

So we wanted to find an easier way that is less resource-intensive, and also that we can do, you know, at scale to measure the privacy risk. And we found that measuring the discomfort associated with the sharing of an information type was one such predictor.

So let's start with what really risk is, and risk has been defined in literature as having two dimensions. One of them is uncertainty, and the other is what are the adverse consequences of an activity? And this defines risk as a function of the likelihood of the adverse consequence multiplied by the magnitude of that adverse consequence.

And there has also been a lot of work and risk in judgment and decision sciences, and there have been like two main concepts. One of them is revealed preferences, and in that paradigm, what the researchers talk about is that we can reveal the risk associated with an activity if we look at historical data.

So for example, if you want to predict what is the risk associated with, say, driving a car? Then you can look at the historical data about how many people purchase cars, and then things like how many accidents were there in a given country, for example, in that period-- in the same period of time, and then you determine what the risk really is, because people accepting that risk when they're driving their car.

This does not really hold true for technology, because technology changes very rapidly. So most of the time, you do not have historical information about an activity. So for example, we have

social media and social media did not exist 20 years back, so we cannot have-- we cannot rely on data that did not exist.

So therefore, the paradigm of expressed preferences came into existence, where Fishoff and Slovic talked about surveying people and actually asking them what they felt about a particular activity-- that is, what is the risk that they perceived.

And one important observation they had from their analysis was that when people see benefits of an activity, they automatically perceive lesser risk, and this is what we confirmed in our studies as well.

This is an outline of the framework. The survey starts with three survey questions about the participant's exposure to online activities. Example questions include like-- we ask participants, how often do you shop online? And then what does-- do you provide your-- for example, do you use social networks? And how frequently do you use them? And then do you provide your personal information when you're using social networks? Or do you just browse and look at other people's posts or information they post online?

The main part of the framework is factorial vignette surveys, which are also defined previously. And in these factorial vignettes, what we try to do is we try to measure the effect of different independent variables on our dependent variable. So what I mean here is that so for example, if you want to say that-- you want to measure how different benefits or how differentiating types of information have an effect on your willingness to share, so you represent that using a scenario which I'll show on the next slide, and then you measure its effect on your dependent variable.

And in the end, we have a demographic survey where we ask people about their gender, their education level, we ask them what their household salary is, where they live, et cetera.

This is an example of the factorial vignette surveys. So here, we have three different independent variables that we want to measure the effect of. And so the first is the data purpose. So the scenario here is that the participant is told that your information has to be shared with the federal government for-- we can show them one of the purposes. So for example, we had purposes like investigating an act of terrorism, or for, say, economic fraud and so on.

And then we insert that data purpose. And then we have the risk likelihood-- that is, how likely is that violation to actually happen? And then we have different information types, as another variable, like age range, your home address, and so on. And then you have different ratings, like, are you extremely willing, are you very willing, and so on. And this goes up to extremely unwilling.

These are our results from the risk likelihood scale. So this likelihood scale is motivated from construal level theory, which says that as the social distance increases, a person is less-- as the social distance increases, people will perceive it as more likely. So for example, if something happens to a person in your family, you would perceive it as more likely as compared as something happens to a person in your country.

And so the results were similar. What we found was that if we told people that the privacy violation affected one person in their family, then people were extremely unwilling to provide that information, and their willingness to share their information increases as the levels go from family to one person in your country.

We also found that providing benefits in these scenarios helped people perceive less risk. So the blue bars that you see are the surveys when we did not show people benefits. So these were their willingness to share the same kinds of information when they were not shown any benefit. And the red bars are the results from the surveys where we showed them benefits.

And for all different scenarios, all different information types, irrespective of what that information was, participants' willingness to share their information actually increases when they see benefits.

So these are results where we try to correlate the usefulness of often information type to help with incident responses with their willingness to share their information. So people were really willing to share that IP addresses and so on, like their device information-- that is, things about who they are. And then they were extremely unwilling to share information about things like their keylogging data, their passwords, and browsing history, and so on.

As I said in the beginning of my talk, that measuring risk is very resource intensive and time intensive, and it's always not possible to survey people for all the different variables and levels for risk, and therefore, we found three different predictors that can predict the privacy risk to-- that can predict an approximate privacy risk, and we found that discomfort was one of them and it was a very strong predictor of privacy risk.

So what it means is that if you're really uncomfortable sharing your SSN or your home address, then that is directly proportional to the risk you experience while you are being asked to share your home address.

And in the end, I would like to conclude by the impact we think this can have. First we think that it can facilitate privacy by design-- that is, it can help our software developers and teams to measure the risk and identify high-risk items that should be assigned more resources to then, you know, they can make sure that they describe those practices in their privacy policies.

Similarly, it can help regulators to identify high-risk items and make sure that at least those-- to pay more attention to those data practices and make sure that the website companies describe them accurately and clearly and not in an ambiguous way in their privacy policy.

And also, this can be used by users, because they can look at these surveys and see what is the risk associated or what other users have felt the risk associated with a given situation before making decisions about using that website or those services. And we also think that it can help with the standardization of privacy impact assessments.

And thank you for being here and being a part of my talk, thank you.

[APPLAUSE]

YAN LAU: And next we have Caleb Fuller from Grove City College.

CALEB FULLER: Good afternoon. So by way of new survey evidence, my paper empirically explores three common claims in the economics of privacy literature, and in so doing, I hope to shed light on a bigger question, which is, is the market for digital privacy a failure?

So the three common claims that I examine are, the consumers are poorly informed regarding firms' information collection practices; the consumers value privacy very highly despite having very little of it; and finally, that consumers stated distaste for information collection results primarily or entirely from harms that originate in the private sector.

And so to examine these three questions, I conduct a survey of over 1,500 respondents who are solicited by a digital panel provider, and 100% of these respondents are Google users who are above the age of 18, and the demographic characteristics of these respondents mirrored the 2010 US Census on the following dimensions-- ethnicity, gender, and religious affiliation.

On each of the three questions that I explore, respondents are queried only with respect to Google, so the results may lack external validity when applied to other firms and contacts as we've heard earlier today.

So it's routinely argued that there is significant information asymmetry between digital firms and consumers. Consumers rarely know that information is being collected or what types of information are being collected, and yet to date, there has been little empirical examination of the extent to which this information asymmetry prevails.

And so I attempt to fill this gap first by simply asking Google users if they're aware of Google's information collection practices. And when my respondents are asked if they are aware of Google's information collection behavior, 90% indicate awareness, which suggests a little information asymmetry at least with respect to the existence of the practice.

It's one thing, however, for consumers to be aware of information collection, it's another for them to know what type of information is getting collected. And so to address this, I presented them with a list of items on which Google collects and some which it does not, and ask them to select which Google does collect.

And again, consumers do well on average. Very few believe that Google collects things it does not. For example, only 6% believe Google keeps a record of driver's licenses, only 7% believe Google somehow gains access to a browser social security number. By contrast, 88% are aware that Google keeps a record of all searches that are made in the browser. 75% know that Google detects a browser's physical location, and so on.

Lastly-- there's something wrong with the slide there. Lastly, when I separate these consumers based on the frequency of their Google use, OK, which is one time a day, a few times a day, or dozens of times a day, I find that the frequent users are the most informed, suggesting that those

who stand to lose the most from potential privacy harms are also the ones who are most informed about those harms.

Secondly, previous surveys of consumers routinely indicate that they would prefer more privacy than they currently experience in digital contexts, and some conclude from this finding that the digital market fails to provide an adequate level of privacy protection. But as this FTC report from a few years ago notes, we know very little about what consumers would be willing to sacrifice to get more privacy in digital context.

So to explore this question, I test whether there is a divergence between the amount of privacy people demand when it's offered to them as free and the amount they demand when they're asked to pay for privacy. And so my first finding on this question is consistent with the broader literature, and what I find is that over 70% of respondents indicate that they would like more privacy, all else held constant when they're using Google-- let's refer to this as the unconstrained demand for privacy.

Now how many of these who would like privacy in the abstract say they'd be willing to pay for privacy, however? A shockingly low percentage. Only 15% of Google users say they'd be willing to pay to use the service. Let's refer to this as the constrained demand for privacy. And so this suggests a significant divergence between constrained and unconstrained privacy demands.

Now the reason this is significant is because previous researchers have identified a privacy paradox in which consumers frequently indicate that they value privacy highly, but they just as frequently forego low-cost methods of protecting that privacy. So because of the privacy paradox, some have argued that consumers' verbalized privacy evaluations represent their true privacy preferences, which they then deviate from in the context of digital environments.

My results suggest a different angle on the privacy paradox. There may, in fact, not be a paradox at all. There may simply be a positive demand for more of an economic good-- in this case, privacy-- everything else held constant.

And so previous literature showed that there was a divergence between consumers' stated preferences and their behavior. My results indicate that we get divergence even between stated preferences themselves, when one set of questions asks consumers about their unconstrained demand, while the other set of questions asked consumers about how much privacy they'd want if they were forced to bear some opportunity cost of additional privacy acquisition.

I think this latter approach seems more appropriate since individuals must incur an opportunity cost to acquire mole of any scarce good, privacy being not an exception to that.

We'd probably get similar results, for instance, if we first ask consumers about their unconstrained demand for more ice cream-- i.e., would you like more ice cream?-- and their constrained demand for ice cream-- i.e., how much would you be willing to pay?

So this finding has relevance in the digital context, because if firms don't earn revenue through information collection, they're likely looking for some other revenue source, like charging a fee.

And since consumer willingness to pay overall appears fairly low, this raises questions about the profitability of digital firms should they attempt to earn revenue through charging a fee or even by switching to the less successful untargeted advertising.

Which brings me to the next point-- when we look just at these 15% of Google users that are most privacy sensitive, how much are they willing to pay to use Google in return for a guarantee to their privacy? Again, not much. Now this part of the survey asked consumers who had indicated a positive willingness to pay to enter a dollar figure for how much they'd be willing to pay to use Google in exchange for complete privacy protection annually.

And so what I find is the average willingness to pay is between $76 and $77. Even this figure is skewed by a few outliers, however, and so a better measure is to take a look at the median. And there, we find that the median willingness to pay is $20 per year, which comes out to five-and-a-half cents per day, again, suggesting a very low willingness to pay for privacy.

Finally, there is some speculation as to why it is that consumers tend to voice an unconstrained demand for more privacy, and an excellent 2016 paper in the Journal of Economic Literature offers a few possibilities. Maybe it's fear of price discrimination or identity theft, maybe it's just kind of the icky factor of not knowing who knows what.

And so what I do is I place these possibilities listed here by Acquisti and his co-authors in front of consumers and then add an additional option, which is the possibility that consumers are fearful of government overreach, an agency perhaps forcing a firm to relinquish information about them.

And what I find is that the results suggest broad support for the possibilities listed by Acquisti and his co-authors. And while it's not the most significant concern, it turns out that the potential for government overreach does contribute to stated consumer dislike of information collection. A full 43% of consumers indicate that that is a concern for them when it comes to collection, even non-sensitive information.

So in conclusion, I found little to suggest that the market for digital privacy, at least in this particular narrow context, is a failure. On average, consumers are highly informed and also exhibit little willingness to pay for more privacy, at least with respect to Google. Consumers, again, would likely voice a preference, say, for more ice cream should we ask them, but that hardly constitutes evidence of a failure in the ice cream market.

My findings suggest that because of low privacy valuations, most consumers won't benefit from outright bans of information. Instead, they may-- of information collection, that is. Instead, they may benefit from policy, which simply informs them of possible privacy harms and then permits them to decide whether the potential benefits outweigh potential harms from their own unique perspectives. Thank you.

[APPLAUSE]

YAN LAU: And lastly, we have Christian Catalini from MIT.

CHRISTIAN CATALINI: Thank you. Good afternoon, everyone. This is joint work with Susan Athey at Stanford and Catherine Tucker, my co-author at MIT. We're back into the privacy paradox regime, I guess, after the previous presentation. And here here's some novel data, experimental data on why we may see that in practice.

As it's been said multiple times today, people say they care about privacy and government surveillance, but then, you know, that's my own inbox, they share personal data with firms, right? On multiple points as if no one was watching. We knowingly use technologies that do not safeguard our privacy when alternatives are available, where stronger encryption and better protection of our private information is essentially just a few scrolls down an App Store chart.

So this paper is about unpacking part of the privacy paradox, and it's based on data from the MIT Digital Currency Experiment that we conducted in 2014. This is not a paper about digital currency, per se, although a big part of the problem is around blockchain and cryptocurrencies, and with all that said, of technologies really to regain part of our digital privacy.

And some of the questions that should be in the back of your mind are, you know, what are we trying to regulate? And what are we trying to protect consumers from? To give you a little bit more background in the little minutes that I have here, the experiment involved 5,000 students at MIT, the entire court in 2014 with a number of randomizations and students have to make their choices in terms of their digital wallets, their privacy, their financial transaction. Of course, we asked them a lot of things about how they think about privacy, how they think about their digital privacy in particular.

And we have three findings, which are all born of three different randomizations. The first one is what we call the kind of the small money treatment. Again, people say they care about privacy, but they're willing to relinquish private data quite easily. And it turns out that what we did was ask one of the most sensitive pieces of information that you could ask. We asked the students for the email of their closest friends.

And at least, if you believe surveys and the rankings that we give when people tell us, you kind of have to rank, what would you share, what would you not share, your closest contacts are right after your social security number. So this is one of the things that at least in surveys and at least in print, it's one of those things that, you know, consumers are really vocal about in lawsuits. Contacts are something that we guard closely because they tell a lot about our lives, about our business relationships, and everything else.

Now it turns out that in our randomization, we added an incentive. So of the sample was offered a pizza-- this is for college students, after all-- In exchange of the emails of their closest friends. And to essentially give you to the punchline, when you add the incentive to the same request, you have a drastic drop in the willingness to protect the privacy of your friends.

Now, you can throw this into regression format, you can look at all sort of permutations, how privacy-sensitive those individuals were, the 50% drop is kind of consistent across the sample no matter who you are.

And we looked really hard for a [INAUDIBLE] of effects. If there was one demographic that would not react to the incentive, but, you know, I have bad news for you, there isn't, and so we slice the data in many different ways, from dimensions-- MIT is known for having some dorms that are really vocal around open source, around privacy, it happens even there. So we didn't notice any difference across our sample.

The second treatment were slightly different and was about introducing small frictions. So the scenario we had in mind is the one that you have every time you download a new app on your phone and your time pressed, you're trying to do something, and maybe you get asked for permission for background location or, like, essentially a prompt where you're giving away your privacy in that moment in time. But your attention span is typically short in that moment and you're trying to achieve something else and you may be willing to give away a little bit more than you should.

So what we did here was ready to take inspiration from the browser ballot. So for those of you that remember the Microsoft antitrust case, Microsoft was forced to give a randomized list of browsers with different features. And so what we did, when students had to choose a digital wallet, we replicated that.

The students had different choices. In one of the treatments, we just listed the four wallets randomly. In another one we added very detailed information about how does this protect my privacy from the intermediary, essentially the startup maybe providing the service, the government, or the public. A lot of details about your cybersecurity, and kind of the trade-offs that you were making when choosing a particular solution-- and again, every student participating in the experiment had to go through this selection.

So what happens, when the wallet that would be the best fit for your state of privacy preferences wasn't first, well, of course you're way less likely to choose it. And, you know, the drop is again, substantial. It doesn't matter if it's privacy from a corporate entity or it's privacy from the government or if it's privacy from your peers, we see the results throughout.

Sadly, adding information-- so providing additional richness to that decision prompt seems to have almost no effect. I think part of that is because, again, our attention spans online and in the digital environment are extremely limited, so additional information is not necessarily better information.

And then the last treatment we did was what we call the smalltalk pretreatment. And again, you see examples of this all over the internet. Simple texts can really distract people from escaping surveillance or escaping additional data collection at large. Here's an example from a service that was supposed to help you unsubscribe from a mailing list, and in the fine print, you were giving up a lot of your private information going forward about all your emails through that service.

So we replicated something along those lines. Students had to learn about PGP-- it was a complex environment because they were trying to create a digital wallet and receive their Bitcoin. And when we described PGP in one of the treatments, we had a lot of information that,

you know, made you possibly feel more protected and your privacy preserved, but really, that text was sort of irrelevant for the privacy of your future financial transactions.

And what turns out to be happening in that regime is that people probably because they feel like, OK, now there's this technology out there that's going to protect my privacy, they become way less likely to do anything on their side to protect their financial privacy in their transactions. And so we see this again across the board, that irrelevant information is kind of swaying people away from at least their stated preferences that we collected too.

So to conclude, again, it's an idiosyncratic setting, but I think what's kind of interesting here is that we were able to test people's choices at scale more than, you know, 3,000 individuals involved and at a very fine-grained level, because wallet choice in this case had different implication, whether you're worried about surveillance by the government, you're worried about a corporate entity, or maybe your peers knowing about your financial transactions.

So again, why people say they care about privacy in the face of small incentives? It's really easy to sway them in the opposite direction, just changing a simple step in the UI. So tiny frictions on how you access a web service have meaningful impacts on the choices people make, and again, people may worry about surveillance, but simple texts can really shift them away from their own stated preferences.

So before they did all of this, we asked them, you know, what is your stand on privacy and all these different dimensions, and they became very inconsistent with what I just told you.

In closing, I think there's a broader debate both in Europe and here around the handling of personal data. I think this is about to become much worse, you know, if you want a competitive market for artificial intelligence in the data licensing and access to data that's more fine-grained than we have today unless if you want more than four players in that space.

What's interesting here is I think that if you're trying to learn about privacy choices and preferences and just using surveys, you may only discover half of the story. You actually have to observe-- at least in our case, you have observe those individuals who make choices in the environment, and also start thinking about how does that specific environment, that moment where I'm asking someone to make a choice-- in this case, it's a digital wallet, so it could have repercussions if you're starting to use that over time, over multiple years. And, you know, that split second when the students made that decision, that long-run implication, that they probably didn't think about it, just that dependency in these choices.

Now I think overall, what the results show is that and notice and choice may be fundamentally broken. You know, you may be asking for consent, but depending on how you ask, you may get the answer that you want. So we showed that privacy preferences are extremely malleable, and especially in an environment where it has limited attention, we may need something potentially like a Creative Commons for privacy where users get a really quick overview of what they're signing up for, and maybe there's consistency across services so that they can start comparing them.

So again, I'm sorry I don't have better news, but essentially, we document multiple market failures, I think, in digital privacy, and a set of open questions that I think will require additional research. Thank you very much.

[APPLAUSE]

YAN LAU: Thank you. So we'll be starting our panel questions. So again, if you have any questions, do write them down on the little sheets of papers going around and they'll be handed up.

So I guess I'll start off with this question about measuring. So measuring how consumers value privacy and data security is important for research into markets in these areas-- for example, how to quantify consumer harm, how to quantify losses to firms or mitigating threats.

So we've seen a mix of approaches today and also in the literature. So how do you think researchers can best approach this measurement issue, especially in light of the privacy paradox?

CHRISTIAN CATALINI: I guess that I can start giving that-- you named the paradox. So I think what we at least learned as researchers in this setting was that we were going into this survey with very low expectations in terms of relying on the survey data to reassess the privacy preferences. To me, this calls for a lot more research in different environments under different conditions live in the real world about the choices that people make.

My understanding, especially looking at the demographic that we were operating with, is that again, people discount the future, and some of the privacy consequences of choices that you make today may not be salient until major privacy failure into the future, and we're really bad at estimating the cost of those major breaches.

I guess, you know, sadly, recent events over time are making us more familiar with that, but there seems to be at least some sort of almost a miscalculation on part of the users on the choices that I make on a mobile phone when I allow a specific app maybe to track me for a specific purpose-- would that app keep tracking me for years?

I don't think consumers sometimes understand that those are kind of the agreements that they're getting into, and so more transparency and a better framework for communicating that to consumers I think would be quite useful.

JASPREET BHATIA: I would like to add to that. So I think I would agree to whatever Christian said, and what we found in our studies was that there was no single factor that was determining the privacy implications or that was determining the privacy risk.

So I think we need better ways of combining all the information that we draw from different studies, like be it surveys or be it like a real-time deployment of apps where we keep asking users about their preferences or what they're experiencing, and then we need statistical ways of combining all that information to determine what the privacy risk is experienced by the user.

And the other thing I would want to say is that in our surveys, we found that many times, people were not also aware of the different meanings of these different factors. So for example, if I ask somebody, what do you think MAC address is? And then, you know, do you think that it's like unique to your computer, do you think you can, like, forge it, or do you think somebody can use your MAC address?

So a lot of time, users are not aware of things, be it from the policy side or be it from the technical side. So I think use awareness is also another thing that we need to take into account.

SASHA ROMANOSKY: I would-- so the notion of measuring privacy harm has been-- right? It's been a unicorn that a bunch of us have been chasing for a while, and I don't think we've really gotten any further. The fact that we're trying so hard to find evidence of privacy harm should tell us something in itself.

I think a lot of people really want to be able to find evidence and strong evidence in lots of different ways, and I'm not suggesting that there isn't, but I don't know. I think just one thing I've learned in doing the research and in seeing the other kinds of research is that we really, really try very hard to understand, to quantify, to capture what these harms are, and it seems in so many cases, the answer is just, well, people get uneasy, people get upset, you don't like when other people have control of our information.

Again, which isn't to say that some people don't suffer legitimate harms, financial, medical, emotional, or otherwise, but I think so much of what we're finding is that it is just so hard to find them, and maybe the answer is because it's not existing in the way we think it is. And I think that's important to recognize, especially when it comes to policy interventions.

Because if it takes a certain threshold of harm regardless of how we measure it, before we want to engage and bring actions against a company or try to enforce any kind of regulation, I think if it's not very evident, if it's not crystal clear what these harms actually are, then maybe we should look somewhere else.

I don't necessarily think the problem is with statistical methods, and I don't think surveys and stated preferences are going to lead us there, right? I don't think that's the answer. Certainly surveys have a place when you really want to understand what people think of things, OK, that's great. And when there's a difference in expectations, you know, that can be a problem certainly. If firms are stating one thing and doing another, that's a problem. I think transparency can get us a long way there in whatever fashion, whatever mechanism that manifests itself.

But, again, I mean, there was a big workshop here a number of months ago. They brought a lot of experts trying to understand how do we measure these harms, and there weren't a lot of conclusions for that, and I think that's an important takeaway in itself.

YAN LAU: I guess related to that, we have panelists and also other presentations today that use sort of consumers as a source-- the users and surveys, experiments. Whereas, I guess Sasha, your paper talked about using firms, right? Which is sort of a very orthogonal source of information.

So I don't know if we can get the panels to discuss maybe this difference, and should we rely on firms to price or value privacy in data security as opposed to, say, just asking consumers, right? Those are very different sources. What are people's thoughts?

CHRISTIAN CATALINI: I'll jump in. And this is somewhat tangential to your question but related to the very last comment. I think it's important to ask, did the consumers have a choice and was that choice informed? What we see in the data is that you give me a sample and I can manipulate that sample away from whatever their privacy preferences are. And even past reveal preferences.

So to give an example, we had people coming to this survey with an open source operating system on an open source browser. So these are people that have made a set of choices that are pretty clear, and we could just get the same results but adding incentives, changing the frictions, which shows that I think maybe consumers don't have that stage of choice and so we need to take a step back. If you want a measure arm, you need to go back and be able to measure that the first stage of the choice.

And in terms of like where the data should come, if it's firms or if it's consumers, I think an important dimension is that never before was so much of our digital life connected and available. So I think we haven't seen yet one of those major failures where some of the more drastic privacy consequences would take place. There's been some isolated cases where people's personal lives have been deeply affected by digital breaches, but I think that's just the beginning just because Facebook and other services didn't exist even 10 years back, right? And so there's a lot more at stake I think for the next time.

SASHA ROMANOSKY: Yeah, I would-- sorry, let me-- so I wouldn't disagree with any of that. I think deception is bad, manipulation is bad, right? And I do think it's true that people want to understand what the rules are so they can play within the rules, right? So we should figure that out. That doesn't mean that something like price discrimination is bad, right? Different prices charged different people and for different reasons, that's not necessarily a bad thing. Student discounts, elderly discounts, right? No one's objecting to those.

And yes, it's true that we don't always have a choice, and certainly, we always want to have a choice and we always want to make informed choices, but I guess my point is just that this notion that if we just inform people, if we just give them all the information, they're going to make the best decisions for themselves, they're suddenly going to care. They're suddenly going to care about all this information that's being collected and they're going to suddenly choose DuckDuckGo.

I just think we need to get beyond that because it's not happening. If people really cared, they would take action. If people really cared, they would stop using Facebook. If people really cared, they would turn off their phones. And I just don't think-- I don't think there is a lot of solutions to be found that way, but certainly forms of manipulation, you know, and discrimination against certain classes, like that, you know, that stuff that we don't want to get into.

To the extent that FTC can help with that, I'm all for it.

CALEB FULLER: Yeah, as a follow-up to that, I think it's important, both as researchers and as policymakers, not to slip sort of subconsciously into normative evaluations.

If our research does show, say, a low privacy evaluation across the different sort of techniques that you've seen here-- and I agree with you about some of the limitations of pure survey-- I think it's important that we don't slip into sort of scolding consumers about what they should value when it comes to privacy, what sort of ends they should hold simply because in a world of scarcity, every time a consumer decides to allocate more time to learning more about these risks, they're necessarily learning less about some other risk, whether that's, you know, eating red meat or driving a car or whatever it is.

And so I think we just have to be circumspect and humble about our ability to sort of tell them what sort of trade-offs they should be making.

YAN LAU: So I have a question from an audience member, which is sort of related. So he or she says, sometimes it is difficult to convince your employees about the importance of privacy. What is your recommendation on how firms and corporations can demonstrate the value of privacy and/or its privacy office to its employees? So you mentioned, we can't really scold consumers, but presumably, I guess, we can scold employees and-- [LAUGHS]

Anyone.

SASHA ROMANOSKY: Is the question about how firms can communicate the values that they hold privacy to their employees?

YAN LAU: Yeah, I think so. I think it's the importance of privacy. You see the experiments that were run, like, you know, do we give them pizza to sort of like incentivize them to care about privacy and learn/

SASHA ROMANOSKY: One thing I might add. Publicly traded companies need to have corporate ethics policies, ethics and governance policies, which lay out many of the practices that the firm embodies. These are supposed be embodiments of the firm's corporate culture and ethics. And often, they state-- basically it's saying that we will uphold the laws, we will not be violating laws or regulations, but a lot of them really see-- those with big brands certainly seem to take-- seem to go the extra steps to really promote what they value most dearly in their culture.

So to the extent that firms really did want to, they could conceivably add a section to their ethics governance policies talking about privacy. The privacy of the firm, privacy of employee information, privacy of customer data. I haven't seen anything like that-- I'm actually engaged in a project now which is on that, which is how I know a little bit about it. And I suppose that's actually something that we could test for to see how many companies have these kinds of sections and to what extent today they describe the kinds of values that are important to them.

So I think the answer to your question is, that could be-- these policies could be one way that they promote their sentiments toward privacy.

YAN LAU: Any thoughts about that? I guess, how do you get your employees, you know, like, it seems like there are different treatments you can do to get your employees to care more about privacy.

CHRISTIAN CATALINI: I mean, that will depend by context, right? I think to some extent, there's a new wave of digital platforms that are being built with a much stronger privacy backbone. The entire activity that's taking place I think within the cryptocurrency and crypto token space, it's about regaining part of the digital ownership of your data of your interactions. So I think you're seeing new types of companies that are built with that in mind.

For existing companies, I think one interesting dimension is that privacy is also social good, right? So if I give access to my data and I'm interacting with my friends, I've also given away privacy of my friends. I think that would be another important dimensions that within a corporate environment will be important to make really salient.

YAN LAU: So I have a question actually for Ying Lei from Twitter. So this is related to sort of firms. Would increasing the ability of the class action bar to bring suits against breached companies or on behalf of consumers incentivize investments in data security? And what are your thoughts on letting more people sue, I guess?

YING LEI TOH: I mean, well, like one of the reasons why firms don't have the incentive to invest is because of this externality problem. So that the firms don't actually take to account the losses that are faced by consumers, for example. And of course, if we had-- or if class action suits were possible, if it's possible for consumers to sue firms and get compensated for their losses, then this is going to increase the cost of a data breach to a firm and that will create incentives for them to invest in security.

But I think like the issue here with at least from my understanding with consumer class action lawsuits is that this problem is standing in contrast to the point of how do you measure the harms to a consumer in the event of a data breach? Of course, with like monetary losses, like financial losses and credit card loss or fraud, those we can easily quantify, but like how do you-- what value do you attach to the psychological stress or the time that is spent on you applying for fraud mentoring and taking actions in order to prevent losses from occurring?

And I think it's important for us to actually try to quantify these costs in order to be able to allow for such litigation to be like a solution-- potential solution to the problem.

YAN LAU: So in light of your research and expertise, do you think there are market failures in regards to privacy and data security-- for example, information asymmetries or externalities? And, you know, Caleb, you have addressed this in your presentation. And so if they exist, how can such market failures be addressed by policymakers, firms, and consumers?

CALEB FULLER: Yeah, so I'll jump in. Maybe this is a less popular opinion, but I'm not sure that the extent of information asymmetry is greater than it is in other markets that we're very familiar with. And I would also say as a follow-up to that, that even if those failures did exist, it doesn't necessarily and logically follow that public policy should be brought to address them. I

mean, what if the size of the failure is very small? What if the policy has unintended consequences that can impose costs that are larger? And what if the cure is worse than the disease?

So as an example of this, I have some potential concerns about the forthcoming GDPR in Europe that's going to be unveiled in May. Among many other things, it makes the collection of non-sensitive information to be a more challenging task, and so I wonder if an unintended consequence of that is that some firms switch to charging fees for some services, and if that's the case, that means there's more transactions in the digital environment that are using credit cards, and if that's the case, it seems like there's more opportunity for identity theft.

So you can think of the cure being worse than disease in that case. OK, if it does-- the GDPR does achieve the end of sort of reducing both the lesser privacy harm while at the same time increasing the possibility for greater harms, then I think that's a real problem. And so we should just be, again, circumspect about whether other policy necessarily is the corrective.

CHRISTIAN CATALINI: I can jump in on that. I have a slightly different take as you can probably figure out. I think, you know, there is enough evidence of market failures, and especially when the externality is not internalized by one single user and there's kind of a network effect playing adversely on that. I think we need to start thinking about better technology. You know, the internet as we know it was built on a technology that was around many years ago, and now we have I think a new wave of innovation that can give you both privacy and the same services and customizations that I think people want.

I would say from a regulatory perspective, the best approach would be to give alternative business model a shot at gaining the consumer or gaining the market share and are competing for those markets. Take an environment like digital advertising. I think there's different models, and we've seen consequences of the advertisement model, even in recent elections and the effect on news.

So if there's different ways for rewarding content creators and for advertisers to pay for attention, I think there's multiple solutions. We landed on one, but I think there's more approaches to really thinking through what consumers are willing to pay for and when they're willing to give up their privacy. I think if they had more fine-grained control, we'd probably land on multiple services and consumers could choose, you know, do I want to be on a platform where it's free and monetized by generating content and by generating data, or do I want to be on a platform where, you know, I'm paying something but I'm getting some other protections back?

YAN LAU: So I have two question cards from the audience on sort of related issues, and I'll paraphrase, but what is the problem about sort of measuring the value of privacy consumer value? Is more of a lack of understanding and a lack of power rather than a lack of caring about privacy? And related to that, might consumers-- another question from the audience is, might consumers be unwilling to pay for complete privacy protection because they don't believe there is actually such a thing?

CALEB FULLER: Yes, I mean, I suppose that's possible. I think on the question of whether or not they're informed, I think we have to ask, if they're not informed, why are they not informed? And one possibility is simply that they, again, believe that the cost of becoming informed is greater than greater than the benefits. What was the first question again?

YAN LAU: Empower. Like maybe it's just all hopeless, right? There's no complete privacy protection and no one can do anything.

CALEB FULLER: Yeah. Again, I guess they don't have a clear answer on that from the research that I conducted. I mean, in the surveys that we presented them with, we basically said, you know, in an ideal world in which everything about your Google experience is identical except for all of your privacy is retained, would you prefer that world or would you be willing to pay for that world? So, again, I don't know how respondents necessarily interpreted that, but that's how the question was set up.

CHRISTIAN CATALINI: So I run through the regressions, but we do actually have a result that speaks exactly to that point. There's this flavor where some of the students in the experiment already assumed that their privacy was compromised, that that data was available anyway, and at that point, I guess you know, what's the additional value trying to protect this little piece when you're already so exposed?

And that goes back, I think, to thinking about what are the initial choices can we roll back in and rebuild in a different way?

YAN LAU: So do you think, you know, how can consumers actually be properly informed of the privacy risks and their potential choices for protecting their privacy? Because there are some-- I mean, we've heard of sort of the Creative Commons or-- sorry, the standardized form where if you standardize it across websites. Are there other ways to kind of inform users?

SASHA ROMANOSKY: So I'm very conflicted over this, because on one hand, I'm a big fan of information disclosure, information awareness by people, right? So these breach notification laws have done an awful lot, and I think one of the unintended consequences is to a spot of all industry of cyber insurance. But in principle, I am a big fan of information disclosure and the economics around that.

On the other hand, it still does seem to be the case that no matter how much we inform consumers of risks, they still don't care. And I think for a lot of people, the reason is because they just don't suffer any kinds of harms, right? I think it's pretty obvious that if they did suffer critical kinds of harm, severe kinds of harms, that would take action, right?

And so again, if really all they can articulate is that they suffer some amount of emotional distress or they don't like the loss of control of their information or that they feel a sense of uneasiness about websites tracking them and present them advertising, like I don't-- right? I mean, there needs to be something more, and they need to be-- we need to be able to articulate something more than just that in order to drive firms to adopt what we think might be better behaviors.

And I think until we can do that, right? As a society do that, I don't I don't think anything is going to change. I think this constant erosion-- we'll call it an erosion, but this movement of privacy norms is going to continue until-- right? Until there is a greater stand that individuals take.

I think one place maybe to look for outcomes, different kinds of risk consequences is in chilling effects, right? If it's the case that people are no longer engaging in socially beneficial activities because of the fear of surveillance or of tracking, either government or corporate, that's a bad thing, right? If people are no longer going to Alcoholics Anonymous, if people are no longer engaging in other kinds of socially beneficial activities because of this, of what they feel is tracking or surveillance or lack of privacy, then that's a bad thing and we should correct that.

And to the extent that there's-- any of us can conduct any kinds of research to demonstrate evidence of that, I think that would be a very valuable thing.

CHRISTIAN CATALINI: I'll disagree on that. And I think the tide has already turned in the sense that if you look at-- take the market for digital assistance. And again, I'm coming from the MIT bubble, so it's a bubble in it's own kind. The market for digital assistance benefits from a lot of privacy data, sensitive data, right? Customization, response time, and companies are taking different approaches.

So I think you're seeing products that are not maybe as competitive in DAI, but are competing on privacy. You're seeing products that you're bringing to your home, some are clearly, you know, on one spectrum of the privacy debate, and others are trying to take a different approach. And consumers will be paying-- or by buying hardware or by purchasing products online.

So I think you're seeing choices, and I hope we'll see more over time, but I do feel like there's an increasing demand for digital privacy, at least in some parts of the economy.

SASHA ROMANOSKY: Let's reconvene in a year, we'll see how it goes.

YAN LAU: Yeah, so I was going to say, we're out of time. But thank you so much to our seekers and thank you so much for your questions and participating.

So just a quick reminder, we're going to have a break right now until 3:30 when Session Four, Tools and Ratings for Privacy Management starts, OK? Thank you.

[APPLAUSE]

[MUSIC PLAYING]