FTC PrivacyCon 2017
January 12, 2017
Segment 5
Transcript

MARK EICHORN: I want to thank everyone for the presentations through the course of the day. It's been a really interesting day for me. I'm Mark Eichorn. I'm an Assistant Director in the Privacy Division here at the Federal Trade Commission. And today we're going to be talking about information security for the next hour. We've got three really interesting papers. And we'll start with Amin Kharraz of Northeastern.

AMIN KHARRAZ: Thanks. Hello everyone. It's a pleasure to be here. Today I'm going to talk about our recent research on ransom work and the paper that we presented at [INAUDIBLE] last year. So, similar to any kind of malware attack, ransomware starts with a successful infection. A user get infected by opening a malicious document, malicious attachment, or visiting an exploit website or a compromised website that is under the control of the attacker. Or, simply by opening an malicious binary.

This is an innocent-looking word file that a user may receive in an e-mail attachment. Once he or she opens the file, the background process will start, connect to a remote server, and downloads the malicious payload. It can be ransomware. It can be any other malware. After a successful attack, the user sees something similar to this, which is called a ransom note. Depending on the ransomware family or the geographical location of the user, the user may see a different one. So after an attack, a user has two options. One is to pay the ransom fee with the hope of receiving the data back. Or not to pay and permanently lose access to the data. So after paying the ransom fee, the user received a decryption key, which is in fact the private key that is generated remotely in the server and is used to unlock the files.

So as you saw, the attack is pretty simple, and it has been in the wild since last decade. But the recent resurgence of this attack has resulted in increasing concern on how to detect this class of malware. So the recent attack on a health care system was on the news. Also, a university in Canada was forced to pay the ransom fee in order to get the data back. So the problem is important.

And the question is how we can detect them. Unlike other classes of malware that try to be stealthy as much as possible, ransomware is different. When a user is infected with ransomware, the user is informed that he or she is infected. Because, for example, encryption is involved in the attack, there is an entropy change. Entropy is actually a measure and score that shows the predictability of the information content. The higher entropy score, the more random the data is. Or, for example, there is a modal dialog or background activity that are predictable.

For example, because the attacker wants to force the user to pay, the malicious process encrypts as many files as possible. So you see repetition other the files. So if we use these high-level ideas and translate them into concrete defense models, it is possible to detect a significant number of ransomware attacks. So one way to do that is by designing a sandbox. So a sandbox is an isolated environment that is used by malware analysts to run the samples and get internal insight of the

sample. For example, in this case, we want to see how the cryptosystem works. Is it a new kind of ransomware attack, or is it similar to the other ones? If its new, what features does it have-- in order to generate signature, in order to do more research from a defense side.

So we proposed UNVEIL which is a sandbox. It's a behavioral-based approach. It sits in the kernel, which is in privileged mode. And it uses Windows kernel development technology like other antivirus products. Or, for example, backup solutions. They sit in the kernel. And we generate a fake but interesting environment for malware to run and sit in the kernel and look at their false system activity of the user mode processes. And we run the sample and look at how the sample interacts with the files and collect the corresponding information.

I'm not going to talk about what we collect, but there is a high-level table here. So this is a typical scenario. And on cryptowall attack, one of the current ransomware attacks, we collect I/O operation, which is the lowest level of-- is IRP, which is the lowest level of I/O operation between the file system and Device Manager. And as you see, for example, the process of opening files, reading their content, and encrypted with high entropy payload, which is an encrypted version of that payload. And then close the file. And as you see, exactly the same I/O operation is performed on another file, which shows that there is some kind of repetition that happens during an attack.

And as you see, svchost is a benign process, but attacker or malware authors, in order to evade detection, kill the original one, and just invoke a customized version in order to do those types of attacks. In order to evaluate our approach, we use the unknown samples. In the wild, we collected 1,200 samples every day. We had an infrastructure worth 56 UNVEIL-enabled virtual machines. We ran the samples in a parallel fashion and collected the results in order to do the tests, the analysis. We cross-checked our results with virus total, which is a collection of AV Scanner-- you can easily [INAUDIBLE] release, submit a sample and get a report about those samples. So in our approach, a new detection is that if a sample is labeled or detected as ransomware in our system, we submit it to the virus total. If none of the AV Scanners in virus total detected that as ransomware or any kind of malware, we label it as a new detection.

So we tested our system for a year with 148,000 samples, and this is the higher-level results of cross-checking with virus total. And we submitted each sample six times in order to see whether the detection ratio changes during this period. And we define a measure which is called pollution ratio that is between 0 and 1. The values close to zero mean that none of the AV Scanners were able to detect the sample, and the value was close to one means that all the AV scanners were able to detect that sample. So as you see, in the first submission, around 70% of the samples had pollution ratio equal to 0. That means that in 70% of the samples that we detected as ransomware, other AV Scanners were not able to detect that.

We continued to submit the samples to see whether this ratio changes. And as you see, it changes, but it's not significant. And this means that most of the samples that are detected are either detected by a large number of AV Scanners or very few of them. This is the deployment scenario or use case for UNVEIL. As you see, it can be attached to a file system of any public or private malware analysis system in order to analyze the sample and give a report on whether it is ransomware. And if it's ransomware how the internal activity of that sample looks like. On the

other hand, we collect a sample and generate a malware dataset. Those who work in malware research know how valuable this dataset is in order to test new techniques and evaluate the detection approaches.

Another deployment scenario is endpoint solution. UNVEIL can be used as an augmented service on Windows machines, on legacy systems, and our evaluation shows that it incurred less than 3% overhead for realistic workloads. And we guarantee zero file loss.

So the conclusion of this work is that ransomware is an important threat. It's a big challenge. But the good news is that it has specific behavior. So our approach was designed to target those behavior by introducing a concrete model. And we showed that our approach works in practice and is useful to be used by the security community. Of course, there is room for improvement because it's adversarial research, and we try to add more features to it, more service to it. And this right now is used as sample-sharing among the researchers, and the cyber defenders, and also reverse engineers. Thanks.

[APPLAUSE]

MARK EICHORN: Our next speaker is Damon McCoy, and assistant professor at the Tandon School of Engineering at NYU.

DAMON MCCOY: All right, thank you. So I'm going to be talking today about some joint research with a whole bunch of people from Google and myself, from NYU, basically looking at one of what we feel is one of the larger distribution mechanisms of unwanted software these days. And for those of you that aren't familiar with the term unwanted software, unfortunately probably most of you at some point in time have encountered unwanted software.

These are things like browser extensions that insert ads into people's pages or do pop-ups on people's browsers or either extensions or applications that change someone's default search engine to some paid search engine that probably no reasonable person would actually use. Or there's another class of these that we call scare-ware which are basically optimization tools or AV tools that try and scare the user into paying for some kind of normally overpriced subscription to these kinds of software.

And so, we've done some previous studies looking at specific flavors of unwanted software like injectors, but we had this big question. We found that there's hundreds of millions of people at any time that are infected with this unwanted software. But how does this unwanted software actually get installed on people's computers. And so this is kind of the question that we set out trying to answer. And when we did our investigation we quickly found what are called these commercial pay-per-install programs. And so on the black criminal markets, this is a common term for installing malware on someone's machine. But there's also this commercial equivalent of these pay-per-install programs. And generally, how this works are you take some software and then you bundle it with some of this typically unwanted additional applications.

And the promotion of this is normally fairly deceptive. So probably, if you've ever run into one of your flash player's out of date, or there's some security problem with your computer, normally

these are trying to deceive you and trying to get you to unintentionally install some of this unrelated software. And some of our work of this year-long investigation is to try and understand the relationship of unwanted software to these commercial pay-per-install programs, understand some of the deceptive promotional materials, and some of the negative impact that this whole ecosystem has had on users.

And so first, I'll kind of take you behind the scenes of commercial pay-per-install programs. And commercial pay-per-install programs work primarily off of this affiliate model, where there's, in this case, advertisers. Advertisers that not just want to show you something, but these are software developers that are willing to actually pay to have their browser extension, their software installed on people's computers.

And so, generally how this works is these advertisers have this unwanted, primarily unwanted software that they're trying to install on someone's computer. And they go to these commercial pay-per-install programs. And they say we have the software. We'll pay you say $1 per install in the US if you can do this. And these pay-per-install programs then accept these advertisers on, and they don't actually install anything typically. What they typically do is they line up what they call publishers or affiliates that sometimes develop software or sometimes they have very deceptive practices on how to actually distribute this unwanted software. And then the publishers then get paid by the commercial pay-per-install platforms and the commercial pay-per-install programs get paid by the advertisers. So this is kind of the flow of money within the system and the set of players.

And so unfortunately, as we see throughout our study, is this decentralized distribution oftentimes lends itself to abuse. And so these middlemen platforms can oftentimes disavow themselves of abuse that's either being done by the advertisers or abuse that's being done by publishers by simply removing that advertiser or that publisher from their network while their network can still persist over the long term.

And so as part of the study, we monitored these pay-per-install networks. And we downloaded some of the unwanted software. It's not hard to find the unwanted software. They make it relatively easy and free. And so we downloaded some of what they call the bundlers in the commercial PPI infrastructure. The Google engineers reverse engineered these bundlers to essentially make it so that we could download the software from the advertisers without actually running the real bundle or installer.

And so we did this [INAUDIBLE] just simply through a series of HTML requests. And so we can look at the initial HTML requests that do some system fingerprinting, figure out what software-- they call them "offers"-- to offer the user and then report successful installs that happens and then optionally splash the user with another screen for some more normally deceptive products to try and hoist upon them.

And so the Google engineers-- why we have all those authors is that we had a fairly significant analysis pipeline that had to be built. And so this took a lot of work on the part of a lot of people at Google to build up this pipeline of things that began by extracting these offers. Then they have a binary execution environment. Then they can annotate what they're getting.

And then finally, we can cluster these things into the various varieties of software. And so we focused on looking at four of what we felt were some of the major pay-per-install networks based on some initial investigations that we did. So these were Outbrowse, Amonetize, Installmonetizer, and OpenCandy. So again, we milk these for about a year. And we, in total, saw about 400,000 offers about a little over 1,000 of them being unique actual offers.

And so what's happening here a lot of times is that the AV companies oftentimes flagged these offers as unwanted software. So similar to the ransomware and other malware families, they're simply doing lots of variance of the same on one software oftentimes.

So now for the analysis. We can do some analysis, and lo and behold, almost all of this is unwanted software. So there's a whole bunch of the major injectors that are being distributed by most of these pay-per-install networks. There's the browser setting hijackers normally that target the search engine defaults. And then there's this large array of cleanup and optimization type tools that are typically the scareware that are being distributed by these commercial pay-per-install networks.

And so we can run these through VirusTotal to look at the labels of these, and on average about 59% of the offers-- the software that they're distributing is flagged as unwanted software by at least one AV company. So confirming our suspicion that most of what they're pushing is unwanted software. It varies by the commercial PPI network. As you can see, some of them distribute less unwanted software. Like OpenCandy is probably one of the ones that are better. And then there's other ones where virtually everything that they push is flagged as unwanted software by some AV company.

The other interesting thing that we did when we were looking at these bundlers is we noticed that they had some ability to check registry keys within the computer. Some this might be legitimate. Some of us might be avoiding them being ripped off by the advertisers by making sure that this is a real actual computer that it's being installed on and a virtual machine or something like that, but it appears that the publishers are also using this to detect what anti-virus solutions are detected on the host. And if a particular anti-virus is installed that would flag that software as being unwanted and potentially remove it, that won't be installed, and something else will be installed that won't be flagged by the AV. So we see some of this evasive behavior and kind of catering almost to the unwanted software as a first class citizen of their pay-per-install networks.

The other thing we note is the price of these installs. So the US generally fetches the highest price in terms of install at about $1 to $1.50 per install. This might not seem like a lot of money, but if you compare this to the black market for malware installs, this is an order of magnitude larger. So normally an install in the US for malware will run you about $0.10 to $0.15. And so they're making an order of magnitude more money on each install than the black market from our installs.

And so finally, I'll talk a little bit about the impact of unwanted software on users. And so Chrome and Google started throwing up warning messages before users visited sites that they had flagged as distributing unwanted software. And also, Google's anti-virus, hooked into Chrome, also started warning users when they were downloading unwanted software from these

programs. And so, Google tracked the numbers of how many of these warnings they showed. And it averaged about 60 million warnings every week. This was actually three times more warnings than they're showing for malware. So they're shown three times more warnings for the unwanted software, showing that this unwanted software-- at least according to Google detection-- is much more prevalent than malware is currently.

And so quickly, I'll run you through some of the deceptive practices. And so, depending on the commercial pay-per-install program, some of them actually give you deceptive materials if you sign up as a distributor of this unwanted software. So this, again, shows that some of these programs are very much treating bad affiliates as first class citizens of their affiliate programs.

The other thing that we can see is that they oftentimes cycle domains. So as Google detects bad domains and starts blocking those bad domain with warnings, they'll cycle the domains for other affiliates of how they distribute software. And then finally, they've done some evasion techniques where Google was inspecting zip files, but they weren't inspecting rar files. And so they're kind of exploiting gaps in Google's detection coverage to try and evade detection by the antivirus. So the takeaway is there's tens of millions of users at any one time that are infected by this. These PPI look like one of the primary distribution vectors of unwanted software. And these misaligned incentives between advertisers, publishers, and middlemen seem to provide a ripe ecosystem for this type of abuse. So thank you very much.

[APPLAUSE]

MARK EICHORN: All right. Our third speaker is Mohammed Mannan, Associate Professor at Concordia University in Canada.

MOHAMMAD MANNAN: Hello everyone. Thank you all for staying so late in the day. So this talk is about our analysis of some leading antivirus products and some parental control applications. This work was done with my peer [INAUDIBLE] Xavier. And in this work, we did not actually analyze how effective some of these anti-virus of parental control applications are. We did not look for their functionality. We looked for how they interfere with the transport layer security protocol, which is the backbone of our modern secure internet communication. So HTTPS is used a lot, and TLS is the underlying protocol which secures HTTPS, which secures your connection from your browser to a website. And now, in recent years, after 2013, HTTPS adoption has increased very rapidly, and more than half of the websites in different estimates now are using HTTPS, which is good news for security.

Now that's not so great news for anti-virus or parental control applications which in some cases want to inspect the traffic, the encrypted traffic. So to do that, they install what we want a TLS proxy in the client side so that they can inspect the encrypted traffic. So to do that, they actually break your HTTPS connection from your browser to a website, and they create actually two connections. One between the proxy and the website, and the other between the proxy and your browser.

So now, from the external website's perspective, like here in Wikipedia, the SSL connection or HTTPS connection actually ends with your anti-virus product. So now it is supposed to actually

verify the site certificate that it is receiving from outside, also securely establish the TLS connection. And then it also has to send or issue a new certificate for your browser so that the browser can also be happy about it. And to make the browser happy and accept that new certificate which is actually issued by the TLS proxy in your AV, the Avi, the TLS proxy, must also insert its own root certificate into the system store which the browser relies on for trustworthiness.

Now, in terms of client and TLS interception, there are some advertisement-related products which also used to do it. And you can, of course, remove them. For anti-virus products, they generally do it to check web malware. Parental control applications-- they do it to check for unwanted content, to block certain websites and certain content.

Now for unwanted adware related products, you can remove them. For anti-virus or parental control applications, they are something that you may actually want. Or they are recommended by experts, not necessarily security experts in all the cases. By the way, can you guys please show a hand if you use an anti-virus or a parental control application? So that's significant.

So this might be related to you in some cases, but not all products actually perform TLS interception. So that's the good news. So we analyzed 14 products in March and August of 2015, sometimes different versions of this same product. And out of these 14, actually 12 were doing TLS interception, and two others were inserting a TLS certificate, a root certificate into the system, but they are not doing any TLS interception. And we are not sure why they are doing it.

And in our analysis of these 14 products, we found that, except one, all of them were downgrading your TLS security. So by downgrading, what I mean is that they may allow a man-in-the-middle attack from a network, who can actually impersonate a website or extract some secrets from your TLS connection. By impersonating, what I mean is-- as you can see here, this is FTC's website, although you don't believe it, because it doesn't look like FTC's website.

But from the browser's perspective, here in Chrome, it sees that this is ftc.gov, https://www.ftc.gov. And it's happy to accept that OK, this is a valid certificate, because it was issued by Net Nanny and because of how Net Nanny was installed in this system, it also inserted a root certificate in the system store which is used by Chrome to validate any certificate that it receives from outside.

So, in this case, I asked my student to create a screenshot. And he did it in the Christmas time, so here is some Christmas gift for FTC, maybe. So it's very easy to do, because if you can somehow extract the private signature key that is installed by Net Nanny in this case, you can actually issue certificates for anyone that you like. But for that, you have to read that private key in some cases. In some other cases, you don't have to run any code in the system.

[INAUDIBLE] that we are going to talk about-- they don't necessarily need you to run code in the target system or need any admin privileges in the target system. So, as I just mentioned, some of these TLS proxies, they necessarily become a local CA to your machine. They can issue any certificate they like. So we wanted to see how they protect the private keys that are installed in

the system, because if you can extract that private key, you can actually send any certificate-- you can impersonate any site to that particular system. So this is kind of important.

So we found that most of them actually do not secure the private key in any sensible manner by some user level code, no admin privilege needed, you can read the private key. Sometimes they encrypt it using a password, but we were able to crack most of them, or all of them actually. What is worse, is that two products actually came with pre-generated certificates, so you don't need to run any code in the target system. You just install it in some place or just get the private key bundled in the installation. If you can extract it, now you can impersonate websites for any user of that installed product.

And an interesting fact I want to highlight from here is that, in many cases, when you uninstall the product, the root certificate that was installed in the system is not removed. And that certificate, in many cases, is valid up until 20 years. So they system vulnerable for many, many years to come, even if you uninstall them.

In terms of site-certificate validation, what your browser always does-- we found that three out of 12 were doing absolutely no validation. So you throw any certificate from outside-- an expired one, an invalid one, a self-signed one, whatever it is-- the proxy will say yes, I accept it. And actually, it will issue something for the browser which the browser will simply trust, because it has the proxy's root certificate installed in the system.

And we found that many of them are using duplicated SSL version, and duplicated cipher [INAUDIBLE] which are unsafe to use now. But they hide this fact from the browser, and they accept any such connection from the outside world.

And, in some cases, they actually hide these facts from the proxy. So whatever TLS parameters they are receiving from the outside world, they are not transparently forwarding them to the browser. So that they are not giving the browser an opportunity to find out these issues. So one example here is if a SSL 3.0 connection is coming from outside, the proxy may actually upgrade that connection to TLS 1.2. So your browser will see that, OK, this is the most secure connection that I can expect. So it will not complain about anything. I'll skip this slide.

In terms of recommendation, the main thing was just don't do it if you cannot do it properly. And there are some proper ways to do it. One we recommended was to use this feature that is now available in Chrome and Firefox called TLS keylogging. And in this one, you don't have to do any TLS validation or anything. You can decrypt the traffic if the user chooses to expose the TLS from the browser, and then you can do your job on with the exposed key.

So what I want to highlight is that our analysis was not about the functionality or effectiveness of this product. But it's more about if they interfere with other security mechanisms, like something as widely relied on as HTTPS. How should we handle such a situation? So this is not a technical problem. This is more of a policy problem. Another thing is it's not only for this product, but any product, if the user chooses to uninstall it, and if that product actually still keeps the system vulnerable for many years to come, how should we deal with such a situation? So maybe the audience here can shed some light on that, and I want to conclude it here. Thank you.

[APPLAUSE]

MARK EICHORN: All right. So we welcome questions from the audience. So come up to the microphones if you have a question. I will sort of summarize a couple of the themes that I took away. Obviously the papers are very different, so it's difficult to come up with themes that go across.

But the main one that occurred to me was that it's a challenging environment for consumers to secure your systems. Right You've got people looking to put malware on your system. You've got a whole industry looking to put nuisance or other unwanted adware or ad-injectors or scareware on your system. And then you've got this irony of, people, when they do want a tool to improve their security, there's this effect of sort of increasing their risk to man-in-the-middle attacks to some extent.

So the other thing that struck me is the theme-- not in all the papers, but in two of them. The first two were just, again, this theme of the technology arms race that we've been hearing about over the course of the day. Like in the previous panel about how there might be a safe browsing countermeasure to some of this software and then there's a countermeasure to the countermeasure. Or similarly with the ransomware-- the paper described certain techniques like the stalling technique and so forth to respond to maybe attempts to analyze the ransomware.

So why don't I dive in with some questions. And I guess I'll start with Amin. The one that jumps out to me from your paper is is this something that-- how will we implement this? Is this something that consumers could use on their side? I know you said that researchers can use this now. Is this something that a consumer would be able to use to sort of help keep their computers clean of ransomware, or is it something anti-virus companies would use?

AMIN KHARRAZ: So thanks for your question. So right now, it's specifically designed for malware research and reverse engineers. In practice, in order to defend against ransomware, there are at least three ways. One is educating the users. Users should not click on every advertisement they see on their browser. Users should not open any attachment they see in their inbox. Or they should use a good backup policy in order to make sure that the data will not get lost. Incremental backup is something that everyone is recommended in order to defend against ransomware.

So this is the most reliable approach to defend against ransomware for consumers. But there are also other ways to minimize the risk and also help the community, for example, that we are working on. One system that we propose actually defines concrete model to detect ransomware. This doesn't mean that current AV scanners cannot do that or don't have it. They may have it, but they don't-- they have it for other malware families, but the thing that they probably don't have is a concrete model to detect ransomware. Because ransomware is doing something in a very specific form.

And referring to your question, the approach that we have, by minimal modification to the tool, we can use it in the end user. This is actually the research that we are doing right now. It's a behavioral-based tool that sits in the user's computer and tries to find behavior that is similar to

ransomware. So there are some challenges, for example, usability is an issue. But we didn't have a huge overhead on this when the system actually is working.

But, for example, something similar like false positive is the thing that may bother people. So we actually work on that and eventually a tool that can be used for the end user will be published in probably two to three months from our lab.

MARK EICHORN: So Damon, your paper talked a little bit about the way that consumers get this-- that usually there's some type of "consent", usually to a deceptive offer or something. So I'm usually doing something to click on and download. And maybe I'm getting a lot more software than I expected. There's also some cases where I might click on something, and I'm then taken to somewhere where it's just a drive-by download on my system. Is there some better technical way-- are there technical controls that can sort of address this better?

And sort of related to that, it seemed like, at least in some of the packages, the paper seem to suggest that the bundles were presented in a way where I could maybe pick and choose which software I got. So I don't know if that's one possible solution.

DAMON MCCOY: I would argue that that'll become another arms race. If customers actually became savvy enough to try and unclick things. So they make it so that if you click, as most people do, the express install, that you get all of the unwanted software out of these bundles. And they make it purposely deceptive. You have to do multiple clicks and read multiple long thick lawyerese type things to understand that you're going to be getting this unwanted software, and that you have to click on this bizarre option to try and prevent getting this nuisance, unwanted software on your system.

So I think that they've already crafted the system in such a way. So I think educating the user might help to warn them to just not install the stuff in the first place and not run programs from this. But I think we've already done enough educational campaigns that people should be relatively aware of that.

And so I think, at this point, there needs to be regulations and lines drawn as to how deceptive can you make this stuff before you've crossed one of these lines. And again, not all pay-per-install programs are equal. So some of them are much more deceptive than other ones, but some of them have just gone down a path of making it virtually impossible to avoid installing this unless you're just extremely diligent and you want to read through lots of really thick text and click on every option before you continue on with your installation process.

MARK EICHORN: So it's not like there's some really promising technical solution to give me, as the consumer, more control over this. The best approach is sort of a more regulatory or policy approach. There's also the option that you see in mobile where the developers are more vetted and things can be removed from the stores and things like that. And there, you do see some unwanted software, but you probably see less and less harmful unwanted software in those kinds of models.

But again you give some freedom now, you're at the behest of the curators of these stores as to what you can install and what you can't install.

MARK EICHORN: Mohammad, so if I'm a consumer going to choose between different anti-virus products-- some of them use proxies and some don't, right?

MOHAMMAD MANNAN: Right.

MARK EICHORN: So are people sort of aware? Is there any way for me as a consumer to know that this product uses a proxy and this one doesn't, just to compare between products on that basis?

MOHAMMAD MANNAN: So whether they do TLS proxying or not, generally this is advertised. So because this is a feature that you may be interested in. So AVs or parental control applications-- they don't hide it. It's kind of a selling feature for some. So you can check it. And in terms of how they implement it-- whether it is safe to use, you can see some of our reports, and then maybe choose which one is best suited for you.

But in general, I don't use any of these products actually. We test them, but we don't use them. So if you have to choose something, my general recommendation would be choose something that does not do it, because it's actually pretty complex to do it properly. Browsers that continually update it. Some of these companies that manufacture browsers, they spend a lot of effort to keep browsers safe. And a significant part of their effort goes into securing TLS. And if you rely on another product which is just a side feature for them, to be as cautious as the browser manufacturers, I really doubt it. So it's probably better to avoid products that are doing TLS proxying.

MARK EICHORN: Damon you talked a little bit about the sort of assistance that the PPI distributors provide to sort of allow the anti-virus detection. And maybe the chrome safe browsing avoidance, and so forth, to sort of facilitate this. But could you talk about the bad intent or not of the other players in the industry? Of the publishers and the advertisers? Especially, at some point, particularly the people whose software is being used as bait, have something that people want, right? Not really?

DAMON MCCOY: Again, this varies case by case. But a lot of times, it's very deceptive. At They'll try and package it as something like a flash update or something like that, where it's not even really a flash update of any sort. It's just completely deceptively marketed. Or they'll have the flashing pop-up that says you're infected with something. And they'll get you to install this.

Or they'll throw out a bait video or some kind. So there'll be some salacious video, and you go to try and play it, and they'll say, oh, you need to download this codec. So the publishers, the distributors, and the affiliates oftentimes have very sharply honed and very deceptive things. And so, oftentimes, they don't have a real, viable software product that anyone would want to download. It's just purely based on deception on their part. So we did some analysis of this to show that a lot of the affiliates are operating in a very deceptive manner.

MARK EICHORN: Mohammad, going back to you on the proxies, is there any information that I would have as a consumer to choose-- as you highlighted in the paper today or in the presentation, one of the products was not as bad as the others. Is there any way for a consumer to assess which ones are doing it properly and which ones may not be aside from reading your paper?

MOHAMMAD MANNAN: I doubt it, because my PhD student has spent a lot of sleepless nights to analyze these products. So I highlighted just some results here. We I did not talk about analysis part at all. It's pretty complicated. So for a consumer to understand it, I don't think there is any chance.

And the one product that I mentioned-- probably I can name it-- so it's actually Avast which was the most clean one, and we did not even contact them. So we actually contacted all these companies where the product was affected. We did not contact Avast, because we found some issues, but it was not big enough. But in their later versions, they actually implemented some of the recommendations that we put forward.

So at least, I don't know whether it's a good AV or not. I don't use it, but at least in TLS proxy, it seemed pretty good.

AUDIENCE: I have a very quick question for Mohammad. So you said that some of these companies leave the root key, which is [INAUDIBLE] for 10 years even after they uninstall. How to get rid of that root key?

MOHAMMAD MANNAN: So you have to probably manually do it. We did not really use any tool to completely uninstall them. I don't know whether they are-- there are some uninstallers that may help, but I'm not sure. The only way that I can think of now is to manually remove that root key from a certificate store. Now, I cannot actually ask a regular user to do it. They will not probably find out where operating system or browsers are storing these certificates, and go there, and find out this specific one which was installed by the AV or the parental controller application, and remove it. So it's possible, but not so easy.

AUDIENCE: I have a question, I guess for all three panelists-- do you have a sense or any thoughts about whether or not operating systems themselves can do anything to help mitigate some of the things you've identified? So can operating systems include some of the detection properties that Unveil uses for example maybe? Or can operating systems identify when bundles are being installed and maybe inform the user or maybe block them in some way? Can operating systems basically help mitigate this in some way?

AMIN KHARRAZ: So for malware research, the main issue that we have is evasion. So once the detection system or detection service gets active, malware authors, cyber criminals, try to think how they can evade it. So if you come up with a service or an operating system that provides a service, if it's based on a certain signature, what happens is that attackers look at it, look at how the countermeasure works, how the security control works, and try to reverse engineer it, and try to bypass it.

So this is actually the main challenge that we have. It's a game. Companies come up with patches, companies come up with new services, but at the end of the day, we see more advanced attacks on exactly the same service or the same feature.

AUDIENCE: Well I don't necessarily mean signature-based detection, but the same way that you and your work have identified, say the I/O properties of ransomware-- can the operating system sort of incorporate identification of those same properties to try to head off ransomware, sort of in the same way? So basically the operating system now has the responsibility to do some of the detection. Or it can identify some of the same properties that you are able to identify in your work and basically help defend against.

AMIN KHARRAZ: At least, I'm not aware of that. And one issue is that once these types of systems get active, usability becomes something that these companies have to make sure of. So, for example, what happens in Chromium is that-- three or four years ago it was not like this, but right now you have automatic updates. It doesn't ask users to update their browser. It actually does it automatically, because the default assumption is that the user does not do it, for example, for many reasons.

And I think the main issue that operating systems may have with this, companies may have with this, is they have to do extensive experiments on the usability. And for the approach that we were proposing, it's like some other anti-virus companies or backup solutions that work right now on top of the operating system. Does Microsoft provide backup service? Yes it does. So it does in a certain level. So it can be added-- you should actually show that it's feasible by doing that. For example, by the prototype that we provided, we show that it's feasible. So it's actually the companies that have to decide, for example, how to proceed in activating these types of services on the operating system.

AUDIENCE: Could I ask-- our time is up for the panel-- could I ask for a round of applause for the panelists?

[APPLAUSE]

MARK EICHORN: We're going to roll right into the final close-up panel. It's my honor to introduce Jessica Rich who's the Bureau Director for the Bureau of Consumer Protection here at the FTC, and my boss. So I'll say a nice thing about her which is that she's been a leading thinker about privacy for about two decades. So So Jessica, take it away.

JESSICA RICH: Hello, you're still here. We're going to try to make this lively. I'll introduce you guys. These guys get introductions, since this is like a discussion panel. So I'm Jessica Rich. I'm director of the Bureau of Consumer Protection. Thank you so much for coming here today. It's been a really great event. So this is our wrap-up panel. We get to talk about some of the themes and challenges we've raised by the research presented today and sort of talk about whatever we want, which makes it terrific. I have a terrific-- that's the word today, terrific-- group of panelists to help me with this.

This is Howard Beales. He's a professor at GW School of Business. And he may be most famous for his three-year stint here as director of the FTC's Bureau of Consumer Protection, the job I now hold. Howard was my boss, a very good one. And during Howard's tenure the FTC established it's Do Not Call Registry, launched the FTC's data security program, and pursued an approach to privacy emphasizing tangible harms, just to give you a little background on Howard, among many other things.

Then we've got Andrew Stivers. He's the Deputy Director of the FTC's Bureau of Economics. In that role he weighs in on every matter, yes, every FTC consumer protection matter, including enforcement reports, and research involving privacy and security. We stole Andrew from the FDA where he led the economic analysis of major regulations including food package labeling and calorie labeling restaurants, which we now all see.

Deirdre Mulligan is a professor at the School of Information at Berkeley. She is a well-known researcher, author, and thought leader in the privacy and technology space. And among many other things, she serves as Chair of the Board of Directors at the Center for Democracy and Technology.

So I'm going to start by pitching a few questions to our panelists. But I really want to open it up to the audience. You can ask about any of the topics we've coveted today or any topics you think these guys might know something about. So please think up some good questions as we're talking here.

So why don't I start with a question about consumer attitudes towards privacy, an oldie but a goodie. We heard several presentations today that bear on the question of whether and how much consumers care about privacy. We heard that consumers are downloading ad-blockers in record numbers. We heard that consumers' level of concern about privacy may depend on context. We also heard that even after Google changed its privacy policy to combine user information across platforms, it had little long-term effect on consumer searches for sensitive information. And we all know that consumers don't hesitate to use websites and apps that collect enormous amounts of information despite their stated concerns about privacy.

So my compound question to my panelists is based on the research you've seen and what you've observed over the years, do you think consumers care about privacy? Do you think this question is important to the privacy debate? And if so, what is the best way to measure whether and how much consumers care about privacy? So why don't I just start with Howard, right here.

HOWARD BEALES: Well I think this has been a very interesting conference. But this has sort of been very much the weeds. And I think we need to sort of think about the forest that this is part of, and how it all fits together.

To answer your question of do consumers care? Well some do, some don't. What do we do about that? Well, I think markets are really good at giving consumers what they want. It's actually why we have an illegal drug problem. Markets are good at giving people what they want. And that's where privacy-sensitive people are going to find the kinds of protections that they want. It isn't feasible to expect, it isn't reasonable to expect consumers to make instance-by-instance decisions

about what they're going to do. The number that stuck out for me was if you approved each mobile app request for data, there'd would be 231 requests per hour.

That doesn't work. Where consumers who care about privacy are going to have to get that privacy is from products that help them protect it that are marketed as privacy protection products. And if people don't buy them, it's because they don't care enough. In the same way as if they don't buy any other kind of a product they don't care enough to be willing to pay the costs.

JESSICA RICH: Well hold that thought, because I want to come back to the arms race, which I think somewhat contradicts a little bit what you just said. That everyone's been talking about. Well, OK, we'll come back to it. But let me get the input on the consumer perceptions from Andrew.

ANDREW STIVERS: Sure, so yes, I would echo much of what Howard said in terms of yes, some people care, some people don't. I do think it's important, and I think, for me, as an economist, the sort of best way to determine how much people care is to look at their consumption decisions.

I think that this field gives us a particularly interesting and complicated problem for consumer choice for a lot of different reasons. One big one is that your choice about privacy is connected to all your other choices about privacy. So if you fool with the very intricate and detailed privacy controls on a particular very popular social website, that may absolutely have no effect on your overall privacy profile, because you share information in all sorts of ways. And it's going to be really hard for consumers to figure out what's the value to me of spending a lot of time investing in privacy in one dimension when there's 16 other dimensions that maybe I either don't know or don't understand or don't care about for whatever reason.

So it's a really complicated question to get a look at the market and understand what people's values are in addition to the heterogeneity that we see. And I always kind of go back to-- and I'm going to give Alessandro Acquisti a bit of a plug for a paper that he put out this fall called "The Economics of Privacy" that was in the Journal of Economic Literature. If you haven't read it, I strongly, strongly encourage you to. And I apologize to him if I paraphrased incorrectly, but what I get out of his paper, from all the literature that's been done, is wow, it's complicated. So that's great for me, because it gives me a lot of job security. But it makes it hard for us to know what's right for consumers in any particular situation.

JESSICA RICH: Deirdre?

DEIRDRE MULLIGAN: So I don't think we need any more research telling us whether or not people care about privacy. That might not be a popular thing to say, but I think we know people care about privacy. I think there's some interesting work to be done about what kinds of privacy-- there are many concepts of privacy. Is it about control over information, or limiting access to the self, or zones of decisional autonomy that might be most relevant to consumers in different settings?

And how important is it, under what conditions, and what time of life, what context? And I think the FTC has made some really important strides in that area to say, in Helen's words, of course privacy matters, but a contextual inquiry is really important to understand what privacy means in that particular moment, in that particular context and the environment.

But I think there's this really important question, which I think is part of what both Howard and Andrew were going straight to is if people care so much about it, why do they find it so very difficult to protect it or to act in a way that would seem to conform to their stated preferences? And I think that's a question that we, one, have lots of answers.

And I would point to the same article and say that not only has Alessandro in that article said it's complicated, but in that research and numerous other papers produced by him and others like Jens Grossklags who was here earlier today, have highlighted all of the ways in which information asymmetries, cognitive biases, the public good nature of privacy, architectures, default settings, policies undermine people's ability to get what they want from the marketplace. And I would point to another paper that was written by none other than Joe Farrell who held, I believe, Andrew's job.

ANDREW STIVERS: My boss's job.

DEIRDRE MULLIGAN: Your boss's job, right. And he was exploring the challenges to the market provisioning of privacy, and he was looking specifically at privacy conceptualized as a final good. Meaning that people are actually shopping for privacy. And as we all know, very few of us rarely do that. Usually privacy is an aspect of a good or service that we're shopping for. Or maybe an instrumental good that we're interested in.

But even thinking of privacy as a final good, he noted that there were some particular reasons to question whether the information and other conditions efficient contracting would hold in the marketplace for people to actually get good policies that were consistent with their expectations. He wrote that if consumers understood the implications of different up-front privacy policies, and the policies are truly effectively disclosed, including drawing consumers' attention, then the demand shift effect-- consumers could basically shop with their feet, because you could discipline market actors. There would be incentives for firms to choose responsible policies.

But given what we know about cognitive and informational barriers from Alessandro's work and Jens's work and others', that those present real barriers to those sort of informational knowledge that consumers need to actually shop with their feet. And Professor Farrell said that one could see-- he said that if these things don't happen, you could in fact have a dysfunctional equilibrium in which few consumers devote much attention to disclosures, disclosures are vague, non-commital, or even if explicit, mostly ignored. And the privacy policies chosen are inefficiently non-protective.

So what this means is that there would be very little incentive for businesses to present consumers with good policies, because they wouldn't be rewarded for them. And that's even assuming that people are shopping for privacy. But if we assume that people are not shopping primarily for privacy, but it's just one aspect of a good, the reason that something like the

Consumer Reports activity might be ever so important is that privacy isn't a good that you can shop for on the front-end. Because even if there's a statement, most people don't have the expertise to actually digest it and understand it.

Second, it's not an experience good. Like I can get this bottle of water and drink it and decide it's good water, even though I couldn't tell from the outside. And then I want to buy it again. But privacy-- I can't actually assess whether or not firms are doing what they say for the most part, because so much of what happens to my data is behind the scenes. So I can't even understand if they're providing good privacy once I've had a relationship with them. So it's not even an experience good.

And so Consumer Reports is stepping in, because they're saying this is a credence good. This is an expert good. This is something where we need expertise to actually assess whether or not the privacy is good, whether or not it's being followed. We might need to look internally.

So I think there are lots of reasons to question whether or not the market is going to adequately produce the kind of privacy consumers want, and it's not because people don't care. And it's not even because companies don't care. It's because it's a really complex good.

JESSICA RICH: So I think everyone here decided it was complex, and we also heard a number of studies that made the same point that we've seen in a lot of other [INAUDIBLE] that it's really difficult for consumers to protect their privacy. It's really difficult. The pay-per-install paper that we just heard about illustrates how consumers are tricked into downloading unwanted software all the time.

The TLS proxy paper shows that consumers who actually buy software to protect themselves may not know how to use it, and may, unwittingly, create additional risks when they download it. We know consumers consent to terms of service and EULAs all the time without reading it. And we're now seeing that households with connected devices like refrigerators and lightbulbs can be used for denial of service attacks, something consumers will most likely don't understand and wouldn't be able to defend against if they did.

So what is the solution? So you mentioned that Consumer Reports is trying to come up with ratings for privacy and security. That's one option. Is the world ready for machine-readable policies? That's something that's been pushed for years but hasn't taken off. Is more education the answer? I think we've sort of said no. Is legislation the answer? What is the solution to this? Howard may say no, the market is the solution, but let him say that. But why don't I let Howard say that.

HOWARD BEALES: Well I was actually going to say something different. There is no solution. And when we started doing information security many years ago, we said, repeatedly, security is a process. Privacy is a process. What we heard about today are what I think you would recognize in a broader context as flaws in software. And all software is flawed. We discover those flaws when good guys like the people presenting here discover those flaws and tell other people about them. Lots of them get fixed, not as quickly as we might like. But a lot quicker than Congress acts.

And it's inevitably going to be a back and forth, because there are smart people on the other side trying to find those same flaws to do bad things with them. And that is not going to go away, and a law isn't going to make it go away. And that's the core problem. It is a process. We have to keep monitoring new practices and new problems as they emerge and respond to those.

JESSICA RICH: Andrew?

ANDREW STIVERS: So I'd say two things. First, I think really very much in the wheelhouse, the FTC is making sure that the information environment is good for consumers, so that, to the extent that consumers want and use the information, it's available and non-deceptive. And that's something that I think is a little bit more clearly defined for us.

The other thing, though, that I think is really important to understand and to be explicit about--and I heard what I thought was sort of some confusion perhaps on the part of some of the speakers today. There are a couple of things going on. One might be the markets aren't working the way we want to. And so we would want to identify what's the market failure there.

How are consumers not being served in ways that they want to be served? And what are the blocks that Deirdre was bringing up that might prevent them from being served those choices. But the other issue that I think people are kind of conflating with some of these market issues is there are some social choices that are being discussed that are perhaps not market issues.

So, for example, in this country, we have a minimum wage law. And that minimum wage law is not primarily discussed as an economic issue. Given the minimum wage law, we can talk about what the costs and benefits associated with preventing young people from working. Maybe that means they go to school longer and there's some trade-offs there.

So we can talk about the costs and benefits of particular social choices, but I think people should be clear if they're going to make a social choice about what they think the privacy environment should be. And they can't trace that back to some kind of market failure. They should be explicit about, hey, I think the world should look like X. And here are the costs and benefits of getting to X if that's not the place we're at.

JESSICA RICH: What's going to drive more power for consumers in this marketplace in your view, Deirdre?

DEIRDRE MULLIGAN: Well could I start at the first question?

JESSICA RICH: Yeah.

DEIRDRE MULLIGAN: I was going to anyhow. So I think thinking about security and privacy separately is probably somewhat helpful. And you were, I think, referring to the IOT botnet attack, and I think it was a really excellent example of the public goods nature of security and the complete problem of suggesting that people are going to be able to solve their own security problems by making their own decisions. That people were attacked because of other people's poor choices in the marketplace. Now why did they make poor choices?

Were they actually choosing devices because they thought I don't want to pay an extra dime or $2 for security? No, they didn't understand that the things we're going to create security vulnerabilities that they were vulnerable to. But they might have if they were given the choice for the less secure or the more secure product. They might have chosen-- many people might be in the position where that extra $2 for the secure device is really more than they have to spend. But the fact of the matter is that those choices aren't going to be borne by them, because their devices were used to attack other people.

And so when we think about security for sure, and the negative externalities of the lack of security properties in devices that are proliferating and managed. Many of them don't even have interfaces for people to manage. That's a huge problem and it would be really, I think, damaging for all of us if we decide we're just going to let people make decisions in the marketplace, because we know even big companies don't always shop with security in mind, right? So I think we need to be thinking about, what are the minimal security properties that should be on those IOT devices? And surely updates is one of them, right? You should either be able to update it, or kill it, right? Because otherwise we're going to have lots of unmanaged, unupdated devices, that can be assembled into a nice little army to attack important assets. And that seems like a really bad outcome.

On the privacy side I think we're facing some really fundamental shifts, in that machine learning, right, is all about inferences. And if we assume that people are going to control their privacy by making decisions about what they disclosed to others, and it turns out that really benign data that you disclose can be mine to infer really sensitive things about you, the whole model of suggesting that individuals are going to be able to exercise control over who knows what about them, right, which is kind of essential to the fair information practices, doesn't really work. And so we certainly need better ways of giving consumers control over their information.

JESSICA L. RICH: Well a--

DIERDRE K. MULLIGAN: We're going to also need controls on what people can do with it. And you can look at something like the way we've handled genetic information, where we understand that the information doesn't just have implications for you. It has implications for other people that share traits with you.

JESSICA L. RICH: Well and adding to that complication is the idea that you don't even know all the companies that are collecting your information. Many of them are in the background. There's no more one on one relationship.

HOWARD BEALES: But that's the core problem with FIPS is that you can't think of this as a problem of controlling information. You can think of this as a problem of trying to control bad things that people might want to do with information. But if you think about controlling it as controlling information, that once you've given it to anybody can be passed, and in one of the papers I'll pass on all my friends information, because I don't care much about my friends. But the--

JESSICA L. RICH: Even me?

HOWARD BEALES: It's an information use problem. Not an information control problem. And if we expect consumers to control that, that isn't going to work. And I think it's inherently not going to work. The decisions are complicated. They require effort and costs of making decisions that most consumers, most of the time, are simply not going to be willing to spend.

JESSICA L. RICH: We have a question in the audience.

SPEAKER 4: Privacy rights we've had a very wide ranging discussion. And this panel's getting to a different aspect of it. I want to ask about the general data protection regulation approach in Europe, and why that hasn't come up, except indirectly? And is it because we're just waiting to see how that works out in mid-2018, or that it's only going to address a subset of the issues that were raised specifically today? Or any comment related to, and I think Deirdre's comments are headed in that direction quite a bit. What have we learned, or should we learn from what happened in Europe, or what's happening in Europe?

JESSICA L. RICH: Do you want me to start?

DIERDRE K. MULLIGAN: Oh, sure. So where to start. So I did some qualitative work looking at how companies in France, Germany, Spain, the US, and the UK, understand the meaning of privacy, and then how they operationalize it, in practice. And you know there is no one kind of implementation in Europe. They vary widely. Even though the overarching directive was the same, right? And so the general data protection regulation is going to create a kind of increased level of harmonization, but a lot of what matters in, kind of, practice, and how rigorously firms handle information, and whether or not privacy is expressed, kind of in paper, and legal documents, or built into technical systems. And our research was the product, not just of the rules, but also of the behavior of regulatory agencies, the actions of civil society, choices within firms. So you can have different regulatory environments that produce similarly good behavior, right? For example.

So the reason I say that is, I'm not-- like if you say, should the United States take the general data protection regulation, and import it into the US? So I mean, we have the political reality of that's not going to happen any time soon. We have, what, how many days? Yeah. But even if we could, right? If you look at the way in which firms behaved in Europe under the Data Protection Directive, I'm not sure that you would have been happy with the extent to which all of them were kind of taking privacy to heart. I think that there were some companies in the US who were doing more under a more ambiguous, consumer protection driven framework, because of things like the relatively activist regulator, that we had here at the Federal Trade Commission, the proliferation of NGOs that work in this area, the ongoing dialogue, privacy research. A whole bunch of things that influence firm behavior.

Substantively though, there is much to like, right, in what is in the general data protection regulation. But I think, while there is much to like Europe, too, even under the general data protection regulation is going to be struggling with how we address privacy issues with machine learning, AI, and robotics. And I think those are really-- because they're not just about data protection. They're about autonomy. They're about inference. They're about statistical learning. And so there's a yes, and not yet, not completely, right?

JESSICA L. RICH: Well, I'd like to add that some of the debate that-- and some of the research that we've been focusing on here-- does seem uniquely focused on a privacy regime that depends on consumers managing their own privacy a lot. Whereas, in Europe, they view privacy as a fundamental right, and there are more substantive requirements, not more enforcement, but more substantive requirements in the law. So a lot of the research today, and the debate around consumers managing their own privacy, is very uniquely American, I think, in terms of what we expect our consumers to do.

HOWARD BEALES: I would just say I think it's telling that in the United States privacy is largely driven by a Consumer Protection Agency, and that's really what it ought to be about, is consumers. It's not about data protection. Who cares? It's about protecting consumers from the possible consequences of that data. And that ought to be the focus.

JESSICA L. RICH: So some of the research today focused on the special privacy challenges created by mobile devices, which had now been around a long time. But they can collect more information, share more information, always with you, always on, all of that stuff. Even as we're grappling with these challenges we now have the explosion of internet of things devices, which we're seeing are not necessarily that secure, or private. And we, of course. Have artificial intelligence now on the horizon. What new privacy and security issues are on the horizon with these new products, and services? Are they the same, or different from the ones we're grappling with today. Who would like to take that? Andrew?

ANDREW STIVERS: You know, I think it's more of the same. I mean, you can potentially draw trends, and lines from kind of pre-internet to today. It's a matter of scale, right? The cost of information collection has dropped dramatically. And so, you know, for a long time the FTC was very interested in flows of information from seller to buyer. But as the cost of information collection gets lower, and lower, and lower, you know we're much more interested also in the flow of information from buyers to seller. So, you know, I see that increasing and the various vectors of information collection, and the granularity of the data, is certainly going to be increasing. And the potential amount of analytics will also increase. I guess it's going to continue to call into question the value of that data. We get more, and more, of this data on people. There's more and more-- just in terms of the privacy side, not necessarily the security side-- you get more and more of this data, are companies going to be more and more convinced that they're going to be more and more accurate? And maybe they are. Maybe they're not. So that kind of tension between accuracy, and privacy, is going to continue to be-- well, it's going to loom larger. But I think it's be more of the same, frankly.

HOWARD BEALES: But there is a market incentive to make better predictions, and because it's the-- I mean, you make money off the prediction, and not off of the algorithm that builds the prediction.

ANDREW STIVERS: Oh absolutely.

HOWARD BEALES: The prediction needs to be right. And if you can make better predictions, I think by and large, that's good for all of us. I mean, I think one of the things that was one of the

interesting findings in one of the early papers today was consumers care about the accuracy of the inference, about the accuracy of the prediction that's being made.

HOWARD BEALES: There's positive privacy rents. And negative privacy rents, right? If I'm the high value consumer, I really don't want greater accuracy. I don't want greater predictions, because I get a lower price, because you don't know I'm the high value guy. So I think there are some tradeoffs. But absolutely, there's more incentives to provide greater accuracy.

JESSICA L. RICH: Dierdre, do you have something to add?

DIERDRE K. MULLIGAN: Yeah, so I want to say, I think, first, I don't think that consumer protection is the only lens through which we should view privacy, right? It's an important human and political right. And while the Federal Trade Commission's activities have been incredibly important, I think that there are other ways in which privacy needs to be protected. And at times the activities that happen in the commercial sector have really important ramifications for government collection of information, as we found out very loudly and clearly during the Snowden revelations. On the, is it just the same? So certainly the data protection issues, and information privacy issues are going to be similar, but on steroids. And part of that is just that many of the devices that are proliferating, they have no interfaces. They don't provide any notice. There's no indication to consumers that data collection is happening. So it's very, very hard to use your exit, or your voice, right, to express a preference when you have no indication of what's happening.

But I think that there are other challenges that are going to arise, that I think are fully in line with the sorts of issues that the Federal Trade Commission looks at, which are things about, as we have all this data, and we start to be able to use it in real time, not just to alter prices, but to alter your experience of an environment to nudge you, and you know, concepts like undue influence, and manipulation. And you know, there's a huge potential increase in information asymmetries that I think are going to create an increasing set of issues for a functioning marketplace, that consumers are simply not going to know how it is that their experiences are being shaped. And so I think that there's another set of issues, that I think one can think of under the concept of privacy. We like to have this idea that there's some autonomous decision making happening, or some decision making that is free of undue influence, right? Or concepts people have talked about, fiduciary duties, and things like that, that I think are related to the increase in data collection, and the way in which it can be brought to bear on people's experiences, and life opportunities, that are going to be much more important than they have been.

JESSICA L. RICH: We have a question.

SPEAKER 6: Yeah, Dierdre, also--

DIERDRE K. MULLIGAN: What's your name?

MARK WEINSTEIN: Mark Weinstein, my company is Mewe, M-E-W-E.

DIERDRE K. MULLIGAN: Right.

MARK WEINSTEIN: And my question is really around this whole idea that privacy-- there's a certain irony to privacy being regulated by a consumer division of the federal government, when you know, it's actually written into our constitution, it's part of democracy, and if any of us here had told our great grandparents what was happening today, and that we allowed it, they would look at us and say you're breaking all, you know, the sort of integrity of what America is about, and they would never have allowed it. That companies can spy on us relentlessly, that the government could-- and you know the NSA program was discovered to be illegal. So my question is, what do we know in this country, right now, about the legislative process for lawsuits?

You know, I know there's a lawsuit in the state of Illinois about facial recognition on Facebook, and things like that. In Europe, the governments aren't regulating much more rigorously, even though, you know, for in the United States it looks like what used to be a Republican issue for privacy, and respect, and all that is out the window. But what can we do, and what do you see that's going on? Because I believe that ultimately it's going to be the courts that somehow going to pull it all back together for us. What do you see happening legislatively, or in the judicial process, or lawsuits around the country, to protect our privacy?

JESSICA L. RICH: So let me just say that, although I said earlier that the whole notion of consumers protecting their own privacy is somewhat American, not European, the FTC in a bipartisan basis, has been supporting for years a data security law so that consumers don't have to make these choices. They should be entitled to rely on the security provided when they entrust their information to a company. And many people at the FTC, not all, have supported privacy legislation, baseline privacy legislation, as well.

We don't know what the position of the incoming administration is going to be on that. We'll have to see. But there has been an ongoing debate about that, and many bills in Congress, the data security bills have advanced much further. There's bipartisan support for data security. But for whatever reason, a law has not passed. So this is not something that is not being looked at. And meanwhile, the states are also very active in enforcing. And the FTC has brought hundreds and hundreds of cases under our unfair and deceptive--

MARK WEINSTEIN: Yeah you guys have been great--

JESSICA L. RICH: And, enforcement is significant. So, you know, Europe has many laws. But we, despite not having a baseline privacy law, have done very, very strong enforcement. So just to lay that out, and be clear. But I'll let you guys-- you were being attacked for saying it belongs in a Consumer Protection Agency, so maybe you would like to respond to that.

MARK WEINSTEIN: Well support it in the consumer age-- like the FTC. Like Chairman Ramirez is heroic.

HOWARD BEALES: Well, I mean there are clearly aspects of privacy that are none of the FTC's business, and that it doesn't deal with. And, you know, what you point to is constitutional privacy. I think that's absolutely right. And I think it's completely legitimate to be concerned about information that goes to the government. But the problem, and this is thinking about the

consequences, not the data, the problem is it's going to the government. It's not that the information exists, or that it was collected somewhere. And I just disagree completely with the premise about your grandparents, because I don't know where your grandparents grew up, but mine grew up in a small town, and everybody knew everything. There was no privacy.

MARK WEINSTEIN: In the privacy of their home, there was privacy. There wasn't a camera.

DIERDRE K. MULLIGAN: On the--

MARK WEINSTEIN: You know, there wasn't-- go ahead, sorry.

DIERDRE K. MULLIGAN: Sorry. My prediction on what you're going to see is I think there will be a lot of activity in the states. California has been a leader in legislative, and, you know AG actions, on the privacy front. And I think you will continue to see them on the government issues. One of the most interesting things that's happening at the kind of city level in California, Santa Clara, Oakland, up in Seattle, there have been the creation of privacy councils that are advisers to the cities, that are specifically looking at the acquisition of surveillance technology by police departments, and things like that. So I think there's a lot of activity to address privacy that will continue regardless of whether or not Congress acts. But certainly there will be continued pressure for the adoption of kind of baseline privacy regulations. And I think that would be a good thing. But I also think most of the ones that have been offered are not sufficient to address some of the concerns that we're going to see.

MARK WEINSTEIN: Thank you.

JESSICA L. RICH: Do we have time for one more question? We'll take one more question, and then we'll adjourn.

SPEAKER 7: Thanks. So, could [INAUDIBLE] of like, the limited scope of the consumer protection powers of the FTC? Do you think that using like a security, like regulatory procedure would be useful for-- like earlier, it was discussed, the D link settlement with the FTC, but that's only done under-- and only they're shipping insecure routers, and the only way the FTC was able to stop that was through a production rule policies. So basically, if they hadn't advertised routers as secure nobody would be able stop them. And that creates a problem for the internet as a whole. Like we thought that the DNS D-DOS thing, like a couple months ago, that took out a lot of websites. What is the prospect, especially with the initiation of using, of giving powers to like DHS, and stuff to regulate these device manufacturers? Or even maybe even doing, wrapping privacy into a security framework, to help have stronger regulatory powers to create a more healthy device environment?

JESSICA L. RICH: Well, I feel the need to say that if additional powers are granted to a government agency to deal with privacy and security, it should be the FTC, because we have a lot of expertise. And we have authority beyond if a company advertises. We have authority beyond that to protect privacy and security. So I just want to be clear on that. And D-Link was not a settlement. It's in litigation. So our authority will be tested.

ANDREW STIVERS: So let me just put in a plug for the kind of authority the FTC has, which I think is really valuable, and an incredibly, incredibly powerful. So one of the advantages of doing the way we do things is that we allow the innovation to flourish in the marketplace. Generally, the kinds of things that we're concerned about, consumer protection, aren't health and safety issues-- sometimes they are-- but they're things that we can often kind of re-address after the fact. So we can allow the marketplace to evolve, and then when we identify practices that we think are harmful to consumers, we can go after that, and send signals to the marketplace that say, hey this isn't right. Taking the kind of [INAUDIBLE] regulatory approach, that I'm familiar with from my work. It has a whole raft of other issues associated with it, and especially when technology is moving very quickly, it's sometimes difficult to craft a regulation that strikes the appropriate balance between allowing new uses of data, and making sure that there isn't any harm there. For food safety, which again was what I was involved with before, that technology is a little bit more stable. It's still evolving. But it's a little bit easier, I think, to regulate that. Plus, you know health and safety violations are a little bit harder to wind back, if you allow them to occur. So there's some differences in what we're looking at here.

HOWARD BEALES: And I guess the other--

JESSICA L. RICH: One comment form each and then-- I've gone 10 minutes over so--

HOWARD BEALES: I guess, the other difference I'd point to from food safety is the germs may evolve, but the threats on the internet, and the threats to hack into internet of things devices evolve a whole lot faster. The regulatory process is just not fast. I think the only way we could build a camera that we could guarantee would never get hacked is it didn't take pictures. No. Other than that, we can't rule out all possibilities because smart people are going to be trying to figure out ways around that if there's something that they can do with it. And that's inevitable.

JESSICA L. RICH: Dierdre, comment.

DIERDRE K. MULLIGAN: So the Federal Trade Commission has done yeoman's work on getting industries attention to the security of their products. You know, I believe that they will continue to do that as they've already started in the IOT space. And there is an education process that happens as companies realize that they have to be paying attention to known vulnerabilities, and making sure that they don't build them in, et cetera. That said, I would agree, if anybody is going to be given authority around security in the commercial marketplace it should definitely be FTC, and to the extent that DHS would get authority, it would be around critical infrastructure. But who knows what the next administration will do? But you know, I would certainly fully support giving the FTC greater clarity around their authority in the security space. But I think what they've done has been both fully supported by the law, and has been very effective and important.

JESSICA L. RICH: But we're working with all those other agencies, and we're trying to strengthen the whole approach to data security across the government. So I've kept everyone really too long. One plug I want to make is that although we have taken great delight in highlighting research the second year here, second PrivacyCon, we want to hear about research all year long. It helps really inform our work. And we also may be able to highlight it for the

public. And so we have a dedicated mailbox, research@ftc.gov, and we hope that you share with us and we monitor it regularly. Thank you so much, to our panelists, and to all the panelists today, and for all of you for coming. Happy new year.