

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR:

Hey, everyone. Good on that? All right, cool. So today we're going to talk about the economics of spam and security in general. And so up to this point in the class, we've mainly talked about the technical aspects of security. So we've looked at things like buffer overflows, the same-origin policy, Tor, and all kinds of things like that.

And so the context for that discussion was that we were looking at how an adversary can compromise a system. We tried to devise a threat model that would describe the types of things we want to prevent, and then we tried to think about how we could design systems that would help us to defend against that threat model.

So today we're going to look at an altered perspective. And the perspective that we'll look at today is, why is the attacker trying to compromise your system? Why is the attacker trying to do these evil things to us?

And so there's a bunch of the reasons you can imagine why attackers might be trying to do these evil things. So some of these attacks are done for ideological reasons. So think about people who perceive themselves to be political activists, or things like that. Or if you think about Stuxnet, for example. Sometimes it's like governments attacking other governments.

And so for these types of attacks money, economics, is not the primary motivation for the attack to take place. And what's interesting is that it's actually hard to make these attacks go away, other than generically making computers more secure. There's not really some financial thumbscrew you can turn to make these attackers disincentivized to do things.

However, there are some types of attacks that do involve a strong economic component, and those are some of the things we're going to look at today. One of the interesting things, though, is that for a lot of these attacks that don't have an economic component, in that we can't use regulations and things like that to try and prevent them. It can sometimes be difficult to figure out how we'd be able to stop them at all beyond, like I said, just trying to make computers more secure.

For example, Stuxnet's a great idea. So this is the malware that was attacking some of the industrial software in Iran, with the centrifuges. So we all kind of know where Stuxnet came from, right? We basically know it was the Americans and the Israelis. Basically. But can we prove that in a court of law? Like, who can we sue, to say You put Stuxnet on our machine?

So it becomes a little bit murky when you have some of these attacks, where it's not clear you can sue the Federal Reserve, or you can sue Israel, for something like this. And furthermore, no one's gone on the record as officially claiming that it was them. So there's some very interesting legal and financial issues that get involved when you look at how to prevent these attacks.

So there are many kinds of computer crime that are driven by economic motivations. So for example, state-sponsored industrial espionage, for instance. So this is one thing that some of our previous speakers have talked about. Sometimes governments try to hack into other governments or other industries to steal intellectual property, or things like that.

And what's interesting is that, like the attacks that we'll look at today, which are spam, you'll see that actually take some money to make some money. Spammers actually have to invest in an infrastructure before they can actually send these messages out.

And so if you have these attacks where it takes money to make money, and you can figure out what that financial sort of tool chain looks like, then maybe you can think about applying upstream financial pressure to stop that downstream malware attacks or security problems. And so I think the take-home point is that if we look at the context of spam in particular, spammers will stop sending spam if it becomes unprofitable.

One of the sad truths of the world that we continue to get spam messages because it's cheap for them to send them, and 2% to 3% of our fellow human beings will actually click on links and look at stuff. And so as long as these costs for sending these messages out are so low, then even if the hit rates are low, people can still make money off that kind of stuff.

So for today we're going to look at attacks that have a significant economic component to them. And so one interesting example which I actually just read about takes place in China.

And so in China they have this problem with what they call text message cars. So the basic idea here is that people drive around with these cars that have these radio antennas attached to the side. And they can essentially do-- think of it almost like a man in the middle between

people's mobile cell phones and the actual cellphone tower. And so they can basically run around in these troll cars, and they can get all of these cell phone numbers, and then use that car to send spam messages directly to the numbers that they've collected using this sort of vehicle take.

So these text message cars can actually send upward of 200,000 messages a day, which is an incredibly high number. And the cost of labor over there is actually very cheap. So it's very inexpensive to hire a driver, drive around one of these cars, and just snoop on people's traffic and send them spam.

So let's look at the economics of this. So what is the cost of the evil antenna, this thing that allows people to take these messages off the air? Roughly speaking, it's somewhere in the order of about 1600 bucks, give or take.

So how much profit can these people make a day? So in a hilarious coincidence, this is also roughly 1600 dollars. So this is very interesting. What this means is that once you buy one of these things, then in a day essentially you've made back your money. So that's great, from the perspective of being a spammer.

Now you might say, OK, but you might get caught by the police and then you might get put in jail or have to pay a fine. So in the case of the fines, the fines for getting caught are less than 5K. And people rarely get caught.

And so these are the types of calculations we have to look at when we're trying to think about how to economically deter these spammers. Because if these spammers only get caught a couple times a year, and they basically make back their hardware costs in a single day, it's very tricky to figure out how we can use financial disincentives to make them stop doing this kind of stuff.

And what's interesting is that in China the mobile carriers are also somewhat implicit in this scheme. So every time you send a spam, you're going to send some small amount of money to the mobile carrier, right? A couple cents. It works that way over here as well.

Now over here in Europe in many cases, the mobile carriers have decided that they don't want angry customers contacting them saying, I'm getting hit by these spam messages all the time. But apparently a lot of the Chinese mobile carriers, at least the top three ones, they're actually seeing these spam messages as a source of revenue. They actually think this is a nice way for

them to get some free money.

So in fact these telcos have set up these things called 106 prefix numbers. I don't know if you've heard of these before.

[BANGING]

But the original-- there's apparently a ghost in the room.

The original purpose of these numbers was to do things for non-commercial reasons. For example, imagine that you run a company, and you want to send a bunch of text messages to all of your employees. You can use one of these 106 numbers, and you would basically be able to send things in bulk. You'd be able to avoid some of the built-in rate-limiting mechanisms they had in the cell network.

So there's this nice thing sitting around that spammers can actually use. And so as it turns out, I think it's something like 55% of the mobile span that gets sent in China comes from one of these 106 numbers. So this is a really interesting case study of how these financial numbers work out, and how sometimes you can actually have these sort of perverse incentives, where in this case the cellphone carriers are just going along with these scams and these schemes. And there'll be a link in the lecture notes. There's an interesting Economist article about this.

[BANGING CONTINUES]

There is like a pan-African drum circle back there. This is super exciting, though. I like it. I am being adversarially attacked. That's OK. We will play through the pain. Perhaps this is the Mossad. They don't want me to talk about Stuxnet.

Another interesting thing about security is that there are actually many companies that deal in cyber arms. So this is kind of something out of G.I. Joe, but there are actually these companies that will sit around and they will actually sell you malware, they will sell you exploits, they will sell you things like this.

So one example is this company that's called Endgame. And so for example for about \$1.5 million, Endgame will give you IP addresses and the physical locations of millions of unpatched machines. So they have sort of vantage points all over the internet, and they know all kinds of interesting information about machines that you may or may not want to attack if, for example, you're a government, or if you're another agency or something like that.

For about \$2.5 million, they will give you what is delightfully called a zero-day subscription package. And so if you sign up for this, then basically you will get 25 exploits a year, they claim, for that much money.

And so you'll get those exploits in your inbox or whatever. Once again, you can do with these things whatever you want. You've clearly got 2.5 million dollars, so you've got a lot of spare time to think about this stuff, presumably.

And so what's interesting is that a lot of people who work in these cyber arms dealers, they're actually ex three-letter agencies. They're ex-CIA, or ex-NSA, or things like this.

It's interesting to think about who are the actual customers of these cyber arms dealers. Some of them are actually governments, like the American government, for example. And they use these things to attack other nations, or whatever.

But some of the people who buy this stuff are actually, increasingly, companies. So one thing we'll talk about a little bit at the end of the lecture is how sometimes companies are now taking cybersecurity into their own hands and sometimes doing what's called hackbacks. So without getting the government involved, companies that are attacked by cybercriminals will sometimes go back and explicitly try to take out people who tried to steal their intellectual property. And they've used some very inventive legal arguments to justify this, and so far it's actually been fairly successful. So this is an interesting aspect of cyber warfare.

AUDIENCE: How is any of that legal?

PROFESSOR: Well, so. I mean, information wants to be free, dude. Right? So if you think about stuff like this, for example. Just telling you stuff isn't necessarily illegal. I mean, it gets a little bit gray.

But for example, if I tell you that look over there, there's a house, and the lock doesn't work on that door. Can I have 20 bucks? That's not necessarily illegal.

Because as it turns out, these companies have, like, hordes of lawyers that look into things like this. But in many cases, if you think about it, you can search for stuff on the internet and go to websites that tell you things like how to build bombs, for example. Just posting that information typically is not illegal, because you're just learning. What if I'm a chemist, for example? Or something like this. So a lot of times, just giving someone knowledge is not necessarily illegal. But you're right that there's some gray areas here, and as we'll talk about with some of these

hackbacks, it's not always clear.

For example, if I am a bank, I'm not a government, I'm a bank. I get hacked. It's not always clear that I actually have the legal authority to go back and, let's say, try to shut down a botnet or things like that.

Companies have done stuff like that. But I think this is an example where the law is lagging behind practice. And so people have used things like, we will use copyright infringement law to attack botnets as a company. Because they're selling legal goods of ours, so we'll use IP infringement. Like, this is probably not what Thomas Jefferson was thinking when he was thinking about how these laws work. So this is a little bit of a cat-and-mouse game. So we'll do a little bit of that later in the lecture.

So, yes, this is very interesting. Basically what this all means is that there's this marketplace for all kinds of computational resources that you might use as someone who wants to launch attacks. So for example, there's a marketplace for compromised systems. So, for example, you can go to the darker places of the internet, you can purchase entire compromised machines that might be part of a botnet. You can actually buy access to a compromised website, for example. You might use that website to post spam, or put up evil links, or things like that.

You can also get access to compromised email accounts, like Gmail or Yahoo accounts. As we'll talk later, those things are very very powerful for an attacker. And you may also just buy sort of a subscription service for a botnet. You'll just have this thing lying around. You can use it to send denial of service attacks or things like that. So there's a marketplace for that.

There's a marketplace for tools. So you can get, as an attacker, off-the-shelf malware kits, for example. You can use perhaps arms dealers like this to get access to zero-day exploits so you can write your own malware, so on and so forth.

And there's also a big marketplace for stolen user information. So this is stuff like Social Security numbers, credit card numbers, email addresses, so on and so forth. So it's all out there on the internet if you're just willing to look for it.

And so the paper that we're going to look at today basically focused on one aspect of this, which is the spam ecosystem. And so in particular, they look at the sale of pharmaceuticals, of knockoff goods, and software. And so they basically break this spam ecosystem into three

parts.

They break it into advertising. So this is the process of somehow getting a user to click on a spam link somehow. And then once they've done that, there's this issue of click support. So this is the notion that once the user clicks the link, there has to be some type of web server, DNS infrastructure, so on and so forth on the back end that actually presents the spam website that the user goes to.

And then the final part is realization. So this is actually allowing the user to say they want to buy something. The user sends money to the spammers, and the user's going to get some product back in the back end. And so this is where all of the money makes place.

And so a lot of this stuff is actually outsourced to what the paper calls affiliate programs. And so you can think of these affiliate programs as essentially doing a lot of the back-end grunt work of talking to banks and Visa and MasterCard and things like this. And so a lot of times, the spammers, they don't want to deal with that stuff. They just want to create the links and do-- you can think of it as the advertising component. And so a lot of times the spammers themselves, they will work on a commission. So they will get, let's say, anywhere between 30% and maybe 50% of the final sale that they deliver to one of these back-end affiliates.

So does that all make sense at a high level? OK. So what we'll do is we'll look at each component of this spam trajectory, and then see how it works, and then maybe think about how we'd to be able to shut down spammers at different levels of this [INAUDIBLE].

So the first thing we'll look at is the advertising component. And so, like I mentioned, the basic idea of the advertising is, how do you get the user to click on a link? That's the primary question we'll be concerned with here.

And so the typical thing, as we all know, is you're going to email spam, although as we discussed at the beginning of lecture, people are starting to use text messages and some of these other forms of communication. You could also imagine maybe here we're going to start using social networks as well. So now when you go to Facebook, not only are you polluted by your real friends' content, you're also polluted by spam messages too.

So this is about economics, this discussion. So one interesting question is, how much does it cost to actually send out these spam messages. And so as it turns out, it's not very expensive at all. For about 60 bucks, you can spend a million spam messages.

So that's a super, super low cost. And this cost is actually much lower if you're directly operating a botnet. You can cut out the middleman. But even if you are renting one of the botnets from one of these marketplaces, this is still super, super low.

AUDIENCE: So how many of those are actually effective? As in, they don't get filtered?

PROFESSOR: Ah, so that's a good question. So that leads to my next point. So you're sending a million spams, but then they're going to get dropped at various points along the way. They're going to get caught in spam filters, people will-- they see it but they just delete it because they know that an email that has, like, 18 dollar signs should just be deleted.

So if you look at the conversion rate, you'll see that the click rates are actually very low because of things like spam filters and stuff like that. And also many users are trained to avoid these things. Click rates are low.

And this is why sending spam has to be super, super cheap, because you will not get a lot of conversions. So for example, there have been some empirical studies that looked at these click rates. And one study found that they looked at 350 million spam messages, and they found that out of those 350 million messages, there was only about 10,000 clicks on those messages.

So there's a massive dropoff here. And then out of these 10,000 clicks there were only 28 purchase attempts. So that's super, super low. And so that's why it's extremely important for this entire ecosystem to be very cheap from the perspective of a spammer. Because I mean, look at these dropoffs here. These are multiple orders of magnitude.

And so that's why one might hope that at least in theory we could squeeze-- like for example, we could drive this number up maybe just \$10. Maybe that has some catastrophic knockdown effect on how profitable this stuff is. So it's very important for the spammers that everything be as cheap as possible.

AUDIENCE: So those 10,000 clicks. Again, how many of those 350 million emails were filtered out of the inbox? I'm just trying to get a sense of out of how many emails those clicks were out of, to gauge how effective spam filtering is versus how silly us humans are.

PROFESSOR: Yeah, that I'm not actually sure. That's a good question.

AUDIENCE: So I was just listening to a talk by Jeff Walker on Friday about this stuff, and he says that on

the order of 20% to 40% of clicks going to one of these websites actually goes from a user's spam folder. So users go in their spam folder, looking for this stuff, and they click on it. So presumably there's a class of customers that are looking for this, and if they're looking for it-- oh, yeah, I'll just go into my spam folder to find this. So it's not clear that things going into spam folders are getting zero clicks.

PROFESSOR: Yeah, I've heard anecdotal reports of that too. Some people, even for legitimate emails, they'll mark it as spam just so that if there's a shoulder-surfer, like at work, who's seeing them go to Gmail, let's say, they won't come and see that you've subscribed to, you know, whatever. And then they can secretly go into the spam folder, they know it's not deleted, and look at this stuff.

This is actually a really interesting point. There's this whole psychology of who it is that actually clicks on these links. And so I think one of the papers that I linked to in the lecture notes talks about why these Nigerian scams still work. Because you'd think that anyone who basically has either common sense themselves, or a friend who has common sense, would never click on one of these Nigerian email scams. Right?

But it turns out that the Nigerian meme is actually useful for spammers to filter out idiots. In other words, if you are so foolish that you would still click on a Nigerian email, then oh, OK, you're going to do one of these conversion things here.

When you think about it, that's one of the key things that spammers need. They need people who are gullible enough or idealistic enough to click through on these things. There's a whole sort of psychology behind this. It's very interesting.

AUDIENCE: So each of these purchases, about how much are they worth?

PROFESSOR: That's a good question. So it actually depends on the type of thing that you're looking at. A lot of these purchases are not actually super high in value. So you're thinking that someone's buying herbal Viagra or they're buying like a knockoff Windows license or things like that.

And in fact, a lot of times when they're buying these knockoff products, presumably the price is lower than what they'd actually get in the real market, because otherwise you could just go down to your local mall and buy these things. So a lot of times these purchases you're actually making are less than 1,000 dollars, and oftentimes a lot less than that. Any other questions?

OK. So these conversion rates are super, super low. So like I said, one of the key things to do

as a defender is to try to basically make spam more expensive for the spammer. So there's a couple different ways you might think about doing that.

One way you might think about doing that are IP blacklists. So maybe ISPs or someone else basically collects this list of IPS that are known to be bad, that are known to come from spammers. And then we just don't let these people send traffic.

So this kinda-sorta used to work for a while. But now it's so much easier for the attackers to use techniques like DNS redirection and stuff like that, that we'll talk about in a little bit, this doesn't actually work out very well. Because now there's a much larger set of addresses that spammers can send spam from, and they can also dynamically switch the binding between hostnames and web servers and all these types of things So this doesn't work out so well.

Another idea that's been around for a long time is charging for email in some way, so each email you send, you have to pay some micropayment. So that currency could be a couple different things. So you might imagine that if I wanted to send you an email, maybe I'd have to pay a tenth of a tenth of a penny. And that's no big deal for me, because I don't send that many emails a day. But if you're a spammer trying to operate at these volumes, then that quickly adds up. That destroys their value chain.

Another idea that people have had is, what if you used computation as a currency? This is the idea that before my email server will accept an email from me, I have to solve some puzzle. I have to do some math trick, or something like that. Once again, that cuts down the rate at which these bulk mailers can send messages.

Also, we're all familiar with CAPTCHAs, too. This is basically the idea that I have to look at some picture of nine animals and find the cat instead of the dog, or type in some weird squiggly number that looks like a migraine, or something like that. So there have been all kinds of ideas for charging for email to stop this kind of stuff from happening.

One of the classic problems, though, with all these schemes, is who's going to be the first one to implement it. And if all the email providers don't move forward at the same time, then of course spammers are just going to migrate to the email providers that don't require these techniques. So there's been the problem of how do we get everyone to upgrade en masse.

And there's this issue of, well, what would happen if a user device is compromised? So maybe if someone breaks into my Gmail account, then maybe they're going to force me to pay 350

million micropayments, which could individually bankrupt me.

And so it's not quite clear that some of these schemes are ready for primetime, but they do represent an interesting thought experiment about how you might be able to stop some of this stuff from the senders' side.

AUDIENCE: So how do they work with mailing lists, where you have these big mailing lists?

PROFESSOR: Yeah, so there's problems with that, and with mailing list aggregation. So it's very, very tricky, because there are actually some bulk mails that you do want to send.

I mean, you might imagine having some heuristic where you look at the size of the mailing list and maybe you scale the payment according to that. So for example, maybe heuristically you think it's reasonable to send email to 1000 folks but not to 350 million folks, or something like this. But you're right that there are a lot of practical limitation issues that come out with this kind of stuff.

So what the adversary can do to get around some of this? There are basically three workarounds that adversaries might try.

So one thing they can do is just use botnets, because botnets have a lot of IPs that the attacker can use. And so for example, even if someone were trying to do something like IP blacklists, then maybe the attacker can cycle through a bunch of IPs in this botnet and maybe get around some of that blacklist filtering.

They can also try to use compromised webmail accounts to send spam. So the reason why these are super useful is because sites like Gmail or Yahoo or Hotmail, those services can't be blacklisted, because they're super, super powerful. So if you blacklisted the entire service, then you're probably going to shut down service for tens of millions of people.

Now of course, these individual services can shut down you. And so that will actually happen once they have these heuristics running that see that you're sending to a lot of people you've never sent before, and so on and so forth. A lot of AI strategy takes place on the webmail server side to try to predict these things.

But these things can be very valuable to an attacker because even if your compromised account is not used to send a lot of emails, it can be used to send emails to people that you know. So maybe it allows the attacker to do things like spearfishing more easily, or things like

that. People are more likely to click on an email that comes from an address that they recognize. So that's a very powerful technique there.

And then attackers can also try to do things like hijack IP addresses from legitimate owners. So as was mentioned briefly in Mark's talk, there's this protocol called BGP that basically is used to control routing on the internet. So there are these attacks that people can do whereby they will essentially say, hey, I'm actually the owner of some prefix of IP addresses, even though they don't actually own it. So all the traffic that's involving those addresses will go in towards the attacker, and then they can actually use those addresses to send out spam from there.

Then once they're done with their evil, they can release the BGP advertisement and then go try to do this somewhere else. There's a lot of research in how you can essentially think of ways to authenticate BGP by advertisement or otherwise prevent these IP address hijacks.

So there's a bunch of different techniques that attackers can do to try to get around some of these defensive techniques. So this can all be done, but still, these defenses, they're not free. So presumably the attacker has to pay for the botnet somehow, they have to get inside these webmail accounts. And so any of these defenses that you can do will help to drive the cost up of generating these spams. So as such, they're still useful, even though they are not perfect defenses.

So what do these botnets look like? So at a high level, you have the proverbial cloud from your cloud diagram. You have your command and control infrastructure up here, and this is the thing that actually sends commands to all of the individual bots down here. So the spammer will talk to the C&C and will say hey, here's my new spam messages I want to send, and then maybe these bots will act on behalf of their command and control infrastructure and start sending emails to a bunch of people.

So let's see here. So why are these bots useful? Well, as I mentioned here, they have IP addresses, which are super useful. But of course they also have the associated bandwidth there. They also have computational cycles. Sometimes these bots are actually used as web servers themselves. So these things are very, very useful.

And they also serve as a layer of indirection. So, as we're to discuss in more detail in a second, indirection is very useful for attackers. That means that if law enforcement or whatnot shuts down this level, well, if the command and control infrastructure's still alive, then maybe

the spammer can just attach this command and control infrastructure to a different set of bots and keep on running.

So that's one reason why these bots are very useful. And these bots can scale to the order of magnitude of millions of IP addresses. So as it turns out, people will click random links involving malware all the time. So these things can get very, very, very large. And so some of these takedowns that these companies get involved in, with trying to take down these botnets, they involve millions upon millions of machines. So they're very technically challenging.

So how much does it cost to get your malware installed on all these bots? Remember, these are all typically regular end-user machines. So the cost for getting your malware on one of these machines, so price per post, is about \$0.10 for U.S. hosts and on the order of \$0.01 for posts in Asia.

So it's interesting there's this differential here. There might a couple of different reasons we can imagine for why that is. It might be that people are prone to think that connections originating from the U.S. are more likely to be trustworthy. It may also be that because there's pirated software running here, stuff that's not actively up to date with respect to patches. It's actually easier to get botnet posts over here.

So you'll see some very interesting statistics about how some of these rates might fluctuate, for example, as you see companies like Microsoft go out and try to stamp down on piracy and things like that. But anyway, this is a rough estimate. Suffice it to say, this is not super expensive.

So what does-- any questions before we continue? OK. So what does this command and control infrastructure look like? So you can imagine that in one substantiation, the simplest substantiation, this is just some centralized setup. And so this is maybe one machine or maybe some small number of machines. The attacker gets to log into those machines and essentially just send these commands out to the botnets from there.

So if it's going to be centralized, then it's going to be very useful for the attacker to have what's known as bulletproof hosting. So the idea behind bulletproof hosting is that you want to put this command and control infrastructure on servers that reside in ISPs that ignore requests from banks or from law enforcement to take down servers.

So there are actually bulletproof servers that exist. They charge a premium, because there is

a little bit of risk involved there. But if you can manage to host one of your command and control centers there, it's going to be very nice. Because then when the American government or when Goldman Sachs or whoever says hey, shut this guy down, they're running spam, the provider will say, how can you make me? I run in a different legal jurisdiction. I don't have to follow your intellectual property laws. So on and so forth.

So this is very useful. Like I said, these types of hosts actually charge a risk premium for running that kind of service.

And so the other alternative for running the C&C infrastructure is, this could be a peer-to-peer network. And so the idea here is that maybe this is sort of-- you can almost think of it as a mini-botnet up there too. So the entire control infrastructure is spread across many different machines, and maybe at any given time there's a different machine that's responsible for sending commands to all of these worker nodes down here.

And so this is nice, because it doesn't require you to have access to one of these bulletproof hosts. You can construct the C&C infrastructure using regular bots. The P2P aspect of it makes it a little more difficult to provide guarantees about the availability of the hosts that are up here, but it does have some other nice advantages. At a high level, those are the two approaches that people can use.

So what happens if the hosting service gets taken down? Well, there's a couple things that the adversary can do. So they can use DNS to essentially redirect requests.

So let's say that someone attacks, or someone issues a takedown for the DNS infrastructure for something like this. As long as the back-end servers are still alive, what the attacker can do is basically-- the attacker creates lists of server IP addresses. And there may be hundreds or thousands of these IP addresses that it collects.

And then it will bind each one to a host name for a very short period of time. So let's say maybe for 300 seconds. And so what's nice about this is that if someone's trying to run heuristics that say, if I see some particular server sending more than 1,000 spam-like messages in a given period I'm going to try to issue some kind of takedown to them, well, these types of techniques will maybe help the attacker fly under the radar of those types of detection techniques. Because essentially every 300 seconds they're saying, OK, I'm going to be serving spam from here, then I'm going to be serving spam from here, serving spam from here, so on and so forth. So this is a nice use of indirection, at least from the attacker's

perspective.

And so, as I mentioned earlier, these types of indirection are one of the key ways that attackers try to evade law enforcement and these detection heuristics. So you might think about, well, what if we just take down the DNS server? How hard is it to do that?

Well, as the paper describes, there are a couple different layers on which you can attack these spammers. So you can try to take down the attacker's domain registration. That's basically the thing that says, like, hey, if you're looking for russianpharma.rx.biz.org, then here's the DNS server that you talk to. You can imagine attacking it at that level.

You could also imagine attacking it at the level of taking down the spammer's DNS server, the thing to which you'll be redirected once you look at that top-level domain. And so what's tricky is that the attacker can use these sort of fast flux techniques at every different level.

So, for example, they can rotate the servers they use to act as their DNS servers. They can rotate the web servers they use to send out the spam. And so on and so forth. So that's just a high-level review of how people can use multiple machines to try to avoid detection.

So as I mentioned earlier, you can use compromised webmail accounts to send spam. And the power of that is that if you can get access to someone's account, then you don't actually have to install malware on their machine. You can actually access their account from the privacy of your own machine, wherever it is that you're located.

And as we were discussing earlier, this is useful for spearfishing attacks, because you can send this spam message as the person whose account it actually belongs to. And so as a result the webmail providers are very motivated to shut this kind of thing down. Because if they don't do that, then they risk being blacklisted as a whole. All the users risk being flagged as spam, which they don't want.

And also the provider actually needs to somehow monetize their service. They actually need real users to be doing things like clicking on ads in the righthand bar of their webmail account. So the higher the proportion of their users which are spamming, the less likely advertisers are to advertise in their webmail system. So the webmail account providers are very incentivized to shut down this kind of stuff.

So how do they try to detect this type of spam? They use those heuristics. They might try to

use CAPTCHAs. If they suspect that you've sent some spam-like messages, let's say five times in a row, they might ask you to type in one of those fuzzy letters or whatever.

Suffice it to say, though, a lot of these techniques don't work very well. If you look at the price per account, so how much you as a spammer would have to pay to get one of these things, it's still super, super cheap. So it's on the order of \$0.01 to \$0.05 for an account on Yahoo, Gmail, Hotmail, something like that. So once again, this is very, very low. And so this does not act as an effective disincentive for spammers to try to do these types of things.

So this maybe is a little bit disappointing, because it seems like everywhere we go, we have to solve these CAPTCHAs if we want to buy things or send emails or do that kind of stuff. So basically, what happened to CAPTCHAs? They were supposed to make all this bad stuff go away.

And as it turns out, the attacker can build services to solve CAPTCHAs. So this can be automated, just like anything else. As it turns out, the economics for this is that if you want to solve one CAPTCHA, then it's approximately \$0.001 dollar to solve a CAPTCHA. Which is nothing. And this can be done with very, very low latency, too.

So CAPTCHAs essentially are not presenting most large-scale spammers with a high barrier for sending these spams. And so how is this being done? If it's this cheap, you might think, maybe it's being done all by computers, by software. But it's not, actually.

So a lot of this is done by humans. In particular, the attacker can outsource this in one of two ways.

So first of all the attacker can just find a labor market where the cost of labor is very, very cheap. So you can employ humans to essentially act as CAPTCHA solvers for you. You, the spammer, are presented with a CAPTCHA by Gmail or whatever. You, the spammer, then send that CAPTCHA over to some human sitting somewhere. They solve for you, they've earned some small amount of money, and then you send their answer to the legitimate site.

You could also do this with Mechanical Turk. Have you guys heard of Mechanical Turk? I've asked the question, my back is turned, [INAUDIBLE].

OK, so Mechanical Turk is pretty neat, I mean neat if you're trying to do evil. So what's nice about that is that you can post these tasks on Mechanical Turk and say, hey, I have a picture-solving game, or something like this. Or you can just come out and say straight up, I've got

some CAPTCHAs I want to solve. You post a price, and then basically the market will match you with people who are willing to do that task. And then they'll do it for you, they'll post the answers.

So this actually automates a lot of actually finding the labor pool for the spammer. The problem with this is that you have more overhead for the spammer, because Amazon has to take some cut of that profit that's generated from that. But that's very nice there.

Another thing that attackers can do is they can actually reuse CAPTCHAs on legitimate sites. So there's some CAPTCHA that the attacker wants to solve. They then have some legitimate site on the side where they present that exact same CAPTCHA, and get a real visitor to figure out what that CAPTCHA is. Then they come back over to the first site and then use that answer as the answer.

And like all these crowdsourcing-type things, if you don't trust your users, then you can maybe replicate the work. So you send the CAPTCHA to maybe two or three people. And then you come back in and use majority voting, take whatever that majority vote was as your CAPTCHA answer. And so these are some of the reasons why the CAPTCHA defenses don't work as well as you might think.

So the providers, so for example Gmail or Yahoo or whatever, can try to implement more frequent CAPTCHAs to try to push the friction level up for the spammer. The problem there is that then regular users will get irritated.

So a good example of this is Gmail's two-factor authentication. It's actually a super good idea. Whenever Gmail will detect that you're trying to use Gmail from a machine that it doesn't know about, it'll basically send you a text message saying hey, enter this verification code into Gmail before you can actually continue to use the service.

And so what's funny is that it's a super great idea, but at least for me, I get super irritated when I have to get that text message. Like, I know it's good for me, but I just get angry. It's frictionful. And so I'll do it if I don't migrate to a lot of different machines a lot, but if I had to do it any more than I did right now, it's unclear that I'd feel as happy about it as I do.

So there's this very interesting sort of tradeoff between the security that people say that they want and the security measures that they're willing to put up with. So as a result, it's very difficult for the webmail providers to increase the amount of CAPTCHAs and still keep users

happy. OK, so any other questions before we move on to click support?

AUDIENCE: So is one of the reasons for the non-adoption of encrypted emails, besides the [INAUDIBLE] is that spam filters have a very, very big part?

PROFESSOR: Ah, because then they can't inspect messages and see what's going on. That's a good question. I think it's actually hard to say. I don't know, because it's a little bit of a chicken and egg problem.

So because there isn't a huge volume of encrypted email, it's unclear whether spammers are actually trying to take advantage of that. But I could see that maybe being a problem. I mean, people have looked at ways to do computation over encrypted data. So maybe you could think about doing something there. But it's always tricky.

So for example, with spam, people have these spam filters that were based on Markov models and things like that. So what do the spammers do? They start making these images that basically can't be seen by the text scanners, but then have the spamming content in there. So it's always an arms race.

All right. So let's move on to click support. So what is this about? So once the advertising step has succeeded and the user is given a link, so these are clicks on that link, so the user contacts some DNS server after clicking on that link to basically translate some hostname that was in that link to some IP. And then after that translation takes place, the user has to contact some web server that has that IP.

So to make all this work, the spammer has to register a domain name. And then the spammer has to run a DNS server, and then they have to run a web server. So this is essentially what the spammer has to do to make this click support thing work out.

So one question you might have is, well, why wouldn't the spammer just use raw IP addresses, for example, like in these spam URLs? And so does anyone have any thoughts about that? Why wouldn't you just have 183.4.4 dot whatever, instead of having something like russianjewels.biz?

AUDIENCE: Because it looks sketchy, it makes it easier to tell.

PROFESSOR: Yeah. So one thing, one would hope, is that a user would look at this thing that just has a bunch of numbers in it, and they'd say, well, this clearly seems weird. As it turns out, this will

only weed out some of the users, but you're exactly right. There's a subset of people you would lose just because nobody wants to click on that.

Another reason is that once again, having this sort of DNS infrastructure up here gives the attacker another level of indirection. So once again, if the legal authorities or whoever shut down the DNS infrastructure but they somehow don't manage to shut down that back-end web server, then the spammer can conjure up a different sort of front end for their service and maybe try to use that same web server on the back end. So that's another reason, I think, that people don't typically put these raw IP addresses in their spam URLs.

So another example of how this redirection comes into play-- how this indirection comes into play, sorry-- is that these spam URLs often point to redirection sites. And so these are sites like bit.ly, or things like that. And so in addition to things like bit.ly, you could also imagine that a compromised website can actually also act as a redirecter. You just put the appropriate HTML or JavaScript in there that when the user goes to that site, it's then going to redirect the user's browser to some other different site.

So once again, this useful because it provides that level of indirection. And it actually acts as a force multiplier, so you have a single spamming web server back end, but then you can name it using different things. And that will allow you to maybe confuse filters who have blacklisted, let's say, 10% of your URLs, but not the other 90% of them. So this is a very, very common technique.

And then another thing is that sometimes the spammers can use botnets as web servers or maybe as proxies, as DNS servers, and so and so forth. We mentioned this a little bit earlier, but this is another example of how the more machines you have as an attacker, the more defense that gives you. Because you can hide your evil amongst a watershed of machines.

All right. So in some cases, one of the things the paper talks about is these affiliate providers. These affiliate providers kind of act as evil clearinghouses. They will help to automate some of the tedium of interacting with the banks, and things like this, on behalf of you, the spammer.

So one thing you might wonder is, well, why can't the law enforcement just take down the affiliate providers? They seem kind of like a choke point. And the thing is that these affiliate providers are kind of like SPECTRE from the James Bond movies. They're very decentralized themselves. So it's very difficult to point to an affiliate provider at this particular machine, and we'll just shut down that particular machine. Oftentimes the affiliate providers are distributed

themselves.

So that means that it's actually pretty tricky for, let's say, the FBI, to just go to some affiliate program and say, thou shalt not do this anymore. Another interesting thing, too, is that the paper mentions that in many countries IP laws are different, for example. So the FBI may not be able to enforce intellectual properties that we have with other countries.

And also, according to the paper, in many of these spam forums, the spammers claim they are providing a useful, legitimate service to Western countries. They say that essentially, prices are too high for some of these things, in these Western countries, and that the fact that people are clicking on demand indicates there's a legitimate need to buy Windows copies that may be riddled with malware.

So a lot of times the spammers themselves don't feel that they're doing anything bad. And as we'll discuss a little bit later, the spammers do often actually give you the stuff that you've paid money for, which for me was one of the most surprising outcomes of the paper. And so we'll discuss why that is in a little bit.

So one thing that the paper talks about is various takedown strategies that you can imagine employing to try to stop a spammer. So one thing it talked about, they said that only a few number of registrars host domains for many affiliates. And so what that means is that most of these affiliate programs are-- there's sort of this one-to-one binding between affiliates and the registrars that are dealing with their domain name and infrastructure. It's very rare that you have a single domain name registrar who's going to be associated with a bunch of different affiliate programs.

So what that means is that in many cases there's not this, like, master decapitation strike you could launch, where you'd take out this particular registrar and then all of a sudden the entire spam infrastructure falls down. They found similar results for things like web servers. It's very rare that one ISP will actually host a ton of web servers for a ton of affiliate programs. This distributed nature, once again, makes it very difficult to say, if we just take out these three things then the whole ecosystem just crumbles.

So that's a little bit disappointing, because one would hope that there'd be one web server in Evildonia, where if we could just take down Evildonia, then people would stop sending us spam. That's actually not true. As we'll see later, though, that may be true to some extent at

the banking back end. And so maybe we can actually put the squeeze on there.

So anyway, I was alluding to earlier about this realization phase. So the realization phase is what happens after you, the user, have decided to buy something. So the realization phase consists of two parts.

The user pays for whatever goods they've bought, or they want to buy, and then the user hopefully will receive those goods. So either in the mail because they're buying some type of knockoff drug, or they get some software download because they want to get some fake version of Photoshop or something like that. And so the money flow looks something like this.

We start with the customer here, and they're going to tell the merchant hey, I want to go buy something. They will send some credit card info here, and then the merchant is going to talk to the payment processor. And this is essentially a middleman that helps the merchant, the spammer, deal with some of the intricacies of interacting with the credit card system.

The payment processor will talk to the acquiring bank. So the acquiring bank, that's the merchant's bank. And then the acquiring bank-- running out of space here. So, violating all good design standards, we will come up here.

So the acquiring bank is then going to talk to-- they call them in the paper the association network, but just think of this as Visa. This is the credit card network up here. And then finally the association network, Visa or MasterCard or whatever, talks to the issuing bank.

So that issuing bank is the customer's bank. And essentially the Visa or whoever is going to go to the customer's bank and say hey, is this a legit purchase? Is this a legit transaction? And if this is a legit transaction, then the money will actually flow through this entire system. So this is what the end-to-end financial workflow looks like.

And so this workflow can actually process a lot of money. So one of the papers that we mentioned in the lecture notes shows that a single affiliate can get more than \$10 million dollars at this workflow here. And so in practice, you might think that oh, why wouldn't the acquiring bank or the issuing bank say, something looks kind of fishy here? As it turns, in many cases, they don't.

And so this gets into this interesting discussion about why is it that these workflows are often tolerated by the financial system. For example, why do spammers properly classify their transactions?

So if you want to send something through this system, you have to tag that transaction with some type of type. You have to say, this is pharmaceuticals, this is software, this is whatever, this is whatever.

So you might think that as a spammer, you wouldn't actually want to do this. If you were selling fake Flintstones vitamins, maybe you don't want to say this is actually a pharmaceutical transaction. And what's interesting is that spammers do actually properly classify these transactions in many cases. And the reason is that there are high fines if you misclassify.

So essentially what happens is that these association networks like Visa or Mastercard, in many cases they are OK, perhaps, with transactions that are slightly shady. But they don't want to be blamed for being a money launderer, or for trying to deceive the authorities. So as long as you properly classify what you do, then in a certain sense this gives the association networks a little bit of, well, listen, they told us what was going on. Maybe the law was a little bit unclear. But we, at least, Visa or MasterCard, did not try to hide the intent of this transaction.

So spammers do oftentimes properly classify their transactions. So that's interesting. It seems like they're playing within the confines of the system a little bit.

So another question I mentioned earlier is, why send anything to users? Because presumably you're a spammer, so you're a criminal, right? So why wouldn't it just be cool if you just took people's money and then ran? I mean, that'd be the ultimate crime.

So as it turns out, they actually send things to users because, surprise surprise, high fines if they don't. So it's this very entertaining system whereby spammers kind of want to do things that are legal, when they actually can't use Bitcoins yet. They actually have to work within the constraints of this pre-existing system.

So as it turns out, there are these high fines if you, and by you I mean the spammer, have too many chargebacks. So a chargeback is essentially when a customer tells their credit card company, hey, I didn't get the thing that I was supposed to get that I bought with your credit card. Or I got it, but they didn't like it.

So if you're a spammer and you have too many customers saying things like this, then you will actually get charged very, very high fines. And as we saw earlier, the clickthrough rates for spam are super, super low. The conversion rates are super, super low. So even just one or

two fines might wipe out your entire profit for a month, let's say, for something like this. So spammers are really motivated to avoid these fines in both cases.

AUDIENCE: Would using Paypal obscure any of that, like the relationship with the bank?

PROFESSOR: Well, typically, yes and no. So you can think of those-- Paypal is in many respects very similar to Visa or MasterCard. So it has very similar regulations that oversee it, because it bears many of the same types of risks. I do think that Visa has slightly stricter restrictions on some of this stuff, as we'll talk about in a second. But for all intents and purposes, Paypal looks very similar.

AUDIENCE: Is there any sort of idea of having a group where you make some sort of account and then intentionally go to a bunch of spammers, buy a bunch of things, and then ask for a bunch of chargebacks whether or not they send it to you? So that they incur these fines. Or report them for misclassifying things, in order to just make them pay these fines.

PROFESSOR: That's interesting. It's like vigilantes.

AUDIENCE: Spam the spammers.

PROFESSOR: Yeah, exactly. I don't know if I've heard anything about that.

I do know that the spammers do try to detect people who are trolling them. So for example, one thing that they talked about in the paper a little bit is that spammers-- so how did the authors of the paper determine all this? They actually got a bunch of spam messages, they clicked on a bunch of stuff. They got a special Visa card they used to purchase this stuff, and then so on and so forth.

So spammers obviously don't like this. And so in the paper they call this test buys. Spammers want to prevent these test buys from researchers who are trying to figure out what's going on. So one thing that some spammers did-- do, I should say-- is they actually require proof of your identity before you can buy something. So they might ask you to send a picture of your photo ID, or something like that.

In particular, some people started doing this after Visa tightened up some of their rules about spam. Now, the problem with this is that most people who would click on spam apparently are still reluctant to send their photo ID to just some random person.

So there's a bunch of-- I've linked one of these articles in the lecture notes-- there's a bunch of

hilarious commentary from a spammer bulletin board, where they say oh no, Visa's cracking down on us. We try to ask for people's photo IDs, but they don't want to send it to us for some reason. And it's so weird that people wouldn't want to do that, but they will give them their credit card number. But anyway, so long story short, spammers are highly incentivized to try to detect that kind of stuff.

AUDIENCE: So for chargebacks, if you don't necessarily want your bank to know that you were buying these completely shady items, do a lot of users actually do chargebacks if they don't get the item? Or are they too embarrassed?

PROFESSOR: Yeah, that's a good question. I don't know what fraction of people are in the set of people who bought herbal Flintstones vitamins, were disappointed by herbal Flintstones vitamins, and then, yeah, told their bank-- but what's interesting, though, is that the bank has to know in the first place that they're going to this place, right, because the thing went through. So avoiding the chargeback, I don't think you're going to-- but by doing the chargeback, let me say, I don't think you'd reveal any extra information to the bank that they wouldn't already know. Because they had to clear the transaction first for you to actually get it and be disappointed.

AUDIENCE: So then roughly how many chargebacks is too much?

PROFESSOR: So some of the figures I've heard here are greater than 1%. So in other words, if you're a spammer and you have more than 1% of your transactions causing these problems, you get in trouble. And I wouldn't be surprised if it was a little bit lower than that, but 1% is the number that I've heard.

All right. So to me, like I said, this was one of the most interesting parts of the paper. Because I would have thought that a lot of spamming just involved straight-up fraud. That people clicked on links, they sent money, they never got anything. But as it turns out, because these spammers have to go through this network which has all these mechanisms to prevent fraud, they end up having to actually ship things over to users. So that's kind of neat.

And so another reason why spammers want to do these things, properly classify transactions and actually send things to users, is that only a few banks are actually willing to interact with spammers. And so what this means is that if the spammer is getting a lot of chargebacks, or getting in trouble with the bank or the credit card company or whatever, and some bank decides, I can't do business with you anymore, there's not a really large set of other banks that the spammer could go to to continue their chicanery.

So one study of this stuff found that there are basically only 30 acquiring banks that spammers were seen to use over some two-year period. That's actually not very high. So there is this other incentive to not be too goofy with the financial system, because you don't really have too many other places to go if you break those relationships.

So it seems like maybe this is a good choke point to try to cut down on spam. So we've already discussed how things like botnets give the attack a lot of IP addresses. There's a lot of different types of hosts who are willing to run web servers, so on and so forth. But this number actually seems small. So maybe we can actually attack spamming here.

But as I alluded to earlier, it's a little bit tricky to do this because of things like differing IP laws, because of things like the fact that it can be sort of tricky to actually say that spammers are doing something illegal. So if you are using spam messages to sell someone-- let's make this up, let's say sugar, sugar's delicious. It's not illegal to sell sugar, even at cut-rate prices. So even though the way that you may have drawn the user to that purchase was sort of duplicitous or gross, it is not in and of itself illegal to sell someone sugar.

And so as it turns out, a lot of spam sort of falls into this gray area, where the things that the spammers are doing are distasteful, but maybe not necessarily as illegal as you'd think. Now, for stuff like pirated software, there it's much more clear-cut.

But suffice it to say, it's not always the case that you can just point to one of these banks and say hey, your customers are criminals. Because that's not always true. Particularly if there's not a very strong paper trail that attaches the financial transaction to some spam URL that was the origin of the transaction. It's often very difficult to prove those types of links.

OK, so since this paper was published, the credit card networks have taken some actions. So this paper actually made a pretty big splash when it came out. And so the association networks like Visa and MasterCard and all of them were wondering, what can we do to cut down on some of this spam? So interestingly, after the paper came out, some pharmaceutical companies and software vendors actually lodged complaints with Visa.

So if you remember from the paper, Visa was the association network the researchers used to make these test buys, these dummy buys. So it's a little bit unfortunate, but that then showed some of these companies that hey, Visa can be used as the association network to fund some of this spam, or to translate some of this spam traffic. So some people complained about that.

So Visa made some policy changes in response to some of the issues that were brought up in the paper and some of the complaints that they got as a result. So now, for example, all pharmaceutical sales are now labeled by Visa as high-risk. So what this means is that if a bank acts as an acquirer for these high-risk transactions, then Visa will have some more stringent regulations they will put on that merchant-side bank. For example, they will require that bank to engage in a risk management program, and they may be audited more frequently, and so on and so forth. So Visa made that change.

And Visa also changed its operating guidelines. So its operating guidelines, now they explicitly enumerate and forbid illegal sales of drugs and trademark-enforcing goods. So the reason why they did this is that by tightening up this language, it is now easier for them to issue more aggressive fines against banks and merchants that they feel are doing things like selling illegal pharmaceuticals or selling knockoff versions of watches or things like that.

So once again, there's still a lot of spam that's in that gray area where it's not necessarily illegal. It's just that the customers were required to do certain techniques. And this is very useful because now Visa can drop some much bigger hammers on folks.

And as I mentioned before, some of the spammers tried to react to this by saying, well, let's just prevent these test buys. Because not only do security researchers do these test buys, but the association networks can do these test buys too. So they did some things like the photo ID type stuff, and that tended not to work out super well.

And so at least a few years after these changes were made, this did have an impact. I'm not sure what the latest state-of-the-art is with respect to trolling these Visa policy changes, but it was kind of cool to see this paper have this impact in real life.

So one interesting thing they mentioned in the paper is they talked about the ethical aspects of doing security research. And in particular, doing this research about the spam chain. To actually understand how some of this banking stuff worked, these researchers actually had to make purchases. They actually had to give money to people in exchange for these products.

And so in the paper they go through this kind of semi-hilarious defensive section where they say, we totally burned everything that we bought. We didn't use it. We talked to the companies whose pirated software we were buying before we got it.

But these things are actually pretty important to go through, particularly if you're within a

university setting. Because as you may know, if you want to do anything that involves-- particularly human research, but anything that might have these ethical sort of aspects to it, you have to get things cleared by lawyers, sometimes by an IRB, and things like that.

So it's actually pretty important for them to jump through these hoops, because at the end of the day they have to at least be somewhat confident that they weren't supporting some deeply nefarious activity in some far-flung corner of the world. So that was another interesting part of the paper, too. And other people have talked in this class about things like, what are the ethics of releasing zero-day exploits if you know they haven't been patched by someone? So it's a really interesting aspect of doing security research.

AUDIENCE: Is there any sort of oversight on security ethics? Because in the paper, they said the IRB wasn't interested.

PROFESSOR: Yeah, so that was super interesting. Yes. They said the IRB wasn't interested, I think, because there was no obvious human subject. But I think that at most universities, you couldn't just say, oh, there's no direct human subject, let me just go buy some stuff from somebody at the end of a spam link. And what they describe in the paper, actually in the acknowledgment section, they thank this whole set of people. Like, Sally at Legal, so-and-so at the Philosophers For Ethical Computing Association, and stuff like that.

I don't think there's actually a, how would you say it, an America-wide standard for doing this type of research. I know that each university's IRB has slightly different policies of what they do and do not allow, but I don't think there's a blanket policy.

AUDIENCE: Out of the 350 million spam URLs they tracked, of the 28 that actually responded, is there any chance that an appreciable number of those 28 spam responses were coming from researchers researching on spam?

PROFESSOR: Well, it's true that this type of calculus is actually one reason why I think the authors went to such lengths to defend themselves. Because if you think about it, the reason why those statistics are so hilarious is that it means that if you were to add five or remove five, that's the difference between a spammer being able to give their kids, like, a real gift versus a piece of coal. Because those numbers are so small.

So with regard to that particular [INAUDIBLE] that I gave you, I don't know how many of those were researchers. But I do think in general-- like I said, the spammers, they want to take your

money. And so if they could find some equilibrium whereby security researchers could do test buys, but that had no impact on their overall sales, they'd be fine with that. They just want the money.

But the tricky thing is that, let's say that-- let's make some number up-- half of those 35 were test buys, and that resulted in people putting pressure on the banks, and then instead of 35 they'd be getting two. That they don't want. So that's why they're so motivated to stop that stuff.

AUDIENCE: How much of this is blind emailing versus any sort of filtering? Because I'm sure they could run some models and get that 350 million down to, like, one page.

PROFESSOR: Yeah, so it's all about the cost-benefit analysis from the perspective of the spammer. So I think that you're right, and there are actually-- there's a marketplace for more targeted stuff. In particular, that's where some of those compromised email accounts can become very useful. But I think what you see is that people tend to go for the more focused stuff, like the more focused spam emails, for what they view as higher-reward targets.

So for example, political groups. People associated with the Dalai Lama, for instance. There, the perceived value of being able to get into that system is so high that people will spend the time to do this kind of stuff.

AUDIENCE: It would be interesting if there was one company dedicated to finding all the gullible grandmas and putting their emails into stuff.

PROFESSOR: Oh, interesting. I see. So basically having some database where it's like, totally send spam to this person, because--

AUDIENCE: It works.

PROFESSOR: I wouldn't be surprised if stuff like that existed, but I don't know if they do.

So one last thing that I wanted to mention is that, and I alluded to this a bit earlier in the lecture, that some companies have taken to doing these things they call hackbacks. So the idea is that, let's say that you're a bank, someone tries to break into your bank and steal your information. That bank will then, of their own volition, go back to those hackers and try to do something. Where something may be as quote-on-quote innocuous as shutting down the botnet, or maybe they try to steal their information back, and things like that.

This has actually become very much more common than it used to be. And one reason for this is that because the legal system has a little bit slow in adapting to some of these threats, some of these institutions, in particular software companies and banks, are tired of waiting for government-- like, their national government-- to deal with stuff.

So what ends up happening is that, for example, there was this big botnet in 2013 that was hosting all kinds of pirated goods and things like that. And so this huge coalition of Microsoft, American Express, Paypal, a bunch of them launched an operation to take down a botnet. They themselves took down the botnet.

They lurked around for a while, they learned about where the command and control infrastructure was. They actually went in there, took control of the command and control infrastructure, identified where all the end-user bots were. And they could send them messages saying, you need to patch your machine.

And so it's a very interesting area of intersection between security and the law. Because what part of American law, for example, gave those companies the right to do that? So what Microsoft lawyers said, at least, is that they said these botnets were violating Microsoft trademarks.

So for example, if you sell pirated goods, and you're saying this is Windows, for example, but it's not actually Windows or it didn't come from an official channel, then Microsoft says OK, you're violating our trademark. Therefore we can hack your botnet. It's a little interesting to see how that leap of logic took place. But the courts allowed it.

And this is increasingly happening more and more. And the banks in particular seem to be pretty upset about this, because there seems to be a lot of state-level sponsorship of some of these banking hacks. And the bankers care about the money, and so when they lose this money, they get very upset about that.

And so it's interesting to see how some of the burden for doing cyber security, in particular offensive operations, has now shifted a little bit more to the private sector. So it's not quite clear what the long-term implications are.

OK. That's the end of the lecture, and I guess we will see you on Wednesday and we'll go through the class projects.