General Discussion

Session 3

Mr. Anderson: Thank you. I will only speak briefly. We have heard some interesting points here, especially about the extra things people want to be able do, such as proving they have discharged their obligation or perhaps even having privacy of some kind against some types of government access issues.

Perhaps a good top-level way of looking at this is that systems engineering is all about managing complexity. Perhaps a third of big IT projects in industry fail and this is the same as it was in 1970.

Have we learned anything from 40 years' worth of studying software engineering and building ever more complex tools to manage complexity? No, we just build bigger, better disasters. You keep on rolling the stone up the complexity mountain, and a certain proportion of them fall off. So how you manage complexity is important. The evolutionary environment of your system also matters.

Now there are a couple of extremes here. One extreme is Odlyzko's Law, which says any system you can program eventually becomes so complex that it is unusable and you want to throw it at the wall in frustration. This happens and it does not matter whether it is a PC or a laptop or a phone or a computer game or whatever. And why? It is simple micro-economics, because whenever anybody suggests a new feature be added, the people who want the new feature are a concentrated and vocal interest, whereas the costs of this—the slightly increased probability of a blue screen of death—fall on everybody. So you end up getting complex and buggy machines for exactly the same reason we end up getting agricultural subsidies.

At the other end, Hal mentioned the Downton Abbey thing. This is actually very appropriate, because the goal of technology is often to enable the ordinary middle-class guy to live the way the upper class did a generation ago. When you think about it, we have laptops to do the jobs that were formerly done by secretaries and we have cars to do the work formerly done by coachmen. In an ideal world,

we want things like payments to be completely painless: we want to be recognized, and we want to be sent the bill—like a 19th century nobleman going in to a tradesman on High Street.

This enormous gap between the heaven of Downton Abbey and the hell of featuritis is what the designer has to somehow navigate. Now the problem is that, for most of the world, you are not in a position of having your own machine made by a company like Apple that was run by somebody who is a maniac for design. Systems come out of a long process of evolution, whereby there are various incentives facing the various players.

When you start talking about the trade-offs, such as fraud versus privacy or speed versus resilience to abuse, then I think the key question is this: What is the evolutionary environment of the mobile payment system? Which are the more concentrated and the more effective stakeholders? Will the environment be entirely molded by the Barclays Banks, the Wal-Marts, and the Googles? Will there be regulatory pressure as well? Will there be pressure coming from the civil court system through tort claims and contract cases and so on? How do we arrange things? How do we do the mechanism design so you end up with a payments system which has a reasonable equilibrium we can live with?

Mr. Fish: I will open up Q&A with a question of my own and then we will take questions from the floor.

I know from my work there is the big BYOD (bring your own device) movement, where consumers want to use their personal devices at work. And organizations are being forced into this, because they need to support that for their employees, but they feel this represents their No. 1 security risk.

You had discussed how payment applications tend to be insecure and, unlike a credit card, this now puts the enterprise at risk. Do you see a situation where an enterprise can now hold a payment provider liable for a breach that occurred because of their software?

Mr. Anderson: A big problem, of course, facing a medium-sized company, like I suppose Cambridge University with a few thousand employees and a few hundred million a year of turnover, is what happens if your finance department gets spear-phished. That is the big threat nowadays, because as a corporate body, you do not have the protections offered to a consumer. You are supposed to be a grown-up. And yet, when we look at the types of compromise that happen nowadays, very often the bad guy manages to get, say, 30 of the 50 guys in your finance division.

Old-fashioned accounting rules do not necessarily help there, because double-entry bookkeeping rules were invented to deal with one dishonest person, or alternatively one compromised machine. Once you have three or four, all bets are off. So there may be a case to be made for diversity of platforms.

Alternatively, you may want to make a case saying now consumer electronics have made devices so cheap—this is something I have actually recommended to organizations—you should see to it that your serious money bank account payments are made on a machine that is never used for any other purpose at all. Have an iPad that is kept in a safe and it runs your bank's app and is never allowed to run a mail client or a browser and certainly not a game. So the falling costs of consumer electronics can be a benefit as well as a problem.

Mr. Acquisti: Something I am seeing happening in this area (and also in the educational sector), is organizations outsourcing some of their services, precisely to avoid those liabilities. But that does not necessarily solve the privacy/security problem. It simply switches it to another party. The outsourced party, because of its specific knowledge and expertise, may be better equipped. But precisely due to its being large, and having lots of data from many different entities, it represents also a bigger target for the attackers. So, by increasing security in some sense, you are also increasing the incentives for the attackers to go after that type of entity.

Ms. Hughes: It would seem to me the kinds of experiences we have had, perhaps as individuals with data security risks, normally do not affect us very much except in the hassle factor. It takes us awhile, unless we have actually had identity theft.

Someone, not so long ago, tried to get a $250,000 mortgage in my name for a location I had never been and somehow had managed to get a hold of my Social Security number. Now somebody had the good sense not to give them the $250,000, but I would have had a terrible hassle. So I am not the university and I am not being drained of $300 million or $300 billion, but nevertheless to unscramble that would be a terrible problem for an individual if in fact the transaction had gone through.

That suggests to me Ross's advice is very shrewd. Certain things really need to be firewalled off, so that you can control some of your risks. And then you are going to have to figure out which other risks you are going to have. The university I work for has just announced that, unless those of us who also have computers at home that can link to the university's systems follow certain protocols, it will simply cut us off and no longer allow us to do that. There will be no telecommuting into the university's main email server, for example, unless we follow certain protocols and on a regular basis.

Getting everybody to do that with their mobile phone, getting your teenager to do that with the mobile phone is really going to be interesting. If you took the PayPal example and you are giving $80 to that teenager, but their phone may not be linked to your phone unless it is in one account, then that causes all sorts of other planning and employee behavior monitoring problems for us. It also will impose on the persons who suffer the attacks, as Ross has suggested, a duty to report fast and loud, so we can keep it from happening to others if there is something catastrophic in the works.

Mr. Fish: We will now take questions from the floor.

Mr. Burns: Dr. Anderson, I have a question for you, if I could. I was very delighted to hear you call for some form of registry of fraud data, payment data, and front-end payment data in this country, because we obviously need it. I am somewhat aware, but not totally aware, of an arrangement in the United Kingdom, where these data are reported on a regular basis and managed.

I have two questions. One, Do you have any sense about why we do not do it in this country? And, two, Is this system or the collection mechanism in the U.K. as comprehensive as you were arguing in terms of the different kinds of fraud, because obviously counting is a problem in many areas?

Mr. Anderson: I cannot really comment on why such an organization has never been set up here. I hear various things anecdotally, but certainly it is a good thing—it is de rigueur and it is being done elsewhere.

Britain was one of the first two countries to start doing systematic fraud reporting; the other was France. We have somewhat fallen behind the French, who have been enthusiastic in leading the European effort.

In the U.K., as you may know, there is the U.K. Cards Association, which gets information from the banks and provides relatively aggregated figures to the outside world. So we know, for example, how much was lost from the post, from card-not-present and so forth. But we do not have it broken down by individual bank, because that would be beyond the comfort zone of the participants.

What the U.K. Cards Association doesn't do is to talk to nonbank payment channels. So it would be great if a U.S. system being consciously designed could do more than either the British or the French systems do now. I am acutely aware of the fact that, if you try to legislate for such a thing to be set up, it would take years and years and years. And we do not have years and years and years.

So rather than doing something by compulsion, it may be better to do something simply by asking people nicely. I favor putting together a multistakeholder agreement, in which hopefully most of the serious players will collaborate and those who do not can over time be nudged and shamed and gently bullied along until they start to join in.

Ms. Hughes: The other thing happening in the United States, which has not been getting a great deal of attention, is the October 2011 SEC Corporate Finance Staff Guidance on Cyber Security Risk Disclosures and Events and what the remediation efforts are, etc. If you are not familiar with it, it is terribly hard to find unless you go to the Corporate Finance Division's own website, because there was no press release and it was not a commissioned statement of policy. It is just staff guidelines for the purpose.

But they go through six or seven different aspects of cyber security, event disclosures, including management and analysis—things that would be classic possible material changes. If the attack were large enough and you were a publicly traded company to affect your bottom line in a material way, the number and disclosures that theoretically could be made or have to be made are quite considerable.

We think they may make people very cautious about disclosing things. They want to know what you did to remedy the problem and they want you to describe the problem you had.

I venture that very few people in the room who are in payments are going to want to explain to the world in their SEC filings how it was they happened to get hacked. I just cannot imagine that is going to happen. My hunch is this is something anyone who is a publicly traded company should take very seriously, but they really need to talk with the person who handles their SEC materiality questions to determine precisely what they have to say. Otherwise—and a colleague and I wrote a very short paper about this about three months ago—there is a risk it will help the hackers more than it will help investors and businesses. The delicate balance is between not helping the hacker too much and helping yourself and keeping the SEC and your investors from suing you. Also in our paper is an argument that you may be road-mapping the shareholder derivative suit when you make these disclosures, which I think also no one in the room will wish to do.

Mr. Sullivan: I have a quick, two-part question. The privacy concerns of all these data being out there are tied, to some extent, to the potential damage that can happen when they get stolen. A large channel for that damage is payment fraud. I am proposing, if we can find a way of approving payments without having to rely on all the information about my background and my location that would be a good thing. I am curious about your reaction to that.

Secondly, there is always hope that maybe there is a hardware solution, like an EMV card. I am familiar with Ross's work, and his important work at showing how EMV has some security holes. A lot of that is simply because of sloppy implementation. If the implementation is right, the hardware could work very well at primarily getting appropriate payments into the system.

It is a two-part question and the parts are interrelated. Can we get a way of separating information from payment approval? Is there any hope for a hardware solution?

Mr. Anderson: Well, Rick, yes, I would agree with you that many of the problems with EMV are down to poor implementation, but not just poor implementation. There has also been a lot of sloppy design work. But the EMV documents are thousands of pages long; they are many shelf feet. When we get a new student onstream, we almost invariably discover a new vulnerability and almost invariably now you have to look in six different places and four different books in order to track it down.

You need to have mechanisms, not only to design systems better, but also to maintain the design of the systems as they evolve. This appears to be a problem with EMV. Back in the 1990s, when it was all new and fresh and bright and interesting and sexy, you could get bright engineers and academics to go to work on this. Now that is all really old and boring and tiring and complex and "crufty," and you have hundreds of different vendors fighting each other and thousands of banks complaining about this, it becomes that much more difficult.

How do you solve this core governance problem? If you can get the technology right, then yes, there are things you can do to make privacy a little bit harder to compromise. What we do, for example, is use the hardware tamper-resistant EMV chips in order to authenticate gazillions of payments. Then, again, there are economics issues of how you go about motivating people to accept a privacy payments option, if it means they do not get any air miles.

Mr. Acquisti: Thank you for the reference, because this is closely related to some experiments we did recently. Your question, Rick, is seminal to a debate every privacy conference ends up talking about—trade-offs or ostensible trade-offs between privacy and security. To have secure transactions, you can go one way, which is gathering more and more data about the individual (where they are, who they are, what time it is, which clothing they are wearing).

Or, you can go the completely opposite route. One example I gave was e-cash, based on blind signatures. I have no vested interest in e-cash whatsoever. In fact, the patent for blind signatures-based payments even expired a few years ago, so there's no money to make there. However e-cash was arguably a pretty secure system with complete authentication, without identification.

To clarify: I refer to an identifier as something like your telephone number. You can make it public. People use the number to connect with you. The authenticator is, instead, the four-digit code number you use when you access your voice mail. No mentally sane person would rationally want to use the same number as an identifier and as an authenticator. In fact, this is the way in which most financial systems use passwords. For instance, Social Security numbers in the United States are used as identifiers and authenticators. Similarly, when you reveal your credit card, you are providing information that can be used later to impersonate you. Not so when using blind signature. Now, of course, cryptologists know you can take a provably secure system and then, when you actually deploy it, you start adding vulnerabilities in the way you deployed it. Fair enough.

But at least in theoretical terms we have alternatives. So, the answer to your question is a resounding "yes."

The next point I would like to make is that we can offer economic incentives for the different stakeholders to use it. Professor Anderson was pointing out research about whether and how much people will want to pay to protect their data.

It turns out that, yes, there is a significant group of people who will pay a little bit more, but it is not a majority of the people.

Mr. DeCicco: Professor Anderson, I want to go back to the comments you made about the U.K. Faster Payments service. You talked about it being a target for phishing gangs to potentially get money out of the market there. The clarity I am looking for is, Is there already evidence that this is occurring or is that an issue or concern the market has and it is something they need to manage against?

Secondly, the U.S. market is currently considering our own version of Faster Payments. It would be a proposal out for same-day settlement in the ACH system. As we continue to debate that in this marketplace from a safety, soundness, and fraud mitigation perspective, are there issues or advice you can give us and points we should consider to stand it up in a correct way?

Mr. Anderson: Well, the Faster Payments issue is an industry concern, which I have heard from a number of firms that are involved in this. We do not have statistics yet, because where there are phishing losses, banks typically eat those. Statistics should feed through via the U.K. Cards Association and so on in a time scale of approximately a year. Given the different implementations by different banks, industry insiders at least should be able to take some view on how bad the problem is, perhaps within a year or two.

Generically, if you look at the paper I brought to this conference four years ago, we found there was a strong correlation between the speed and the energy with which banks go about stopping, revoking, and recalling stolen money, and inversely with their vulnerabilities to phishing. And it was those banks that were not very vigorous at stopping suspicious transactions and clawing them back that ended up taking most of the losses.

As far as the implications for same-day settlements in ACH are concerned, I would be most concerned if same-day payments could be used from accounts likely to be compromised and, in particular, to send money out of the country or to places where it could be effectively laundered.

The working assumption to make an engineering list is that you should assume something like 5 percent of all consumer PCs are compromised with malware. Before people do work to take Zeus down, you must always assume perhaps 1 percent of your clients' PCs will actually be running evil software on them. So you have to take a view on what sort of scams are likely and whether it is worth taking the risk of allowing people to move money out of the country on a same-day basis. Then again, what business benefit do you get from it? Normal consumers do not need to do that perhaps.

Mr. Fish: That was our last question. Thank you, panel. I thought that was a great conversation.