Monitoring Payment Fraud: A Key Piece to the Puzzle

Commentary

Chris Hamilton

We are going to change accents now for a little while. First, I am filled with envy for the quality of the material the Observatory collects and publishes. We are still a far cry from that in Australia. It is wonderful to see that kind of quality of data available. I do not want to spend a lot of time on how we do what we do in Australia. In fact, I am going to draw into a statistics presentation without talking about too many statistics. The sheer depth of what Alexandre Stervinou presented to you is a testament to how interesting and potentially useful these data are. But I would really rather talk about the whys and the politics and policy behind this kind of data collection because I think it is more relevant to coming to grips with the public policy implications and what should be done by the industry. Let me start with an anecdote about my past life.

A long time ago, when I was a much younger man, I used to work for the Australian Stock Exchange. You probably know that more than 20 years ago stock exchanges around the world went from being what is called "open outcry," where everyone yells at each other in a big room, to being electronic, where they all sit at computers and do not talk anymore and just tap the keyboard all day. Some stock exchanges still have a bit of theater around them; the New York Stock Exchange is an example. One of the side effects from going from open outcry to computerized trading is that you go from a situation where the information that is known about the stock market, who is doing what where, the speed of transactions, what stocks are moving, all that is being picked up at the event. If you really want to know it, you have to stand in the room. That is open outcry. We have gone to a world where the entire performance of the stock market is available, down to keystrokes at the hundredth of a second level to anyone who wants it as long as the stock exchange is prepared to give it to them. You go from a situation of quite limited data about what is going on in a very complicated

human environment to where you have almost unlimited data. And that has some very interesting effects on how things are done. This is the analogy I am trying to draw. When the Australian Stock Exchange computerized, which it did relatively early by global standards, insider trading became extremely hard. Although you cannot always tell when it is happening in the market, surveillance experts say an electronic record of trading can always tell you if someone is insider trading because you can see them moving before the announcement. If you have keystrokes down to the hundredth of a second, it does not matter how clever they think they are. You can work it out from the data. What you need is a good surveillance unit that puts two things together—detailed information about trading on the marketplace and key events in a company's history. The trouble with insider trading is that sooner or later the event has to come out, you have to know, and so you catch the crooks that way. One consequence of really good data is a completely different approach to enforcement and quality of law enforcement. But another very important consequence of that change was that a whole academic discipline and tradition grew up around analyzing this volume of data about the trading market to understand how markets work. As we heard this morning, the application of game theory to how stock exchanges work has become an enormous academic growth industry and people understand much more deeply now how markets work because they have this detailed information. There are, however, all sorts of unintended consequences from being able to capture this data, some positive, which are worth bearing in mind.

When I moved to work in payments, about 10 years ago, I felt like I had been blindfolded. We are lousy at data, and we should be ashamed. The quality of detailed data about performance of the payments systems around the world is really lacking and someone should do something about it. The information that we have is after the event. We have publications; I did my publication a couple of weeks ago and Alexandre is doing his in a week. We have data coming out six months after the relevant period. We have relatively high level data about how things work, and we are only able to draw very broad inferences, which we then need to explore further. So, the first thing to say is we should do this a whole lot better than we do, and there is no technological reason why we cannot. As always, it is the human, the economic and social organization part of it that is the challenge.

I want to talk about that. What is it we are trying to capture, and why? Why is that a good idea? Who should capture it, and who benefits from

Chart 1
Australian Card Fraud by Type, 2009-14

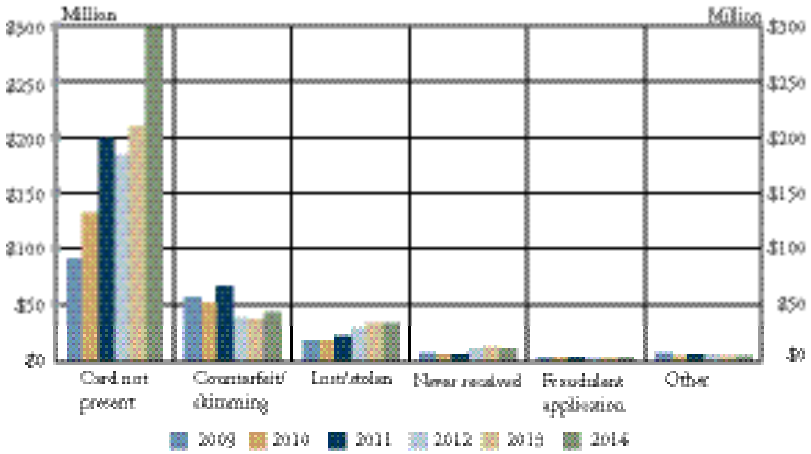

Source: Australian Payments Clearing Association.

that capture? Those are the things I want to address, and I will try and draw some reference points from the French experience.

What we want to capture is reasonably clear, and there is an endless further level of detail you can go down to, but the Observatory gives us a very good starting point in terms of what are great things to capture. You want to know about the sheer rate of fraud, the prevalence, and have it broken down in as many different categories as you can. In Australia, we do something similar. We recently published our 2014 numbers (Chart 1). We have been tracking fraud data for about 12 years. This is just a five-year horizon to give you a sense of what is happening, and you can see very starkly the kind of experiences you see in the French data. Card-not-present (CNP) fraud is the big problem of the day. Everything else is nearly solved. It is either flat-lining or dropping. But CNP is the big problem of the age on card data. There is another story elsewhere. Not only is CNP the problem, but offshore CNP is the big problem in Australia (Chart 2). That differs from the French experience just because we probably are on a cycle that lags Europe by a couple of years. I have observed that before, the cycle happening in Europe and then coming to us. That is another good thing to bear in mind as you look at these numbers. And of course, the consequence is, and this is again very similar to the Observatory's experience, over a five-year cycle we have gone from CNP fraud being half the fraud problem to being more than three-quarters (Chart 3).

Chart 2

Card-not-present Fraud in Australia, 2009-14

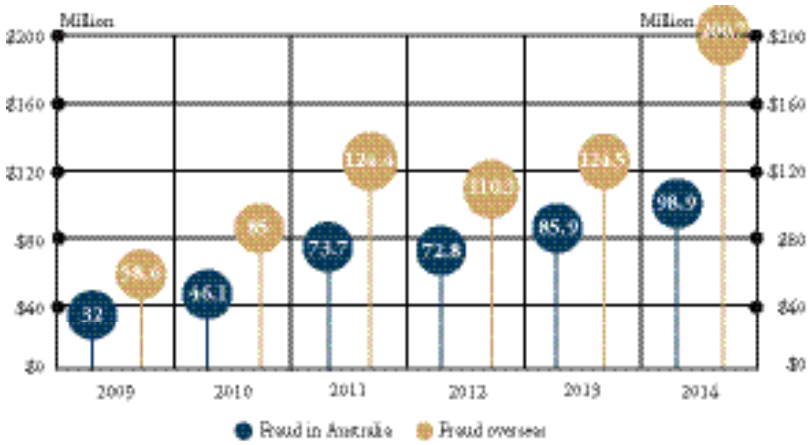

Source: Australian Payments Clearing Association.

Chart 3

Growth of Card-not-present Fraud in Australia, 2009-14

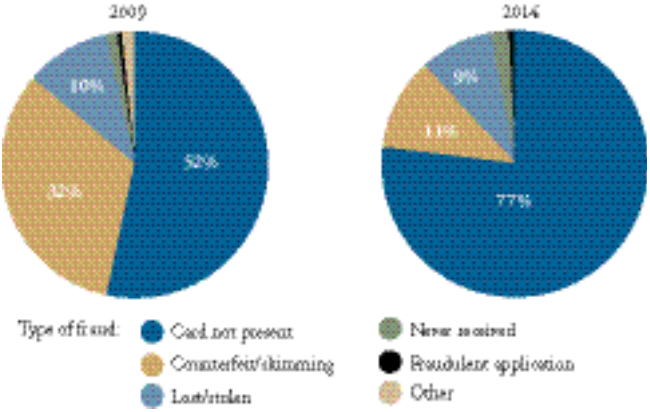

Source: Australian Payments Clearing Association.

Capturing the prevalence, the trend line, is really important, but it is only the beginning of the challenge. The other thing the Observatory does well, and which we do but in a different way in Australia, is capture the threat matrix to determine the upcoming problem—what Alexandre called the technology watch. In Australia, we do that in a much more informal way, sort of a clearinghouse approach where you get the large organizations involved in comparing notes on fraud events. They take away the raw data of observations and do their own analysis. It is a much more decentralized process. You can argue it is both more and less effective for different purposes. It probably is better if they are looking specifically at protecting their own shops because they will have much more detail on the standing of their own customer environment and their own particular risks and vulnerabilities. On the other hand, it is not very helpful for looking at the global picture and seeing what is happening in a broader sense. One thing that has started in Australia is the formalizing of a longstanding informal structure called the National Fraud Exchange, which is sort of a clearinghouse of ideas. The major participants will all fund and provide threat information and use that as a shared resource across the industry. So, formalizing and automating that process is one of our current priorities.

The third thing, which none of us does very well, but which is actually really important, is impact analysis. What happens when fraud happens? Who actually loses, and what are the costs both of prevention and of the actual event itself? And this is really hazy. We saw some of that in the first series of presentations. Is it really right that the consumer does not bear the fraud? Is it really right that the issuer does? In Australia, officially the issuer bears the fraud, but in practice the great bulk of the fraud is probably borne by merchants because of the various liability shifts. That has very big impacts on their incentives to change and the way they are going to work or not work with the industry. For me, that is the least well-developed of data areas that we should be working on. What are the real costs of this stuff? I am sure the global cost of EMV implementation dwarfs the actual savings in fraud. There is no question that we have all spent a great deal more putting the EMV chips in cards than the fraud that we have saved from doing so. That does not necessarily mean it is a bad idea, but it probably is a useful thing to know. There needs to be much more on that work. If that is what we are trying to collect, then it is worth thinking about the whys. What are we going to do with this when we get it, and who might benefit?

There are several very good reasons. I come at this from an industry perspective. The Observatory thinks about things from a public policy

perspective—what is in the best interest of the community? I am coming at it from a slightly different perspective and I should explain what the Australian Payments Clearing Association (APCA) is. It is not a government body; it is completely privately funded. The nearest equivalent in the United States is the National Automated Clearing House Association (NACHA), but we are not that much like NACHA; APCA is an organization that administers the rule books on behalf of the financial institutions in payments. So, we are only as good as the collaboration we can persuade our members to perform in improving the overall payments system. Our goal is to improve the payments system, but from the industry perspective of how do we work together as a community on what is important to all of us to make the payments system better, rather than what is the public good. Public good clearly comes into it; it is clearly a big factor. But we need to marry that with the collective industry of the community. Coming with that lens to this fraud data, why would you voluntarily publish fraud statistics? In many countries, that does not happen and there appear to be good reasons why. People do not want their brands associated with large reported frauds. People do not want to scare off customers with stories of fraud. But that is a shortsighted view; the much better path is to think about the long-term gain for the industry. So, forget the public good for a second.

The people mostly affected are our collective customers, the consumers and businesses of the community. There is a sort of moral dimension here where they have a right to know so they can do their own risk assessment. That is one reason why it probably is a good idea, but there also is a practical one, which is they need to be participants in the fraud-prevention process. Consumers and businesses all can do fairly basic sensible things to minimize their own risk and prevent fraud. They cannot, however, solve the problem by themselves. There are many other things other people have to do, but it would be nice if they were active participants in that process. You start doing that by educating them about fraud, by giving them a clear picture of what it is (Charts 4, 5). So, that is a good, practical reason for industries to do this work voluntarily. The other obvious benefit relates to a point made before—what gets measured gets managed. Unless we know what the fraud is, we do not know where to focus our limited dollars on trying to prevent it and improve it. It is very important to have that kind of data when you are arguing the case for whether we should do EMV, or go to two-factor authentication or 3D Secure. And not having good quality data is one of the things that makes that process quite hard. In Australia, we had an initial go at EMV, at chip cards, more than 10 years ago; not as far back as the French. That effort

Chart 4

Gross/Net Fraud Values by Fraud Method, April 2015

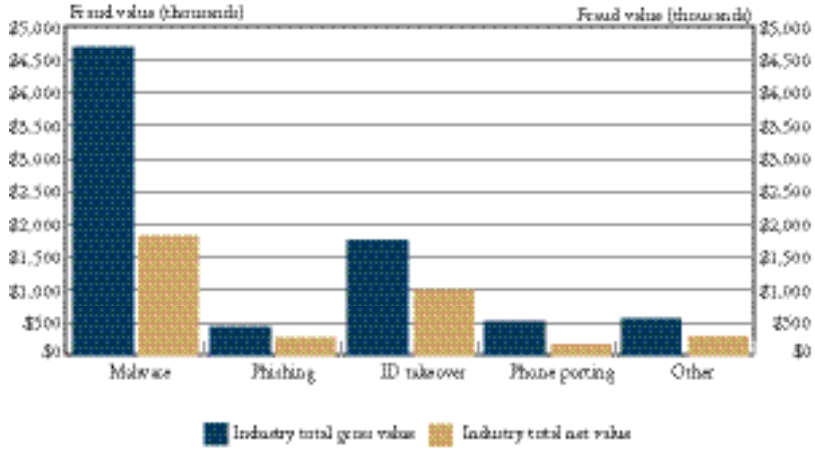

Source: Australian Payments Clearing Association.

Chart 5

Net Fraud Value by Fraud Method, July 2014-April 2015

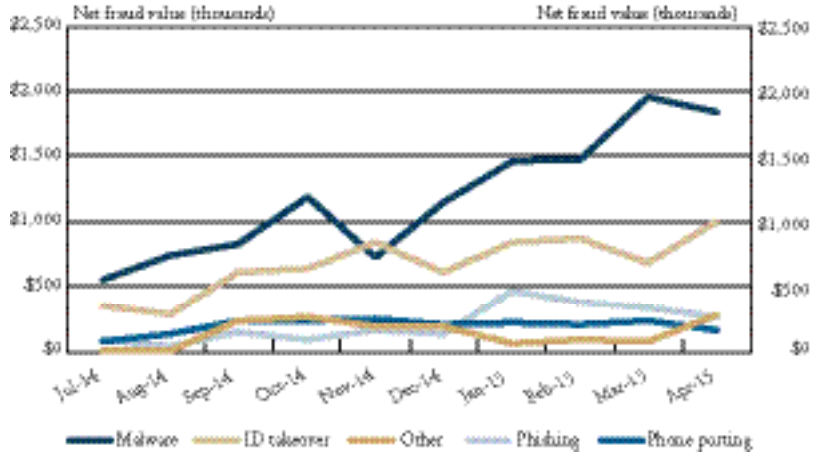

Source: Australian Payments Clearing Association.

failed through lack of articulation or a strong enough case for change. I think if we had had the quality of data and the trend lines we have now about fraud, you might have gotten a different result. Indeed, the second time around, having the benefit of that information was at least as important a factor in what has been a very successful chip conversion.

Having a grip on that helps the industry work out what it should and should not do collectively to improve the system. The data also give organizations a much better risk management capability within their own shops. All large banks around the world now are scoring approaches, doing risk approaches to fraud—some are really good at it and some not so good—but they all would get much better if they all had all the data. Seeing their own data is not enough, and having the benefit of detailed information about data is potentially extremely valuable.

If that is what we are trying to achieve, then the last point I want to cover is who needs to do this, and how they should go about it. And I am going to give a slightly different point of view. I do think that this generally is actually better done by industry. I *would* say that, would I not? I work for industry. Natural bias. And yet, my experience is that work to improve the overall payments system, which is done collaboratively by the institutions that work in it, when they are convinced there is long-term benefit both for their customers and for them, is much better done than forced compliance as a consequence of regulation. It is hard to pull off. It is much harder to do. So, compliance in a way is easier. What happens is the banks have outsourced to the regulator the problem of deciding what should be done because the compliance rules tell them what should be done. They can comply and they get to bellyache about it at the same time—sort of a win/win. But in the long run, these things work a lot better if, having been convinced of the need to actually make the change, they then implement it because they will do it in a cost-effective way. They will do it in a way which fits with their business, but still meets the public policy goals.

The last thing I want to talk about is this Australian way of having a go at the public/private partnership. Let me observe that in relation to Adam Levitin's distinction between public ordering and private ordering, I am suggesting that is a bit of a false dichotomy, or at least it should be. What we really should be doing is finding a way of marrying the public and private methods of doing things, and the public and private interests to get the best possible outcome. And I think that is possible, if you can get the industry convinced of the value to them, which is also in the public interest, you can then get a willing, collaborative approach to solving the problems we are talking about. And in fraud, that actually works better than many other areas of changing the payments system because it is easier to convince people that fraud is everybody's problem. It does not tend to have a major comparative element to it. It sometimes does have little bits of competitive tension among the banks, but in general, people agree that if I am lax on

security it is going to affect you and vice versa, and so it is easier to get that collaborative agreement. My suggestion is that in the long run, we need to gather this data because it is in the interest of the industry. But then we need to work on it together to find the best way of improving the payments system using the data itself.