HOUSE OF COMMONS SCIENCE & TECHNOLOGY SELECT COMMITTEE
INQUIRY INTO MALWARE & CYBERCRIME

Written evidence of Dr Richard Clayton

1. I am currently a Senior Research Assistant in the Computer Laboratory of the University of Cambridge. At present I am engaged in a 3-year collaboration with the National Physical Laboratory (NPL) to develop robust measurements of Internet security mechanisms.
2. I have a particular research interest in cybercrime. My research falls mainly under the heading of Security Economics – a field based on the premise that it is easier to explain security issues with an economic analysis rather than simply using a technical or 'computer science' approach. I am particularly interested in measuring criminal activity rather than merely describing it.
3. I have been using the Internet since the early 1990s, ran a software house that created one of the earliest mass-market Internet access products, and worked at Demon Internet, then the largest UK ISP, from 1995 until 2000. In October 2000 I returned to Cambridge to study for a PhD. My doctorate was awarded in January 2006 for my thesis, "Anonymity and Traceability in Cyberspace".
4. I have continued to work in the Computer Laboratory doing academic research. On several occasions I have acted as specialist adviser to House of Lords and House of Commons Select Committees in inquiries into cybercrime and Internet security.
5. I have written, or co-written, over 40 peer-reviewed professional publications. My main research interest over the past few years has been into the criminal activity called "phishing" – the theft of financial credentials by impersonating legitimate websites. More recently I have been starting to look at the role of malware in the criminal eco-system and I have published work on how malware clean-up should be approached from a security economics standpoint.
6. I should also declare that in addition to my employment at Cambridge and my past association with Parliamentary Committees, I am a director of a small consultancy company that sells my time and expertise. Additionally, I am presently employed by Yahoo! in a part-time capacity within their security team.
7. This document expresses my personal opinions, and is in no way the expression of an official position held by the University of Cambridge, NPL, or Yahoo!

Q1. What proportion of cyber-crime is associated with malware?

8. I have been pointing out for years there are almost no reliable figures about cybercrime. In a report I wrote with colleagues for the European Network and Information Security Agency (ENISA)[1] we set out the problems in detail in section #4.2.
9. Summarising 14 pages of densely argued analysis for this submission is impossible, but in section #4.4 we made two recommendations, both of which I would commend to this Committee:

   We recommend that the Commission (or the European Central Bank) regulate to ensure the publication of robust loss statistics for electronic crime.

   We recommend that ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs.

[1] R. Anderson, R. Boehme, R. Clayton, T. Moore: *Security Economics and the Internal Market*. ENISA, Jan 2008. http://www.enisa.europa.eu/act/sr/reports/econ-sec

10. Until we have reliable data we will not be able to assess the size of the cybercrime problem nor whether we are making any impact on it. Of course, assessing that impact in purely monetary terms is simplistic and the Committee ought to go beyond what we recommended to ENISA and require the recording of all electronic crime incidents, not just those resulting in monetary loss.

11. For example, the UK banking industry already publishes fraud loss figures – but these do not actually document how much money has been stolen, but rather how much the banks end up out of pocket. The bank has no loss if they detect the crime promptly enough to undo electronic transfers before the money leaves the banking system (which we understand is achieved in about half of all cases).

12. Additionally, banks regularly attempt to dump their losses onto their customers, personal and business, by suggesting that the failure of security mechanisms is the customers' fault, despite those mechanisms having been specified by the bank.

13. In particular, to return to this inquiry's focus on malware: banks and others have chosen to rely on general purpose browsers and they have chosen to rely on identifying users simply by their ability to regurgitate a password. Unfortunately, when user machines are infected by malware this reliance is misplaced.

14. Most modern malware includes a 'keylogger' – functionality to record all the keystrokes typed by the user and relay them to the attacker. In response, the banks have moved to systems that prompt on the screen for just a few characters from the password. There is now malware that snaps a copy of the screen area around the mouse and the criminal learns the password one character at a time.

15. More sophisticated software performs "man-in-the-browser" attacks by intercepting legitimate interactions with the bank – perhaps paying a gas bill – and replacing this request with a transfer of money to the criminal's account.

16. This type of malware operates in 'real-time' and will defeat the protection provided by the 'CAP readers' (the calculator-like devices that many of the banks have issued). This is because the user will type in the numbers from the screen still believing that they are paying their gas bill. Even after the fraud is complete the malware will keep the user from realising they have been defrauded by rewriting onscreen bank statements to continue the pretence of paying for gas.

17. One could carry on for many pages in discussing numerous different types of malware and explaining all the different types of criminality that it underpins. Unfortunately, this descriptive approach is pretty much all that we have – we have almost no reliable quantitative information.

18. Hence it is not really possible for anyone to give an accurate answer to the Committee's specific question about the proportion of cybercrime that is associated with malware. All that can be said, in the most general terms, is that the eco-system for mass-market criminality is based on spam sent by botnets, and those botnets are constructed by compromising end-user machines with malware. Furthermore, the majority of specialist attacks on high-value targets – performing industrial espionage or compromising finance departments – are also based on malware.

Q4. What is the cost of malware to individuals and how effective is the industry in providing protection to computer users?

19. The committee asks a number of questions about malware authorship and the cost of protection which other experts will be able to address. What I can discuss, from my own research, is the ineffectiveness of protection – and, rather unusually, I even have some detailed numbers about this relating to the activities of one particular criminal gang.

20. First some generalities. Systems such as spam filters act to protect individuals by preventing them from ever coming into contact with malware. However once an email evades those filters and arrives in the inbox with a malware attachment or a link to a bad website then there is almost no further protection at all. Of course some people will see through the 'social engineering' and will not be fooled into clicking the malware into action, but now that the criminals understand what is too enticing to ignore (and now they have fixed all their grammar and spelling errors) clicks are extremely common.

21. I have spent the past year tracking 'Instant Messenger worms' – malware that is spread between Instant Messenger buddies. What happens is that users receive a message over Skype, Yahoo! Messenger, Microsoft Messenger, Facebook Talk etc. which says something like:

    `foto ☺ http://ofacebooks.net/album.php?your@email.addre.ss`

22. If the user clicks on the link in this message then Windows will put up a warning message asking whether you wish to run a program from ofacebooks.net. Most people, I believe, are so eager to see the promised photograph that they will immediately press OK and thereby become infected by the malware.

23. Once the malware is running on a new machine it contacts its command and control system (C&C) to determine what it should do next. The C&C will generally instruct it to send a message to all of the new victim's buddies (saying **`foto…`** etc.) to garner new recruits. The C&C will then download specialist malware (keyloggers, vulnerability scanners, spam senders, etc.) and the machine will be mined for financial data and turned into a resource in a botnet.

24. At the time of writing, my research shows that the malware for the most active worm is being downloaded just over 70,000 times a day and the number of victims, worldwide, is now well into the millions. This research is currently unpublished – but I expect it to be of significant import, not least because for once we have some accurate numbers to work with.

25. One might expect anti-virus software to detect the downloaded malware and hence provide protection. However, the criminals tweak the malware on a daily basis and only deploy it once it is passed as safe. Then of course the anti-virus software is updated, but too late to protect anyone.

26. To take just one example of the how ineffective anti-virus software is: consider the specific version of the malware that the criminals were using between 10:27 and 14:23 GMT on the 5th September. It was tested at 16:54 (90 minutes after the criminals stopped deploying it) and by that time it was detected by only 7 of 44 anti-virus products; and those 7 did not include any of the top 3 products by market share. Even 24 hours later, only 11 products reported this particular malware sample to be bad.

27. Of course, not all malware gets onto people's machines because they click on a link and are 'socially engineered' into ignoring warnings. Some infections result from exploiting software bugs – for example in the add-ons that automatically play videos within the browser.

28. The large software companies such as Microsoft and Adobe provide automated patching systems to correct bugs. However, modern computers are running software from dozens, if not hundreds, of companies – and most of these companies do not have sophisticated patch distribution mechanisms. It would be desirable for companies such as Microsoft to open their patching platforms to third parties so that users could have a fully integrated way of staying up-to-date.

29. Other companies are just as slow at deploying patches, and in particular the mobile phone companies can be years behind at pushing out patches to their subscribers' handsets.[2] This is a classic failure that is easily explained by 'security economics': the people in a position to fix the problem are not those who would suffer a loss. We often have to resort to fixing such problems by regulation – and this Committee should recommend that subscribers should be entitled to claim damages from their network provider if their phone (or their data) was damaged as a result of an unpatched vulnerability for which they have delayed rolling out a fix.

Q5. Should the Government have a responsibility to deal with the spread of malware in a similar way to human disease?

30. Another way that industry fails to protect Internet users is by failing to act when their users are known to be compromised.

31. It is often possible to record the unique IP addresses of machines that are contacting a C&C system. Additionally, when a botnet is shut down it is now usual practice to set up a 'sinkhole' that will log the identities of the compromised machines which continue to try and make contact with the disabled C&C.

32. The operators of the sinkhole are unable to communicate with the owners of the compromised machines directly – they can only identify the ISP that is providing Internet connectivity. So it is up to the ISP to pass the bad news on to the relevant customer, because only the ISP knows who was using the IP address at the relevant time. In practice very few ISPs relay information and almost none go looking for further sources of this type of data.

33. We can see how poor the data passing is by examining the data collected by the Shadowserver Foundation, who operate a sinkhole for Conficker – malware that infected 7 million machines worldwide in November 2008 and which still poses a threat to the infected machines. The Shadowserver data[3] shows that infections have dropped from 5.5 million in September 2010 to 3.5 million now; the worst affected UK ISP has seen a reduction from 7000 to 5000 infected machines over the same period. The best ISPs completely eradicated the problem, and ensured their customers were safe, two years or more ago, and I suspect that the drop in numbers is as now much to do with old computers being scrapped as customers being told of their problem.

34. The reason that ISPs discard notifications is because contacting their customers is expensive – the standard meme is that one phone call to a customer wipes out the profit made on them for a year. This makes a good sound-bite – but it is roughly correct. My own analysis shows that the cost equates to 8 months of profit, so the ISPs are indeed acting rationally in so far as their own self-interest is concerned.

35. Financial concerns are the basis of the industry-wide agreements (in Germany, Australia and The Netherlands) in which all the ISPs promise to pass on malware infection notifications. The idea is to ensure that no-one can steal market share by undercutting prices by failing to incur the cost of contacting customers.

36. This committee should recommend just such an ISP industry-wide agreement in the UK. However, the recommendation should go further and instruct the ISP industry to explicitly seek out sources of data, from sinkhole operators and others, so that UK Internet users have the best possible chance of being told if their machines are harbouring malware.

---

[2] R. Lemos: Fast phone patching still a fantasy. CSO Magazine, 7 Apr 2011. http://www.csoonline.com/article/679205/fast-phone-patching-still-a-fantasy

[3] http://www.shadowserver.org/wiki/pmwiki.php/Stats/Conficker

37. The Committee should pay particular attention to the system being operated by Comcast – the large cable provider – in the United States. They monitor traffic to their domain name servers – the machines that convert human-memorable hostnames into the IP addresses needed to communicate across the Internet.

38. Comcast use a datafeed from Damballa, a specialist anti-malware firm, to identify when hostname lookups are performed by malware that is attempting, for example, to locate C&C servers. When the presence of malware is deduced then the customer is informed, usually by means of a pop-up message when they next use their browser.

39. One of the many reasons that ISPs fear talking to customers about malware is not just that they want to avoid delivering bad news, but they fear being pressured into having to explain all of the detail – and then being roped in to fix the problem. What Comcast have done to avoid this is to provide substantial online help and links to free online clean-up tools. Further, they have done a bulk deal with a specialist company who will, for an $89.95 fee, give personal help to customers.

40. I considered the economics of this type of clean-up operation in a paper that I presented to the Ninth Workshop on the Economics of Information Security in 2010. This peer-reviewed conference is the leading forum for work in the Security Economics field. A slightly revised version of the paper was subsequently published in Volume 81 of the Communications & Strategies journal.[4]

41. My paper,[5] "*Might governments clean-up malware?*" supposed that the government would subsidise the cost of malware clean-up, and modelled what the costs might be. I considered a world in which ISPs passed problem reports on to their users, but if the user could not fix the problem they would be referred to a standard clean-up service. The users would pay a nominal sum ($30 (£20) perhaps) to avoid any moral hazard, and the government would subsidise the rest.

42. The thrust of my argument is that this is not as expensive a scheme as it might at first appear because the contractor would be able to sell other services off the back of their interaction with users. Hence they would swallow some of the subsidy costs themselves in order to land the government's contract. My modelling suggests that the actual cost for such a scheme would be less than £0.50 per citizen per year – comparable with the costs of fluoridising the water.

43. There are of course numerous details and assumptions in this scheme, and I refer the Committee to the full paper for all of the details, and a discussion of the advantages of involving the government in such a scheme. The Committee might also note that the German malware clean-up initiative[6] is partially funded by the German government.

Q6. How effective is the Government in co-ordinating a response to cyber-crime that uses malware?

44. The government has not dealt with cybercrime effectively, whether it involves malware or not. Successive administrations have failed to provide adequate funding to grow and develop the specialist police units who work in this area. A very small number of officers have practical experience of tackling cybercrime and this has given them a rarity value in the job market, so that personnel retention is a significant issue.

---

[4] R. Clayton: *Might governments clean up malware?* Ninth Annual Workshop on Economics and Information Security (WEIS10), Cambridge MA, US, June 7–8 2010.
R. Clayton: *Might governments clean up malware?* Comms & Strategies, 81, 2011, pp. 87–104.

[5] http://www.cl.cam.ac.uk/~rnc1/malware.pdf

[6] https://www.botfrei.de/en/index.html

45. The Committee should be recommending more resources – if only because cybercrime is volume crime that affects very large numbers of citizens. We have (a rarity as ever) some good data on credit card fraud, much of which is Internet related. A supplementary document to the British Crime Survey was published by the Home Office in May 2010. It looked at data from 2008-09 and found that 6.4% of credit card owners were aware of fraudulent use of their card over the previous 12 months. Victimisation rates were higher at 11.7% for incomes over £50,000/annum. If the Internet had been used at all (irrespective of income) the rate was 7.7% and if the Internet was used "every day" then it was 8.9%. In contrast, the 2010/11 British Crime Survey found that burglary affected just 2.6% of households and thefts from cars affected 4.2% of households.
46. There has also been a complete failure by government to even start to address the need for effective international responses to cybercrime. Police work needs to be coordinated at the international level, because otherwise committing a crimes in another country will make you untouchable.
47. In the US when 1930's bank robbers used the new-fangled automobile to flee across state lines, the solution was to make bank robbery (along with auto-theft and other related offences) into federal offences rather keeping them as state-specific infractions. However, this solution does not look to be practical for cyberspace, because there is no global body with the equivalent reach over the world's countries that the US federal government had over the individual US states.
48. We are not going to see cyber-police operating across borders in the near future, but we should be looking to see substantially more international cooperation in pursuing criminals in one jurisdiction who have committed crimes in another.
49. The best solution that I and colleagues have been able to suggest (in the ENISA paper already mentioned above in paragraph 8) is a liaison system such as Eisenhower developed in 1943 within SHAEF and which morphed into NATO. In such a system police forces would dispatch trusted officers to formulate pan-European (or preferably global) strategy for dealing with cybercriminals. Their role would be to represent their country's police forces, and within the global strategy they would make tactical commitments to deal with criminals on their own soil and would ask for help with pursuing those who targeted their citizens but were based abroad.
50. We need proper international cooperation – to move beyond the current approach where every national police force targets the same, biggest, multi-national criminal gang and no-one worries about the rest of the top 3, let alone the top 10. We must end a situation where cybercrime is a lucrative career choice with a miniscule risk of ever being chased after, let alone caught.

Dr Richard Clayton
7th September 2011