Driving Factors

CDT Threats & Risks: Session 7

Matthew Edwards

Nov. 28th, 2019

Section 1

Orientation

Structure: Fortnights

Threat Modelling

Unknowns

Risk Management

Driving Factors

Scaling Analysis

Today's Goals

- Introduction to cybercriminology, major criminological approaches.
- Tentative summary of evidence on cybercriminal characteristics.
- Introduction to security economics.

Section 2

Cybercriminology

The critical question

Why do (some) people commit (cyber)crime?

Theories

1. Neutralisation theory
2. Self-control theory
3. Social cognitive theory
4. Routine activities theory
5. The theory of planned behaviour

Neutralisation theory

Core concept: people alter their perception of the seriousness of (their) cybercrime, to make themselves happier with their own actions.

Neutralisation theory

Tell a lie

Most people feel uncomfortable when telling a lie, because they think of themselves as honest people. This feeling is an example of cognitive dissonance. It also arises when people learn new information that challenges something they believe.

Neutralisation theory

a) Hacking is unethical.

b) I just hacked someone.

$\therefore$ I did something unethical

Neutralisation theory

a) I'm not a bad person.

b) I just hacked someone.

$\therefore$ Hacking is sometimes okay.

Neutralisation theory in practice

1. Denial of responsibility ("I had no choice").
2. Denial of injury ("It doesn't hurt anyone").
3. Denial of victim ("They deserve it because they're...").
4. Condemnation of condemners ("They did it as well!").
5. Appeal to higher loyalty (For a cause/principle.).

Neutralisation theory in cybercrime

Attitudes towards digital piracy.

'Illegal access' amongst students.

Use of 'DDoS-as-a-service' tools.

Discussion

What would neutralisation theory suggest we do about cybercrime?

Self-control theory

Self-control is a psychological trait linked to the ability to resist temptation and impulses.

The Marshmallow Test

`https://www.youtube.com/watch?v=QX_oy9614HQ`

Self-control theory in cybercrime

Low self-control is associated with a number of negative life outcomes, and a lot of traditional crime.

Associations with cybercrime as well, including identity theft, online drug trade, distributing malware.

Discussion

What would self-control theory suggest we do about cybercrime?

Social cognitive theory

We learn our behaviour from the people around us.

The behaviour of our role models, and the normal behaviour of people around us (perhaps even fictional), decide what kinds of behaviour we consider acceptable.

The same process also affects the skills we learn.

Social cognitive theory

Engagement in digital piracy has been linked to peers engaging in it, and to the perception that important others expected it to happen.

Cyber-bullying, much online ‘mob behaviour’.

‘Fake news’

Discussion

What would social cognitive theory suggest we do about cybercrime?

Routine activities theory

Three things must align for crime to occur:

1. The existence of an attractive target (e.g., credit card details on a server).
2. The presence of a motivated offender (e.g., hacker who needs to pay rent).
3. The lack of a capable guardian (e.g., poor technical security).

Routine activities theory

Anything that supplies opportunities for motivated offenders to come into contact with targets lacking guardianship would tend to increase crime.

The Internet provides boundless opportunity for offenders to come into contact with targets.

A general prediction borne out in studies: greater internet use associated with greater cybercriminal activity.

Discussion

What would routine activities theory suggest we do about cybercrime?

The theory of planned behaviour

Three mechanisms by which intentions to engage in a behaviour are influenced:

1. Their attitude towards the behaviour / those who engage in it;
2. Perceived social norms related to the behaviour;
3. Perceived ease of engaging in the behaviour.

The theory of planned behaviour

Predominantly used to model secure online behaviour, but also used to understand cybercriminality.

Combines elements of social cognitive theory with **rational choice theory**.

Discussion

Design a study that would identify the common traits of cybercriminals.

Demographic Observations

Cybercriminals are:

Age More likely to be in younger age groups, particularly late teens and twenties. (Caveats: existing biases, generational effects, poor study)

Gender More likely to be male. (Consistent with both crime and technology).

Education More likely to be highly educated. (Caveat: compared to non-cyber crime, evidence weak).

Tech Spending more time on the internet/using computer more/being more technically competent.

Employment More often:

a. Employed in computing/technology profession.
b. Unemployed.

Self-control More likely to have low self-control.

Motivation Observations

Best supported motivations are (all supported by multiple studies):

'Addiction'	A combination of low self control and rewards from the activity, sometimes discussed with reference to 'flow' states.
Curiosity	Intellectual curiosity about security of systems, the old-school hacker ethic.
Enjoyment	Cybercrime is fun (challenging, forbidden).
Money	Financial reward, important for organised criminals but also opportunists in need.
Status	Culture around cybercriminality rewards exploits with social credit.
Ideology	Information should be free; hacktivism.
Revenge	Victims of cybercrime become cybercriminals in response.

Section 3

Security Economics

Why Information Security is Hard[1]

Traditional, technical view:

> *"Given better access control policy models, formal proofs of cryptographic protocols, approved firewalls, better ways of detecting intrusions and malicious code, and better tools for system evaluation and assurance, the problems can be solved."*

[1] Anderson, R. (2001). Why information security is hard – an economic perspective. In Seventeenth Annual Computer Security Applications Conference (pp. 358-365). IEEE.

Why Information Security is Hard

The problem of security is a problem of **incentives**, which can be addressed through the tools of economic analysis.

- network externalities;
- asymmetric information;
- moral hazard;
- adverse selection;
- liability dumping;
- tragedy of the commons.

Example: ATM fraud

Legal precedent in the US: *A bank customer's word that they have not made a withdrawal is found to outweigh the bank experts' word that they must have done.*

At the time, no corresponding precedent in the UK.

In the US
Onus was on banks to prove customer defrauded them. Systems were better protected against fraud.

In the UK
Onus was on customer to prove bank was wrong – basically impossible. Banks were careless, poor fraud security.

Examples

"In general, where the party who is in a position to protect a system is not the party who would suffer the results of security failure, then problems may be expected."

Medical payment privacy If systems are paid for by insurers rather than hospitals, patient privacy is not protected whenever this conflicts with insurers wanting data.

Digital signatures Risk from a signature being forged is transferred from the bank (building the system) to the customer.

Tragedy of the Commons

Traditional Grazing of sheep on the common land. Overgrazing degrades the land, but individuals get rewarded for adding a sheep to the commons. *Solution: locally-decided communal rules about grazing rights*.

Cyber Preventing your computer from joining a DDoS. DDoS causes harm to others, but preventing your computer joining these botnets costs you. *Solution: ?*

Network Externalities: Background

1. Technology has high fixed costs (develop the first copy) and low marginal costs (cost to produce copies).
2. As a general rule, in this situation price competition would drive the price towards the marginal cost of production. The marginal cost of production for software $\sim= 0$. So businesses need to sell based on 'value' rather than their production costs.
3. The value of networked IT depends on how many other users adopt it.
4. There are large costs to users for switching, leading to lock-in.
5. Even if a competitor would be cheap to set up, the market remains profitable for a big, early player, leading to a 'winner takes all' effect that favours first-movers.

Network Externalities and Security

- If first-movers win the market, spending time on secure design could doom the business.
- Microsoft software won success by appealing to developers, so externalities for users (poor usability and security) were rational.
- Administration of security pushed to users, even if less effective, because the designers don't want to shoulder those costs.
- Companies go for obscure/patented approaches to increase lock-in and make it harder for competition to arise, regardless of whether security of these systems is tested.

Adverse selection

The market for lemons

Where buyers don't know the quality of the product, there is severe downward pressure on both price and quality.

"Plum": \$3000
"Lemon": \$1000
Equal-odds pricing: \$2000

Application to information security. . . ?

Cost asymmetry

Why do attackers find bugs first?

Bugs to find: 1,000,000

Mean time-to-find 1,000 hours.

Attacker investment: 1,000 hours/year.

Defender investment: 10,000,000 hours/year.

In one year, the attacker might find 1 bug, while the defender has found 10,000. Yet the probability the defender has found the attacker's bug is only 10%.

Measuring the costs of cybercrime[2]

Governments and organisations need to invest in information security.

How much should they rationally invest?

Asking the security industry directly is like asking the car dealership how much you should spend on a new car.

[2] Anderson, R., Barton, C., Böhme, R., Clayton, R., Van Eeten, M. J., Levi, M., Moore, T., and Savage, S. "Measuring the cost of cybercrime." In *The economics of information security and privacy*, pp. 265-300. Springer, Berlin, Heidelberg, 2013.

Cybersecurity confusion

Many sources of data, all insufficient and fragmented. Some crimes under-reported, some over-reported. Errors can be both intentional and unintentional.

Questions we struggle to answer:

- How many phishing websites are there?
- How many different attackers are out there?
- How many different types of malware?

Overestimates lead police forces to believe they cannot do anything, even when strikes against a small number of gangs could be far more effective than public information campaigns.

Differentiating cybercrime from other crime

Working definition from the European Commission (2007):

1. traditional forms of crime such as fraud or forgery, though committed over electronic communication networks and information systems;
2. the publication of illegal content over electronic media (e.g., child sexual abuse material or incitement to racial hatred);
3. crimes unique to electronic networks, e.g., attacks against information systems, denial of service and hacking.

Cost framework

Criminal revenue The money made by a cybercriminal, excluding 'lawful' costs.

Direct losses Monetary equivalent of losses, damage or suffering felt by the victim (e.g., Money lost, time and effort spent resetting account, distress).

Indirect losses Costs imposed on society because this cybercrime exists (e.g., loss of trust in online banking, reduced uptake of electronic services, remedial programmes).

Defence costs Monetary equivalent of prevention efforts and indirect costs of prevention systems.

Table 1: Judgement on coverage of cost categories by known estimates

Type of cybercrime	UK estimate	Global estimate	Reference period	Criminal revenue	Direct losses	Indirect losses	Defence cost
Cost of genuine cybercrime							
Online banking fraud							
– phishing	$16m	**$320m**	2007	$\times^{?}$	$\times^{?}$		
– malware (consumer)	$4m	**$70m**	2010	$\times^{\downarrow}$	$\times^{\downarrow}$		
– malware (businesses)	$6m	**$300m**		$\times^{\downarrow}$	$\times^{\downarrow}$		
– bank tech. countermeasures	$50m	**$1 000m**	2010				$\times^{?}$
Fake antivirus	$5m	**$97m**	2008–10	×	×		
Copyright-infringing software	**$1m**	**$22m**	2010	×	×		
Copyright-infringing music etc	$7m	**$150m**	2011	$\times^{\downarrow}$			
Patent-infringing pharma	**$14m**	**$288m**	2010	×			
Stranded traveller scam	$1m	**$10m**	2011	$\times^{\downarrow}$			
Fake escrow scam	$10m	**$200m**	2011	$\times^{\downarrow}$			
Advance-fee fraud	**$50m**	$1 000m	2011	$\times^{\downarrow}$			
...							
Cost of transitional cybercrime							
Online payment card fraud	**$210m**	$4 200m	2010		(×)		
Offline payment card fraud							
– domestic	**$106m**	$2 100m	2010		$\times^{\downarrow}$		
– international	**$147m**	$2 940m	2010		$\times^{\downarrow}$		
– bank/merchant defence costs	$120m	**$2 400m**	2010				$\times^{\downarrow}$
Indirect costs of payment fraud							
– loss of confidence (consumers)	**$700m**	$10 000m	2010			$\times^{?}$	
– loss of confidence (merchants)	**$1 600m**	$20 000m	2009			$\times^{?}$	
PABX fraud	**$185m**	**$4 960m**	2011	×	$\times^{\downarrow}$		
...							
Cost of cybercriminal infrastructure							
Expenditure on antivirus	**$170m**	$3 400m	2012				×
Cost to industry of patching	$50m	**$1 000m**	2010				$\times^{?}$
ISP clean-up expenditures	**$2m**	$40m	2010			$\times^{?}$	
Cost to users of clean-up	**$500m**	$10 000m	2012			$\times^{?}$	
Defence costs of firms generally	$500m	**$10 000m**	2010				$\times^{?}$
Expenditure on law enforcement	**$15m**	$400m	2010				×
...							
Cost of traditional crimes becoming 'cyber'							
Welfare fraud	**$1 900m**	$20 000m	2011	×	(×)		
Tax fraud	**$12 000m**	$125 000m	2011	$\times^{?}$	(×)		
Tax filing fraud	-	**$5 200m**	2010	×	(×)		
...							

Estimating costs and scaling: Figures in boldface are estimates based on data or assumption for the reference area. Unless both figures in a row are bold, the non-boldface figure has been scaled using the UK's share of world GDP unless otherwise stated in the main text. Extrapolations from UK numbers to the global scale should be interpreted with utmost caution. A threshold to enter this table is defined at $10m for the global estimate. **Legend**: × : included, (×) : partly covered; with qualifiers $\times^{\uparrow}$ for likely over-estimated, $\times^{\downarrow}$ for likely underestimated, and $\times^{?}$ for high uncertainty.

Takeaways

- Traditional tax/welfare fraud costs a few hundred per year/citizen, and defences and enforcement would be much cheaper.
- Payment card fraud costs tens per year/citizen. Defence costs are about the same, but indirect costs due to fear of fraud are several times higher.
- Fake antivirus, etc. net operators some tens of pence per year/citizen, but indirect costs and defence costs are an order of magnitude greater.
- Spend less on anticipating crime, more on catching and punishing perpetrators.

The economics of malware[3]

Malware has to be understood as a product not just of criminal actors, but a range of others:

- ISPs
- Software vendors
- Hardware manufacturers
- Domain registrars
- End users

Perspectives and security decisions of these players may be in their rational self-interest, but impose externalities.

---

[3] van Eeten, M. J., & Bauer, J. M. (2008). "Economics of malware: Security decisions, incentives and externalities." *OECD Science, Technology and Industry Working Papers*

Three situations

1. *No externalities*: Very rare due to nature of Internet. Example: 'good' end users who prevent their machines being compromised.
2. *Externalities borne by those who can manage them*: Someone else is handling this cost. Example: ISPs managing the security problems caused by their customers; financial services compensating for fraud outside of their remit.
3. *Externalities not being managed*: Prominently, costs to society of lax end-user security practices.

It's hard to improve end-user security

Malware is written to minimise its impact on the end-user[4]

The end-user therefore sees little benefit to investing in efforts to prevent malware, creating a large negative externality for other parties that have to suffer costs of malware.

Changing the perceived or actual costs of malware to end-users would be instrumental in altering this outcome.

[4] This has changed recently, as ransomware takes the exact opposite approach

Section 4

Next Week

Flipped Session

As a group (task allocation up to you), present (∼20 min.) the paper, covering

- What it's about, what's the point.
- What the method of analysis is, especially wrt. economics.
- What the results were, key takeaways.
- What the limitations of the paper/analysis/data are.

Also, a **citation exercise** (∼10 min., different for each group).

Group 1: Changing Costs of Cybercrime

Robert, Hannah & Tobias

Anderson, Ross, Chris Barton, Rainer Bölme, Richard Clayton, Carlos Ganán, Tom Grasso, Michael Levi, Tyler Moore, and Marie Vasek. "Measuring the changing cost of cybercrime." *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, 2019.

Citation exercise: For 5 academic citations *from the paper*:

- find the cited paper's full text
- check and briefly state what the paper is about
- explain how & why Anderson et al. cite the paper.
- is the usage of the citation in Anderson et al. a correct interpretation of this paper?

Group 2: Role of ISPs in Botnet Mitigation

Soo Yee, Priyanka & Manolis

van Eeten, Michel, Johannes M. Bauer, Hadi Asghari, Shirin Tabatabaie, and David Rand. "The role of internet service providers in botnet mitigation: An empirical analysis based on spam data." *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, 2010.

Citation exercise: For 5 papers *that cite* van Eeten et al.:

- find the citing paper's full text
- briefly explain what the paper is about
- explain how & why they cite van Eeten et al.
- is this citation usage a correct interpretation of van Eeten et al.?

Request Session

Suggestions:

- Economics of vulnerabilities
- Economics of privacy
- Cyber-insurance
- Porn/security ecosystem
- Attitudes to cybercrime
- Extended flipped session
- ...