Attack and defense

Simona Fabrizi[1] Steffen Lippert[2] José Rodrigues-Neto[3]

[1]Massey University
[2]University of Auckland
[3]Australian National University

2nd ATE Symposium
University of New South Wales Business School
December, 2014

This talk

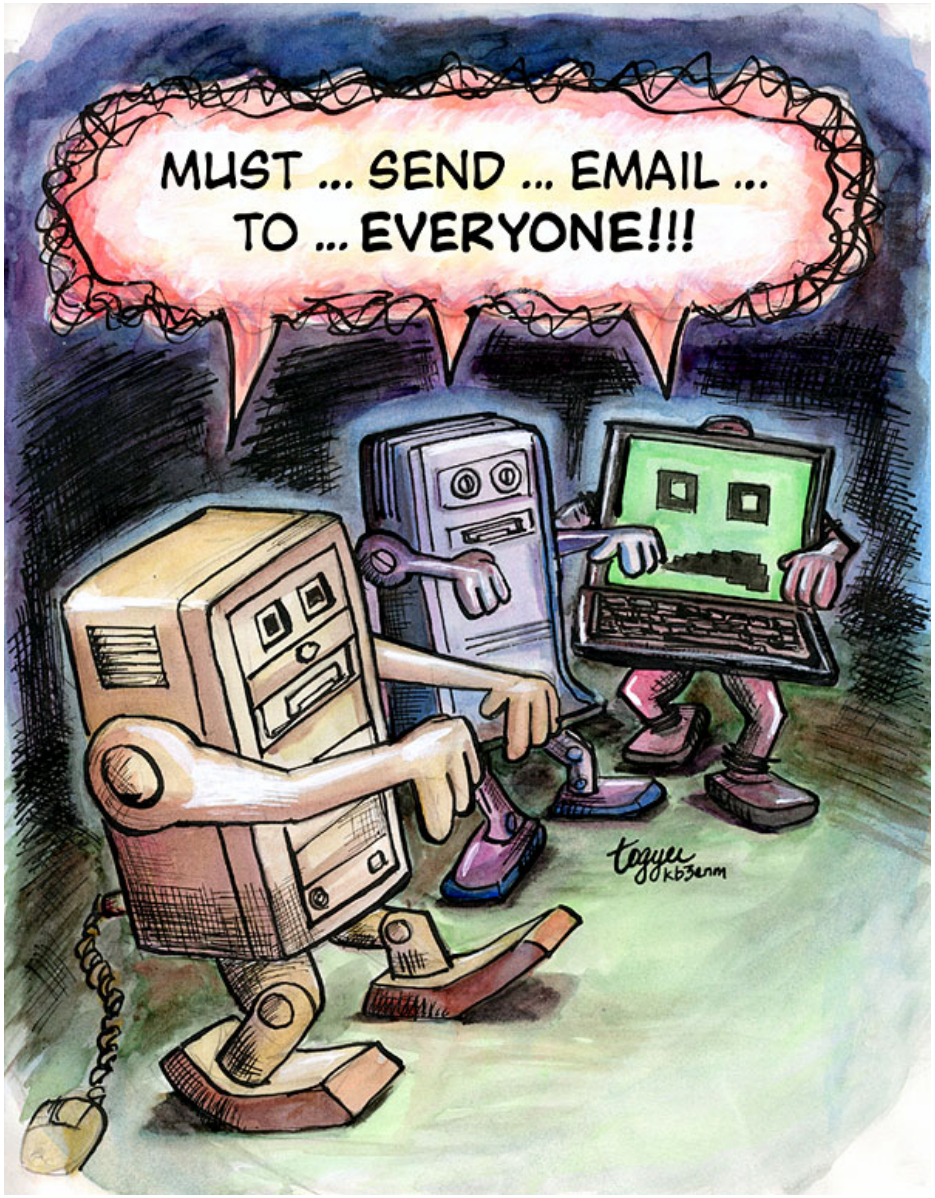

http://uvmzombies.blogspot.com.au/2013/02/computer-zombies.html

Botnets

- Sophisticated distributed systems comprising millions of computers with decentralized control.
  - → Network of "zombie" computers infected with malicious programs ("malware") that allows criminals ("botnet herders") to control the infected machines remotely without the users' knowledge.
- Used to
  - execute Distributed Denial of Service (DDoS) attacks.
  - harvest credit card information, personal data, financial information, email passwords, etc.
  - carry out phishing attacks, send out spam, carry out search engine spam, install adware, engage in click fraud.
- Sometimes they are leased out to others, who use them for the above causes.
- If you have a pulse, you're a target. Anybody's information has a value.
- Any, even "non-sensitive", information is valuable. Names, addresses, contacts can be monetized, e.g., sold for social phishing attacks.

Botnets

- There is a well-organized industry behind this with advertised prices for both outputs (e.g., credit card information) and inputs (e.g., malware-as-a-service).
  - Prices that depend on quality.
  - Try-before-you-buy offers.
  - Bulk offers.
  - "Google Analytics" for the bad guys, etc.
- Some organizations behind this are really big.

Example (Rock Phish)

- High-tech phishing. Practically undetectable & unblacklistable.
- Huge: Peter Gutmann (UoA) estimates US$0.5 – US$1B/year revenue.
- Scary: Joseph Menn writes about Rock Phish as organized crime, including kidnapping of anti-crime investigator's daughter.

Protection

- Some targeted at large institutions: Companies offer banks and other organizations likely to suffer from phishing attacks round-the-clock services to monitor, analyze, assist in shutting down phishing websites, or to implement two-factor authorization, which is being used increasingly.
- Some targeted at end-users: Spam filters target phishing email, firewalls, switches, routers.
- Properties of protection
  - It is privately costly to invest in protection.
  - There are positive externalities from investing in protection.
  - It affects the optimal choice of attackers.

Biological attacks

- Some features of malware attacks are present in biological attacks:
  - Contagion. Possibility to protect. Externalities of protection. Indirect effects through choices of attackers.

The project

- Try to understand more of the Economics underlying the malware economy, including the impact of market power.
- Build stylized models of attack and defense with heterogenous populations of defenders and attackers.

Model

Model

Populations

- Continuum of attackers in population I, mass $\mu > 0$. Choose whether to attack.
- Continuum of defenders in population J, unit mass. Decide whether to pay for protection against attacks or risk suffering loss from attack.
- Attack is successful if and only if the defender did not pay for protection.
- Attackers cannot observe whether defender has protection.

Attackers

- Attacker i obtains payoff of x_i from a successful direct attack.
- x_i is continuously, atomless distributed, CDF F_X, $F_X(+\infty) = 1$.
- Attacker i also obtains payoff of x_i from indirect attacks on all unprotected defenders his target is connected to during the attack.
- Abstract from exact process for now.

Attackers

- Utility of not attacking is

$$U_i(\text{no attack}) = 0.$$

- Let mass of defenders not taking protection be $\lambda \in [0, 1]$.
- Model expected utility of attacking as

$$U_i(\text{attack}) = \alpha(\lambda)x_i + \beta(\lambda),$$

- $\alpha(\lambda)$ positive and increasing; $-\beta(\lambda)$ positive and increasing.

Attackers

- Attack if
$$U_i(\text{attack}) = \alpha(\lambda)x_i + \beta(\lambda) > 0 = U_i(\text{no attack})$$
or
$$x_i > \frac{-\beta(\lambda)}{\alpha(\lambda)}.$$
- Increase in λ means fewer protected defenders, should make attack more profitable
$$\frac{d}{d\lambda}\left[\frac{-\beta(\lambda)}{\alpha(\lambda)}\right] \leq 0.$$
- Proportion of attackers choosing not to attack:
$$\chi = F_X\left(\frac{-\beta(\lambda)}{\alpha(\lambda)}\right).$$

Defenders

- Defenders have a choice between cost of protection and the chance of suffering a loss.
- Denote the loss if she is directly attacked and does not have protection by $S_j > 0$.
- S_j is continuously, atomless distributed, CDF F_S, $F_S(+\infty) = 1$.
- Cost of protection $c(\chi) > 0$ with $c'(\chi) \leq 0$.

Defenders

- Utility if invested in protection

$$V_j(\text{protection}) = -c(\chi).$$

- Fraction and mass of attackers that choose attack: $1 - \chi$ and $\mu(1 - \chi)$.
- Attackers do not target; there may be indirect attacks; abstract from exact process for now.
- Expected utility of an unprotected defender

$$V_j(\text{no protection}) = \delta(\chi)(-S_j).$$

- $\delta(\chi)$ is positive, decreasing, with $\delta(1) = 0$ and $\forall \chi \neq 1$, $\delta(\chi) \in]0, 1]$.

Defenders

- Invest in protection if

$$V_j(\text{protection}) = -c(\chi) > \delta(\chi)(-S_j) = V_j(\text{no protection})$$

or

$$S_j > \frac{c(\chi)}{\delta(\chi)}.$$

- Mass of unprotected defenders:

$$\lambda = F_L\left(\frac{c(\chi)}{\delta(\chi)}\right).$$

Equilibrium self-protection

Equilibrium self-protection

Proposition

Suppose that $c(\chi) > 0$, *for every* χ. *Suppose that* $\beta(\lambda) \neq 0$ *for all* λ, $\alpha(0) = 0$ *and* $\alpha(\lambda) > 0$ *for all* $\lambda > 0$. *Suppose* $F_X\left(\frac{-\beta(1)}{\alpha(1)}\right) < 1$. *Then, the game has a unique Nash equilibrium such that* $0 < \lambda^* < 1$ *and* $0 \leq \chi^* < 1$. *Moreover:*

$$\begin{aligned} \chi^* &= F_X\left(\frac{-\beta(\lambda^*)}{\alpha(\lambda^*)}\right), \\ \lambda^* &= F_S\left(\frac{c(\chi^*)}{\delta(\chi^*)}\right). \end{aligned}$$

- $\alpha(0) = 0$ means attackers cannot gain anything from attacking if all defenders are protected.
- $F_X\left(\frac{-\beta(1)}{\alpha(1)}\right) < 1$ means if no defenders protect, then there must be some active attackers.

Equilibrium self-protection

- $\alpha(0) = 0$ would not apply if protection was not perfect. Indeed, our formulation allows for less than full protection. Then c would combine the price paid for partial protection with the expected damage from successful attacks.
- Denote $\epsilon_c = \frac{dc}{d\chi}\frac{\chi}{c}$ and $\epsilon_\delta = \frac{d\delta}{d\chi}\frac{\chi}{\delta}$.

Proposition

Suppose that $c(\chi) > 0$ *for all* χ, *and* $\alpha(\lambda) \neq 0$ *and* $\beta(\lambda) \neq 0$ *for all* λ. *Suppose* $F_X\left(\frac{-\beta(1)}{\alpha(1)}\right) < 1$. *Then, the game has a Nash equilibrium. This Nash equilibrium is unique if* $\epsilon_c \geq \epsilon_\delta$, *for all* χ. *In this equilibrium, a proportion* χ^* *of attackers do not attack and a proportion* λ^* *of defenders do not pay for protection (as defined above). This equilibrium is such that* $0 < \lambda^* < 1$ *and* $0 \leq \chi^* < 1$.

Equilibrium self-protection

Inefficiency of equilibrium self-protection

- Marginal defender's choice to invest in protection lowers the mass of active attackers.

$\rightarrow$ Positive externality onto other unprotected defenders.

$\rightarrow$ If $dc/d\chi < 0$ also positive externality onto other protected defenders.

Market for protection

Market for protection

- Assume protection is sold at a price p.
- Given the above propositions, for any p, there exists a unique equilibrium with

$$\chi^* = F_X\left(\frac{-\beta(\lambda^*)}{\alpha(\lambda^*)}\right),$$
$$\lambda^* = F_S\left(\frac{p}{\delta(\chi^*)}\right).$$

- Demand for protection: $D(p) = 1 - \lambda^*$. Under reasonable assumptions on $\delta(\chi)$ and $\alpha(\lambda)$, $D(p)$ and χ^* are decreasing in p.

Market for protection

- Assume price-taking firms that provide protection to a measure q of defenders incur a cost $C(q, \chi^*) = qc(\chi^*)$.
- Then, the allocation in the competitive equilibrium coincides with that in self-protection.
- There is too little protection.

Welfare loss with market power

- A monopolist would incur $C(1-\lambda^*, \chi^*) = (1-\lambda^*)c(\chi^*)$.
- Price decrease has two effects

$$\frac{d}{dp}\left[(1-\lambda^*)c(\chi^*;\omega)\right] = \underbrace{(1-\lambda^*)\frac{d}{d\chi}[c(\chi^*)]\frac{d}{dp}[\chi^*]}_{\text{indirect effect, } >0} \underbrace{-c(\chi^*)\frac{d}{dp}[\lambda^*]}_{\text{direct effect, } <0}.$$

- Monopoly solution satisfies

$$-\frac{1}{\varepsilon_\lambda} = \frac{p - \left\{c(\chi^*) - (1-\lambda^*)\frac{d}{d\chi}[c(\chi^*)]f_X(\cdot)\frac{d}{d\lambda}\left[\frac{-\beta(\lambda^*)}{\alpha(\lambda^*)}\right]\right\}}{p}.$$

(1) Welfare loss from externality onto unprotected defenders is compounded by monopoly mark-up.

(2) Monopolist internalizes part of the externality: A decrease in the price increases demand and decreases attacks, making protection cheaper.

To do

- A lot.
- Networked defenders and attackers that can distinguish defenders with different connectivity.
  - Who will be attacked?
  - Who will be protected?
  - Relative size of externalities?
  - Pricing of protection?
  - Cross subsidies between defender groups?
- Organised attacker that hires subset of attackers at a price for a DDoS attack.
  - How does price for botnet rental depends on decentralized equilibrium?
  - How does it change with connectivity and the implied change in protection levels?
  - ...