

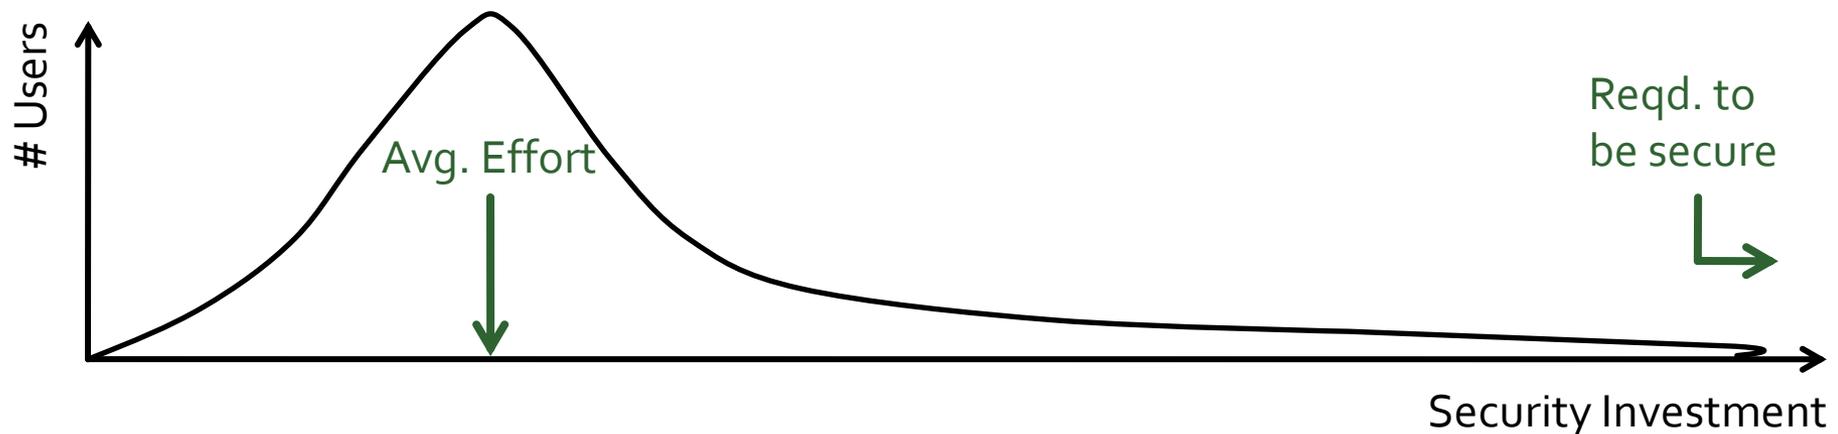
Cormac Herley

Microsoft Research, Redmond

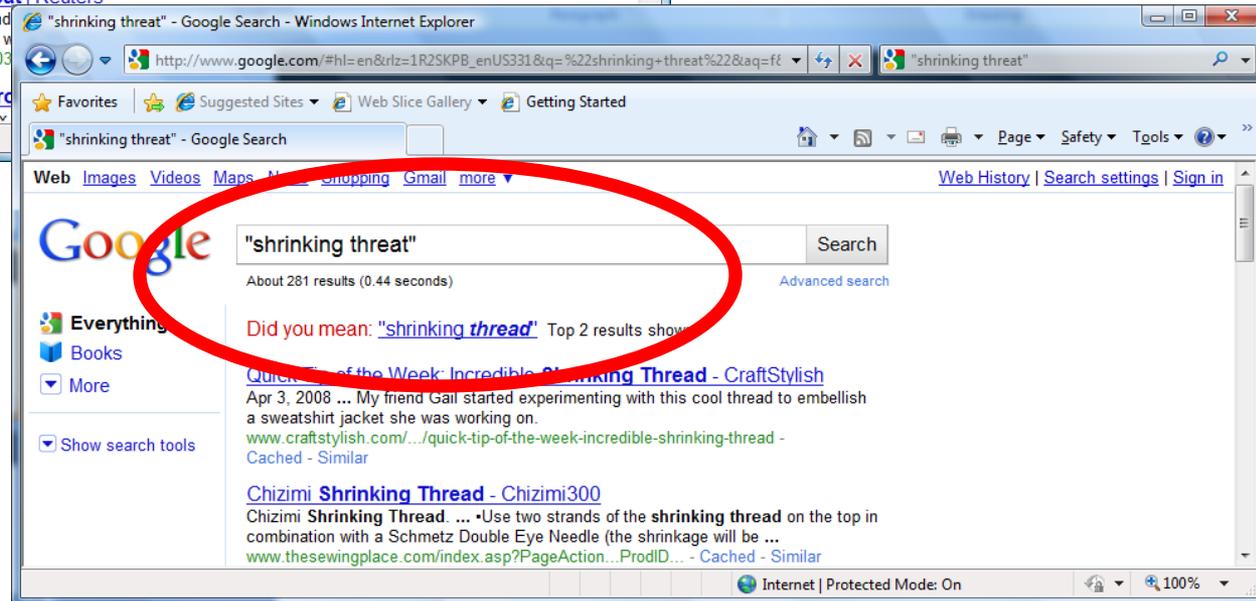
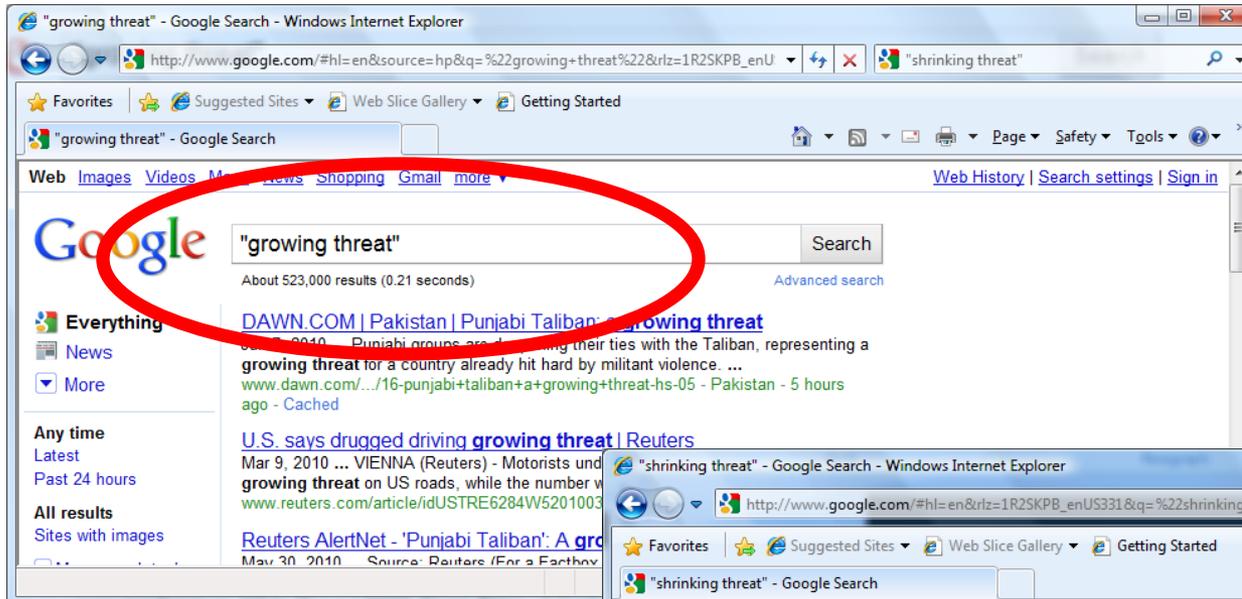
The Plight of the Targeted Attacker in a World of Scale

Puzzle: Where Do All the Attacks Go?

- **Observation: ~2 billion Internet users**
- **Most ignore most security investments**
 - Weak passwords, expired AV, password re-use, obvious secret questions,
- **Amazingly sophisticated attacks**
 - LCD screen reflections, hash collisions, realtime MITM
- **Life goes on. (Obla-di, Obla-da)**



New Threats Every Day



Common Threat Model



- Alice is an internet user
- Charles has ever-increasing number of attacks
- If Alice neglects any defense Charles wins

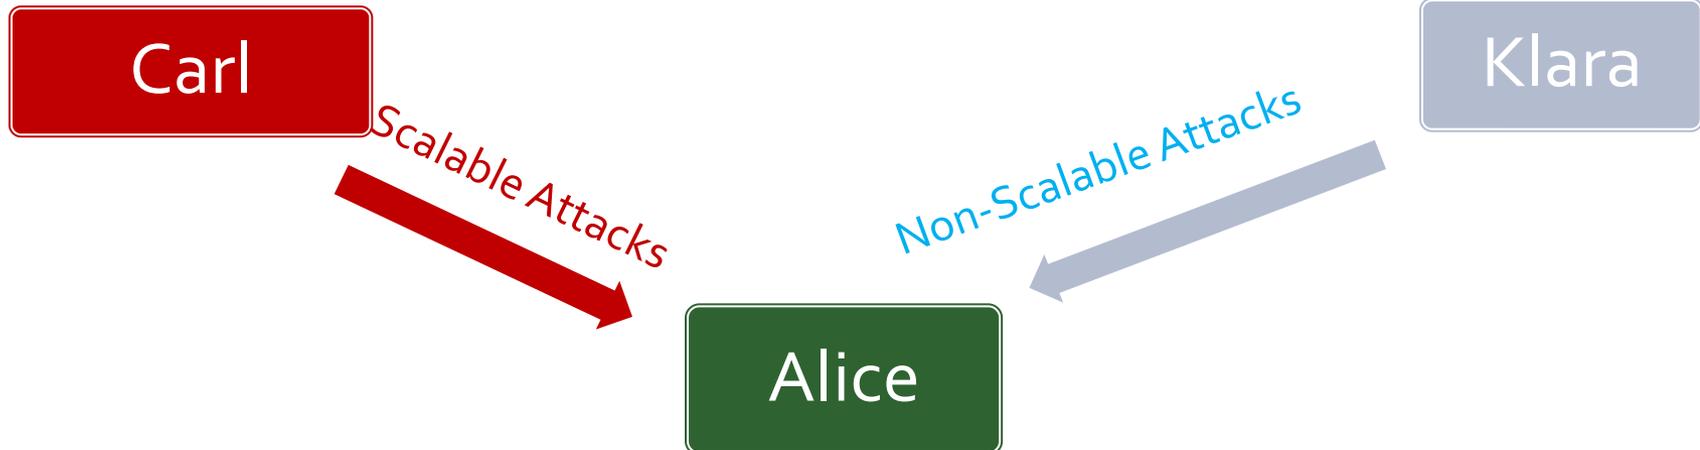
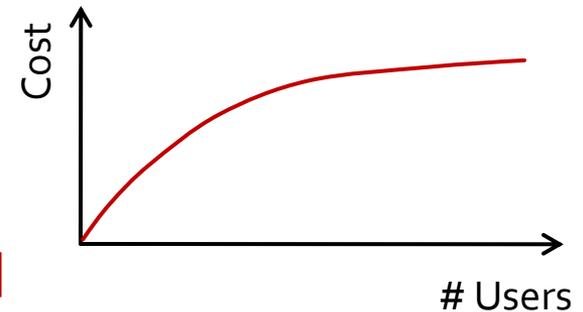
New Threat Model: Scalable Attacks

Carl: Scalable Attacks

- Sub-linear Cost Growth

$$C_s(2N) \ll 2C_s(N)$$

- E.g. spam, phishing, anything automated



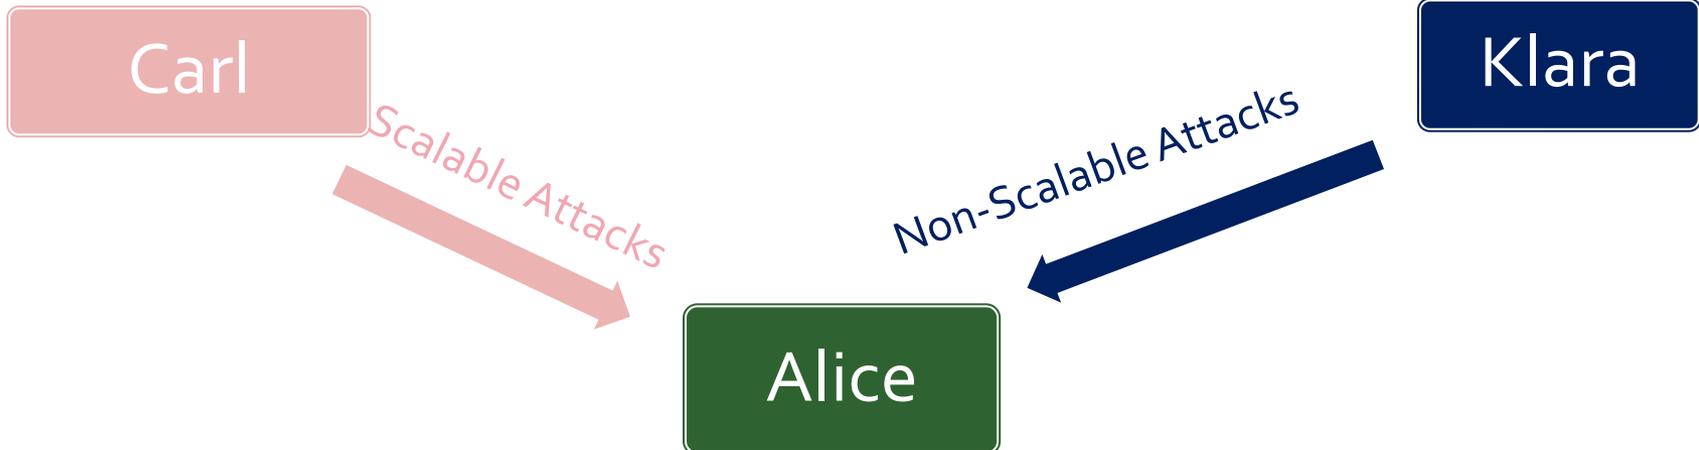
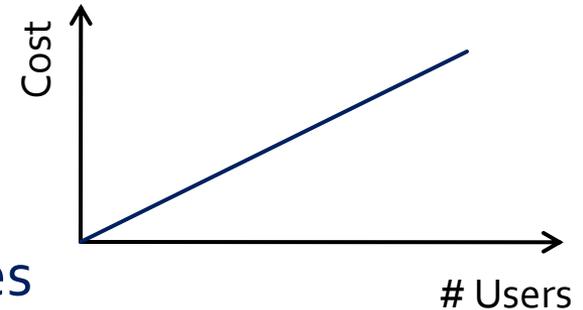
New Threat Model: Non-Scalable Attacks

Klara: Non-Scalable Attacks

- Linear (or worse) cost growth

$$C_n(2N) \approx 2 C_n(N)$$

- E.g. spear phishing, anything that involves per-user effort, knowledge of victim, proximity etc



Threat Model

- Two Attackers, two cost models
 - **Carl achieves economies of scale**
 - **Klara has per-user cost**
 - No loss of generality
- Rewards:
 - $\text{Reward}(N) = N Y V$
 - $N = \#$ attacked users
 - $Y = \text{Yield}$
 - $V = \text{Average } \textit{Extracted} \text{ value}$

1. Scalable Attacks Reach Many More Users (for same cost)

- **Scalable Attacks** : Profit improves with scale
 - $\text{Profit}_s(2N_s) = \text{Reward}_s(2N_s) - C_s(2N_s)$
 - $> 2 \text{Reward}_s(N_s) - 2 C_s(N_s)$
 - $> 2 \text{Profit}_s(N_s)$
 - Attack everyone, as often as possible
- **Non-scalable attacks**: profit constant w/ scale
 - $\text{Profit}_n(2N_n) \approx 2 \text{Profit}_n(N_n)$
 - Be selective

2. Scalable Attacks Produce Commodity Goods

- **Scripted => Anyone can do**
 - Commoditization
 - Tragedy of the Commons
- Competition drives $V_s \rightarrow 0$
- **Data:**
 - Spam: \$2800 for 350e6 emails [Kanich etal 2009]
 - Price of CCNs, creds falling [Symantec 2009]
 - Captcha Solving: [Motoyama etal 2010]

	Captcha/ 1000
2007	\$10.00
2008	\$1.50
2009	\$1.00
2010	\$0.75

How do Carl/Klara compete?

- Carl reaches many more users ($N_s \gg N_n$)
- Economies-of-scale businesses are tough on non-scaleable actors
- Klara should switch to scaleable strategy if she can't match Carl's return

Non-scalable vs Scalable

- **Reward(N) = N Y V**
 - N = Users Attacked
 - Y = Yield
 - V = Extracted Value/Successfully attacked user

- **At Equal cost to beat Scalable Return:**

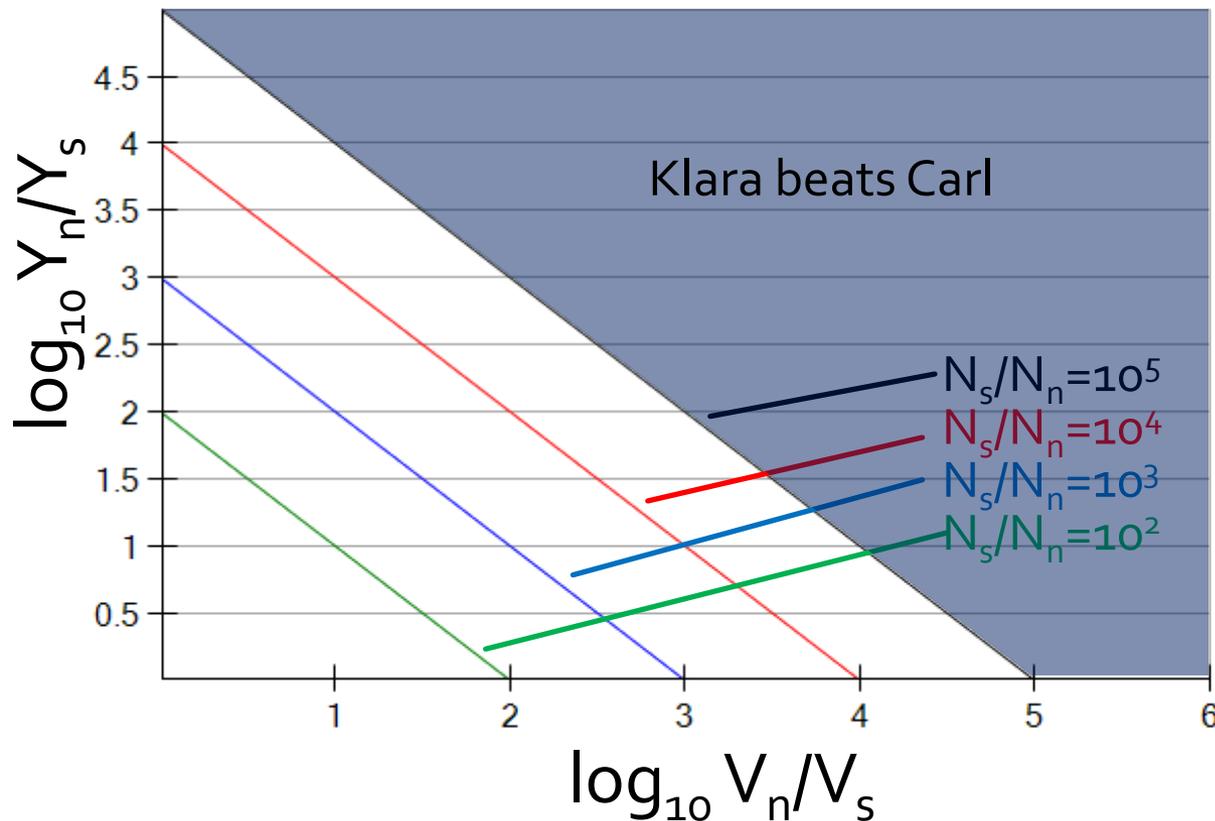
$$N_n Y_n V_n \geq N_s Y_s V_s$$

$$\Rightarrow \log(Y_n/Y_s) \geq \log(N_s/N_n) - \log(V_n/V_s)$$

Profit Frontier:

$$\log Y_n/Y_s \geq \log N_s/N_n - \log V_n/V_s$$

Non-Scalable needs: beat scaleable Yield-Value by as much as beaten on reach.



Competing on Yield Alone makes no sense

- $V_n = V_s$ then Klara competes on cost

- Klara now needs:

$$N_n Y_n \geq N_s Y_s$$

- Since $N_n \ll N_s$ this is hard:

- $Y_n \approx 4.5 Y_s$ [Jagatic et al Spear Phishing '06]

- Also, recall $V_s \rightarrow 0$ due to commoditization

- Reward decreases, but costs do not

- $V_n = V_s$ gives Klara difficult task

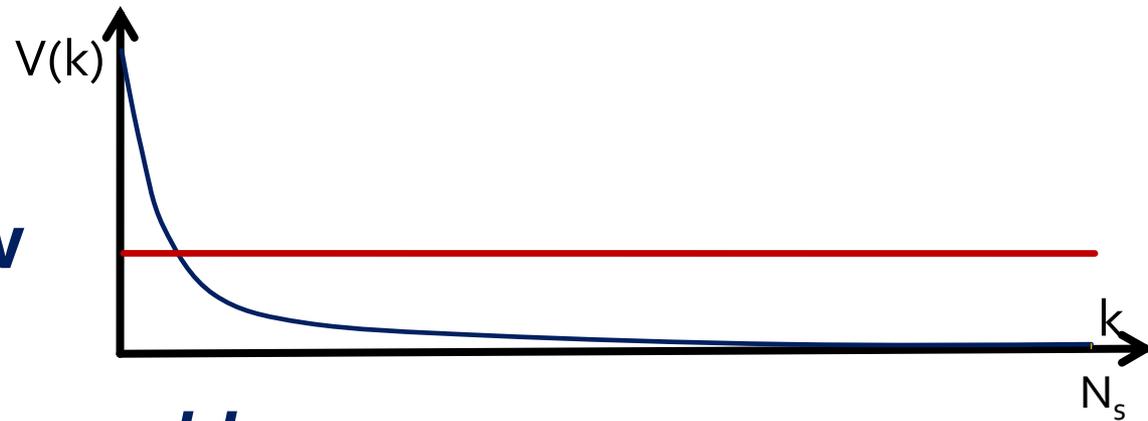
Seeking Higher-Value Targets

- Klara needs: $N_n Y_n V_n \geq N_s Y_s V_s$
- Since $N_n \ll N_s$ must have:
$$Y_n V_n \geq Y_s V_s$$
- So, higher yield, or higher value, or both
- Competing on Yield Alone Makes no sense
$$\Rightarrow V_n \geq V_s$$
- **Needs at least higher-than-average Value**

Klara needs longtail distribution of value

- At very least need $V(k) > V_s$
- Easiest when few users have high value, and most have low value

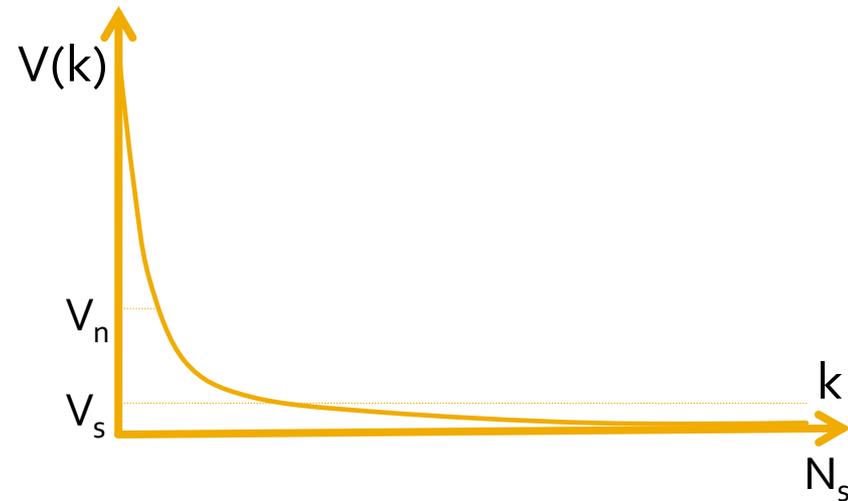
- **Worst: uniform**
- **Best: power-law**



- Must also be *observable*
 - Klara must be able to see who has high $V(k)$

In longtail distributions most Users have below average value

- Power-laws are everywhere
 - Wealth, fame, website popularity
- Mean \gg Median
 - Most users have $V(k) < V_s$
- Example concentrations:
 - US Wealth: 1.8% above avg.
 - Fame: 2% above avg.
- **98% of users worthless to Klara**
- **Attacking them hurts rather than helps.**
- True no matter how many Klara's there are



The Plight of the Targeted Attacker

- To equal Carl: $N_n Y_n V_n \geq N_s Y_s V_s$
- Competing with $V_n = V_s$ makes no sense
 - => Klara seeks high-value targets
 - => Klara needs longtail, observable distribution
 - => In longtails most users have $V(k) < V_s$
 - => Most users not attacked by Klara

On the Internet Nobody Knows You're Not a Dog

- Alice's Bank Backup auth questions can be determined with 1hr effort from facebook
- Acct yields \$200.
- Is this \$200/hr for Klara?

- No. Unless this always succeeds
- Klara's reward depends on:
 - Y = fraction of bank accts hackable from facebook
 - V = Average extracted value
- Alice's ~~security~~ avoidance of harm depends on
 - Worthlessness of average facebook account

What does Klara Attack?

- PC's for Zombie use?
 - Value as Zombie is close to uniform
 - Value of creds on box unobservable
- Email, social networking?
 - Sarah Palin's email, U East Anglia climate researchers
- Bank Creds?
 - Carl bulk-produces consumer creds
 - Small biz creds

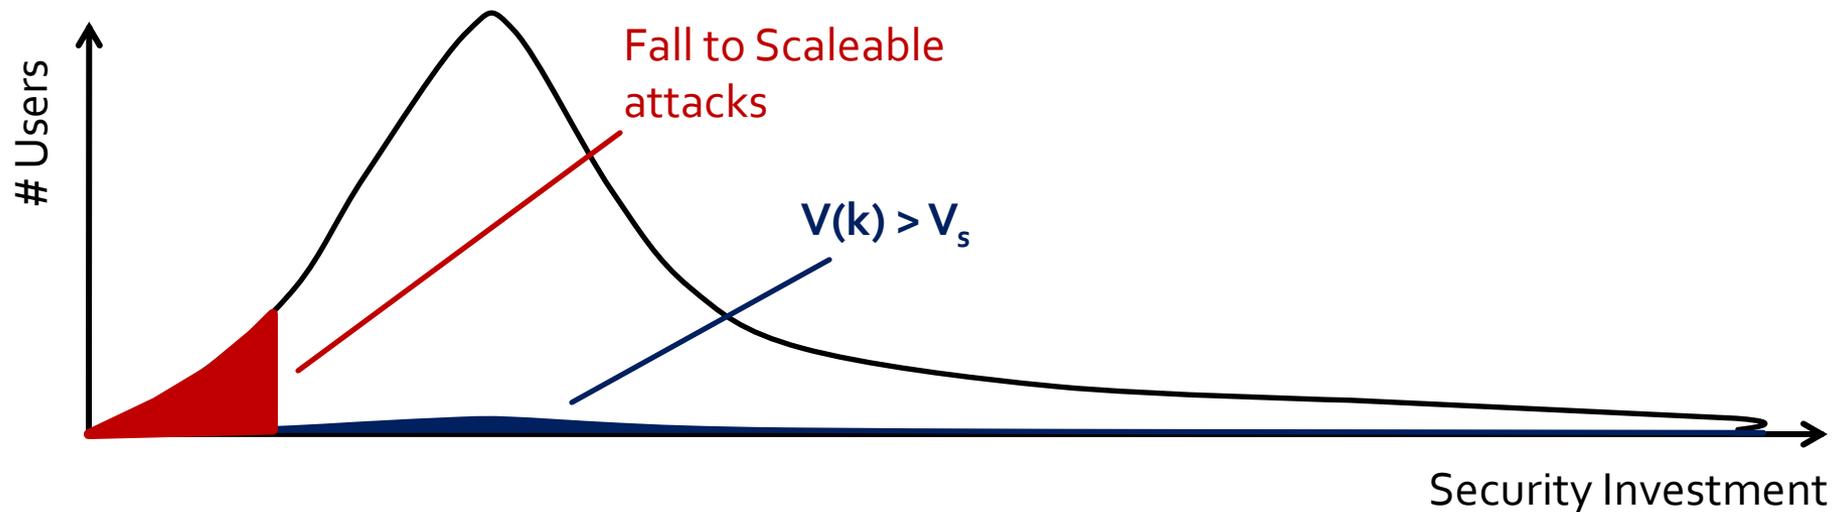
Concentrated/Observable

	Not Observable	Observable
Not Concentrated		Value Generic (PC for zombie, email for spam)
Concentrated	PC for credentials Sloppiness (Hi/Lo value acct. password sharing) Gullability* (responds to 419 scam)	Fame: (Sarah Palin's email) Closeness (jealous ex-SO)

*Gullability not observable. Nigerian 419 email is a scaleable attack which renders gullability observable. Carl/Klara cooperation

Security Investments

- Non-scaleable attacks are common, scaleable rare
- How much you must invest depends on whether anyone is targeting you



Conclusions

- How much should invest depends on targeting
 - Visibly in most valuable few percent for some asset?
- Elaborate non-scaleable attacks fail to happen
 - Benefit (to attacker) < Cost (to attacker)
- Most users never see most attacks