Security Economics:

From Game Theory to Field Measurements

Nicolas Christin
Carnegie Mellon University
nicolasc@cmu.edu

SIGMETRICS 2017 Tutorial

"Traditional" view of security

- Attackers are
  - Bound by computational and mathematical limitations…
  - … but by little else
  - High expertise assumed

(from my intro. security class)

Active attackers

- Or, what can Mallory do?
  - Can eavesdrop on all protocol runs
  - Can **replay** messages at will
  - Can **inject** fabricated messages in the network
    - For instance fabricated from pieces of old messages
  - Can **modify** a principal's message
  - Can **initiate multiple parallel protocol sessions**
  - Can perform **dictionary attack** on passwords
  - Can perform **exhaustive attack** on non-random (or poorly random) nonce

- Sound security engineering shouldn't rest on assumptions about possible attacker's weaknesses
- Likewise, defenders are assumed to be security-conscious

Security in practice

vs

Security in practice

vs

(Most) attackers in practice

- Most security attacks carried out by entities that are
  - Financially interested
  - Economically rational
  - Not necessarily overly sophisticated

  - Heavily reliant on commoditization
    - Purchase "services" from others

(Exception: the (still fairly rare) nation-state actors that are not outsourcing to criminals)

Defenders in practice

- Defenders have been assumed
  - knowledgeable
  - interested in security, and
  - altruistic
- But in practice,
  - Generally self-interested
  - Rarely fully informed
  - Subject to behavioral biases

Research agenda

- **Incentives and biases of both attackers and targets are critical to improving online security**
  - Useful to find where to target attackers
  - Useful to find how to deploy defenses
- How to discover and model these incentives?
- Assortment of different techniques
  - **Game theory**
  - **Behavioral economics**
  - **Network measurements**

Part I: Modeling defenders

Game-theoretic analysis

Related papers

1. Grossklags, Christin, and Chuang. Secure or Insure: A Game-Theoretic Analysis of Information Security Games. *WWW 2008*
2. Grossklags, Christin, and Chuang. Security and Insurance Management in Networks with Heterogeneous Agents. *EC 2008*
3. Grossklags, Johnson, Christin, and Chuang. When Information Improves Information Security. *FC 2010*
4. Johnson, Grossklags, Christin, and Chuang. Are Security Experts Useful? Bayesian Nash Equilibria for Network Security Games with Limited Information. *ESORICS 2010*

Homo economicus

- Assuming a rational, self-interested agent
  - Rational: Individuals are able to estimate the benefits and costs of a particular action (i.e., are able to estimate the net benefit)
  - Self-interested: Agents engage in an activity if the benefit is greater than or equal to the cost (i.e., the net benefit is greater than or equal to zero)

"It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest." (Adam Smith, The Wealth of Nations, 1776)

Network effects and externalities

- Terms often used interchangeably
  - Effects: Benefit, or cost, that an agent derives from a good when the number of other agents consuming the same kind of good changes
  - Externality: Participants in the market fail to internalize these effects
- Relationship to public goods
  - An externality occurs when a decision causes costs or benefits to third party stakeholders, often, although not necessarily, from the use of a public good
  - E.g., is identifying information and shopping data a public good?

Networking, security and economics

- Problem well known in economics and game theory

Economics	Networks	Security
Rational players competing in a market	Selfish nodes competing for network resources	Selfish agents whose security impacts others

- Can use game theory as a tool
  - to determine likely user (nodes) strategies given the context (network topology, network protocols, policies)
  - to design mechanisms (network topology, protocols, policies) which yield desirable strategies

Game-theoretic model overview

- Set of players in a network
- Utility function: value each player extracts from the network
  - Given by a cost model
- Strategies: Actions each player can take
- Equilibrium concept: situation where all players are content with their utility and don't change their strategy

Important equilibria concepts

- **Social optimum**: set of strategies that maximizes total $U = \sum_{i=1}^{N} U_i$ network utility
  - Ideal configuration for the community
  - What a benevolent government would want
- **Nash equilibrium**: set of strategies in which no individual player can increase their individual utility U_i by changing their strategy
  - Selfish equilibrium
  - Best response to others' actions

Limitations in a security context

- Asymmetric games: attackers vs. targets
  - Different motivations, utility functions…
- Incomplete information: Are parameters of model known to agents?
  - Can attackers infer defense posture?
  - Can defenders predict likelihood of attack?
- Information asymmetry: Does one party know more about parameters of interaction?
  - Lemons market: Are security products of high or low quality?

Network security games

- Variety of security threats and responses
  - Model most security interactions met in practice
  - Finite number of canonical security games
- Decouple security strategies
  - Self-protection investments (e.g., setting up a firewall)
  - Self-insurance coverage (e.g., archiving data as back up)
- Consider **network externalities**
  - Choice of strategy by a network participant affects other participants

General defender utility model
[GCC, WWW'08]

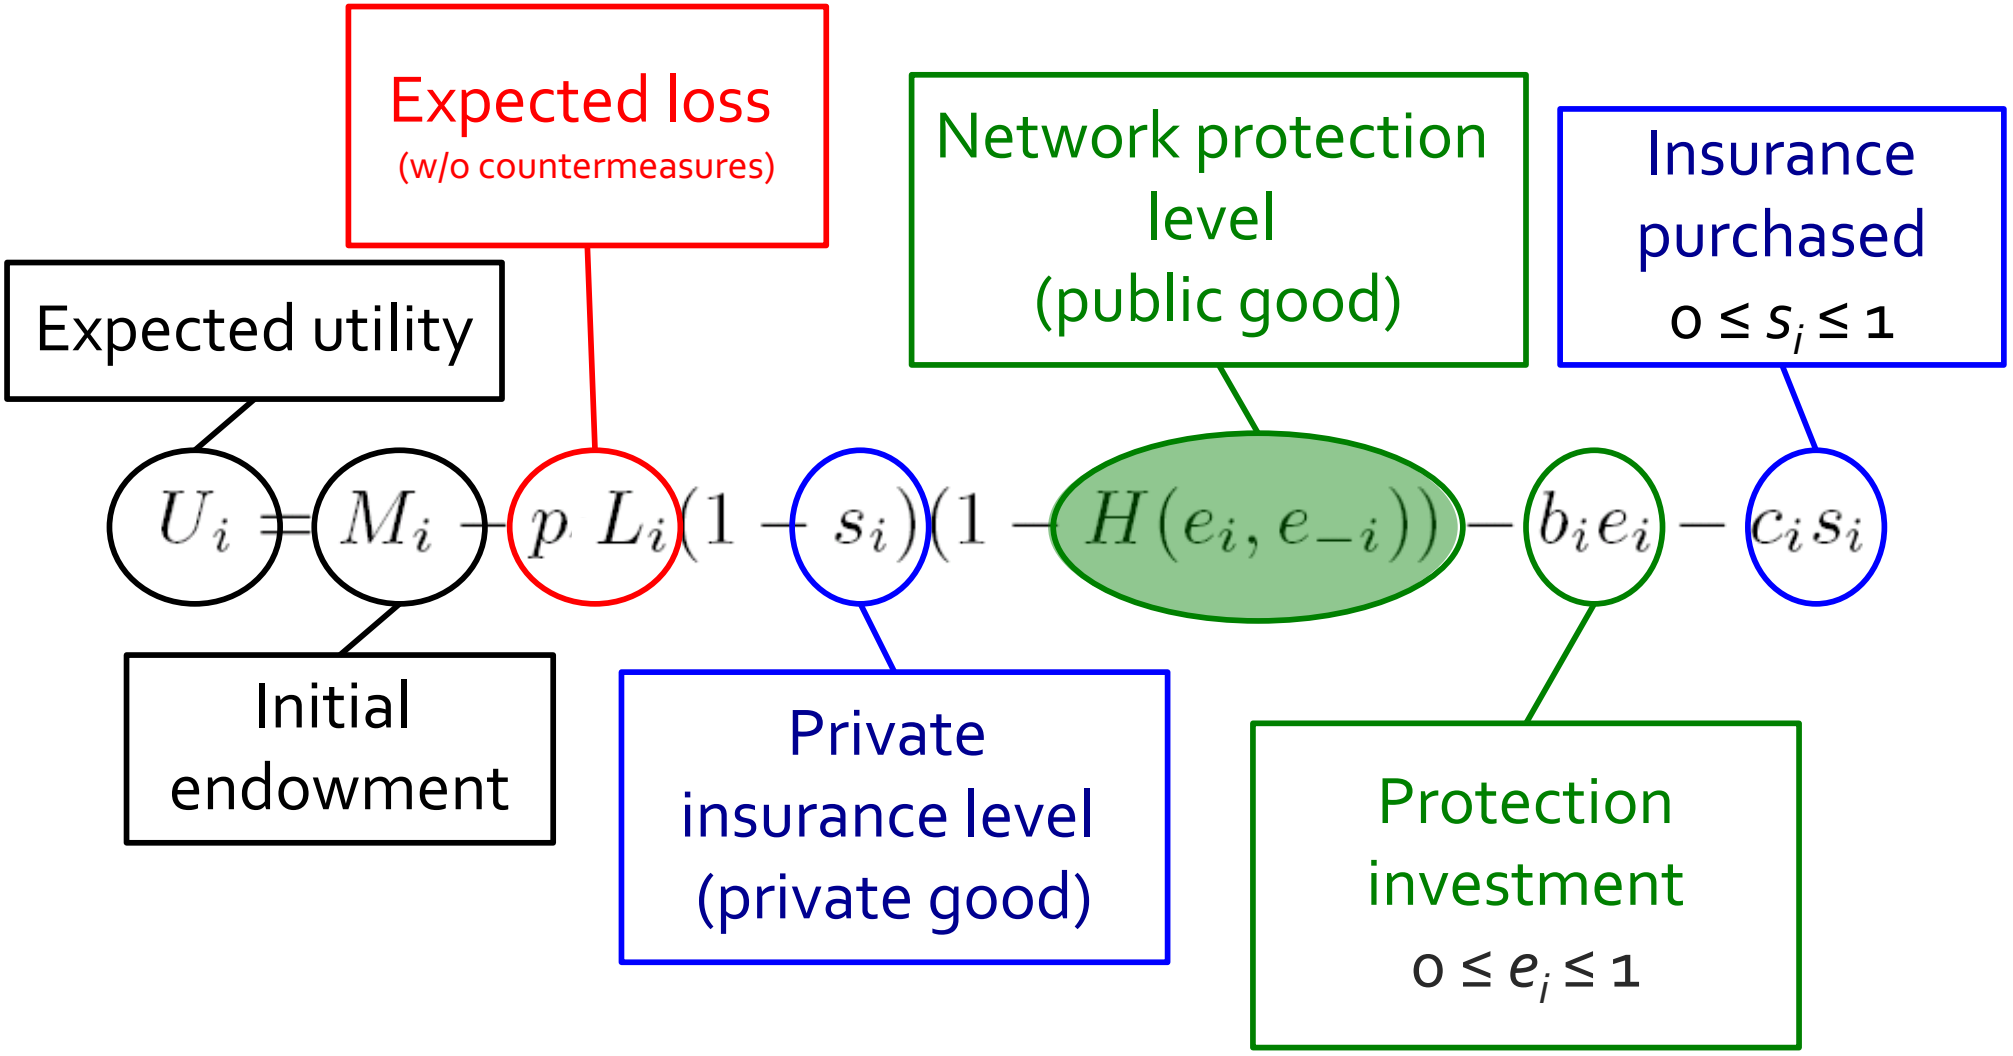

Contribution functions
(or: how is the network protected)

- Tightly coupled networks
  - Total/average effort $H(e_i, e_{-i}) = \frac{1}{N}\sum_i e_i$
    - Example: Distributed transfer of a file on a p2p network
  - Weakest-link $H(e_i, e_{-i}) = \min(e_i, e_{-i})$
    - Example: Corporate network penetration
  - Best shot $H(e_i, e_{-i}) = \max(e_i, e_{-i})$
    - Example: Censorship resilient networks
- Loosely coupled networks
  - Weakest target $H(e_i, e_{-i}) = \begin{cases} 0 & \text{if } e_i = \min(e_i, e_{-i}) \\ 1 & \text{otherwise} \end{cases}$
    - Example: Potential bots
  - Mitigated variant of the weakest link

Intuition behind Nash equilibrium outcome

- 3 types of pure Nash equilibria in our games
  - Protection only $(e_i, s_i) = (e^0, 0)$ (w/ $e^0=1$ fairly common)
  - Insurance only $(e_i, s_i) = (0, 1)$
  - Inactivity $(e_i, s_i) = (0, 0)$
- Increasing network size *N* affects Nash existence/nature

Summary of homogeneous results
($L_i = L$, $b_i = b$, $c_i = c$, $M_i = M$, pure Nash)

	Protection	Self-Insurance
Total Effort	$pL > \mathbf{bN}$ and $c > b + pL(N-1)/N$	*Other cases with* $pL > bN$ *or* $c < pL < bN$
Weakest Link	***Multiple*** symmetric protection equilibria	pL > c and too high protection cost
Best Shot	***No symmetric Nash***	Does exist if $b > c$ and $pL > c$
Weakest T w/o M	***No Nash***	*No Nash*
Weakest T with M	*Full protection if* b ≤ c	*No Nash*

Role of a social planner

- To achieve a social optimum
  - Sum of all players' utilities is maximized
  - Benevolent dictator
- Total effort:
  - More self-protection eq. ($pL > \boldsymbol{b}$ ~~$\boldsymbol{M}$~~$\boldsymbol{)}$)
- Weakest-link:
  - Planner would choose highest protection level
  - Pareto-optimal
- Best shot:
  - Planner now selects full protection for exactly **one** individual
  - In Nash eq. individuals frequently failed to protect
  - Insurance not needed
- Weakest target:
  - Sacrificial lamb
    - E.g., Honeypot
  - With or without insurance

Limited knowledge & information

[JGCC, ESORICS'10]

- Expert vs naive players
  - Expert players know the contribution function H and understand its effects.
  - Naive players are myopic; they behave as if $H(e_1, \ldots, e_n) = e_i$
- Complete vs incomplete information
  - An expert with complete information knows the expected losses for all players.
  - An expert with incomplete information knows her own expected loss L_i but does not know the expected losses of other players.
  - Experts assume that expected losses are independently and uniformly distributed in [0,1].

Best Shot and limited information

- In the Best Shot game, experts have a strong incentive to free-ride (Tragedy of the commons). Adding experts decreases the likelihood that the network is protected.

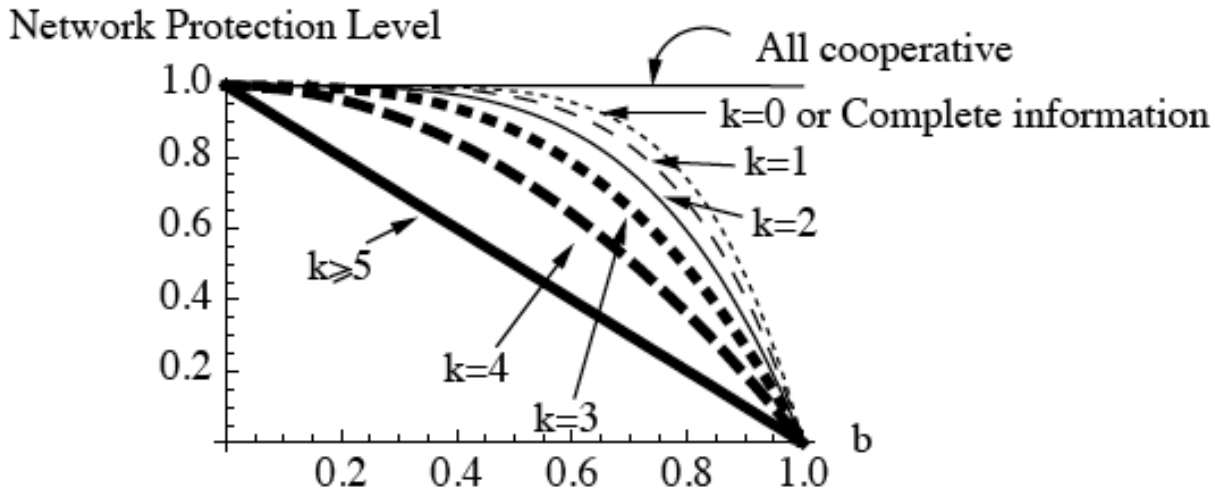

Fig. 2. Best shot. Evolution of the network protection level as a function of the protection cost b. The different plots vary the number of experts k in a network of $N = 6$ players. We observe that the fewer experts participating in the game, the higher the network protection level is, on average.

Weakest Link and limited information

- Protection equilibria in the Weakest Link game only exist when protection costs are small; and the problem is exacerbated by the addition of expert players.

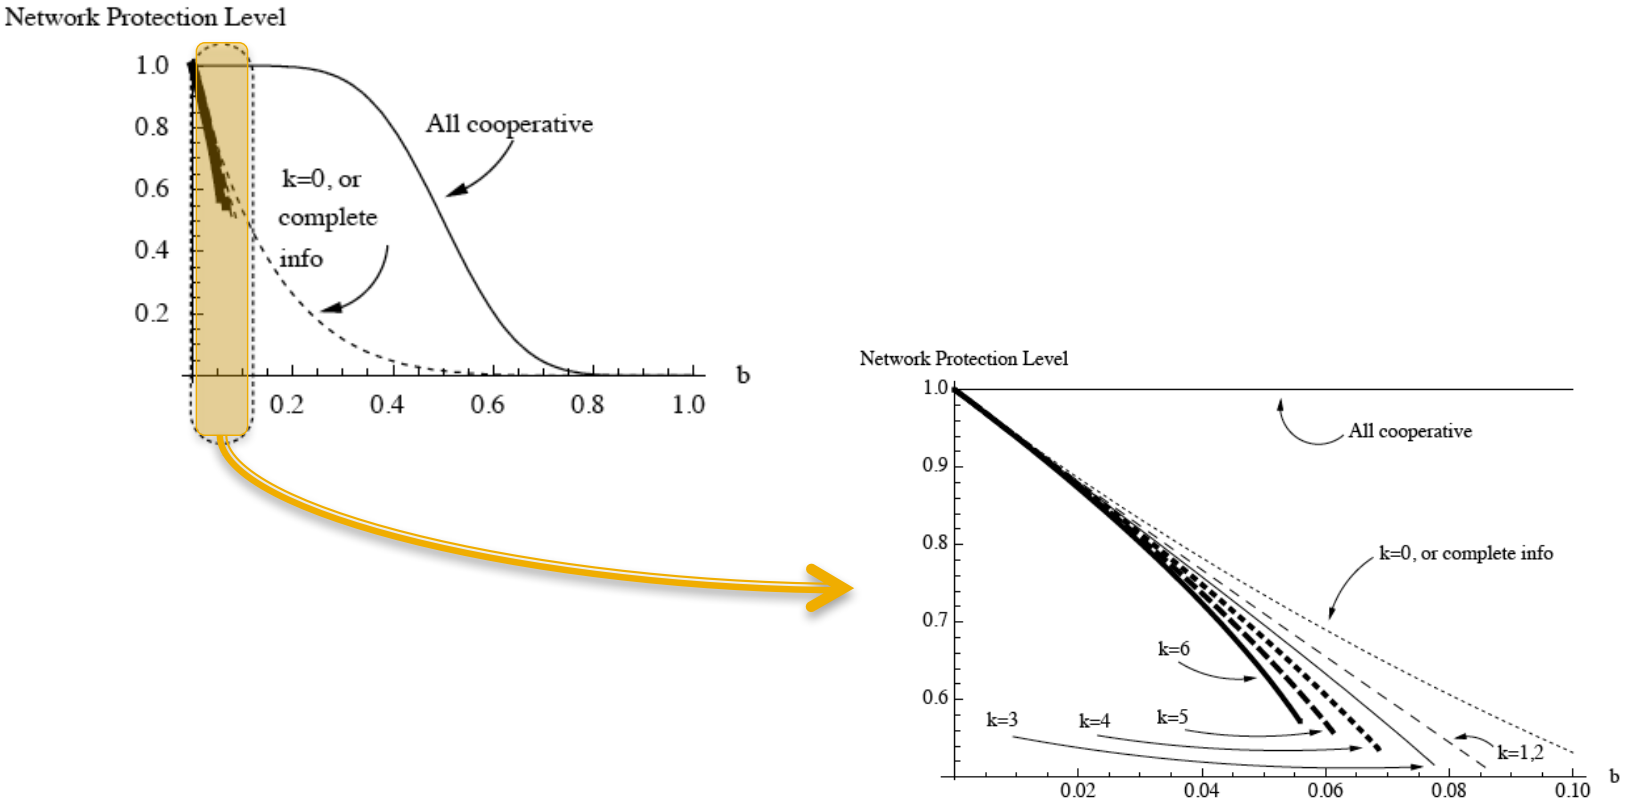

(b) Zoom on small values of the protection cost b .

Total Effort and limited information

- In the Total Effort game, the individual benefit of an investment is always proportional to a 1/N fraction of the investment's cost, regardless of the actions of other players. Experts understand this feature and do not protect very often.

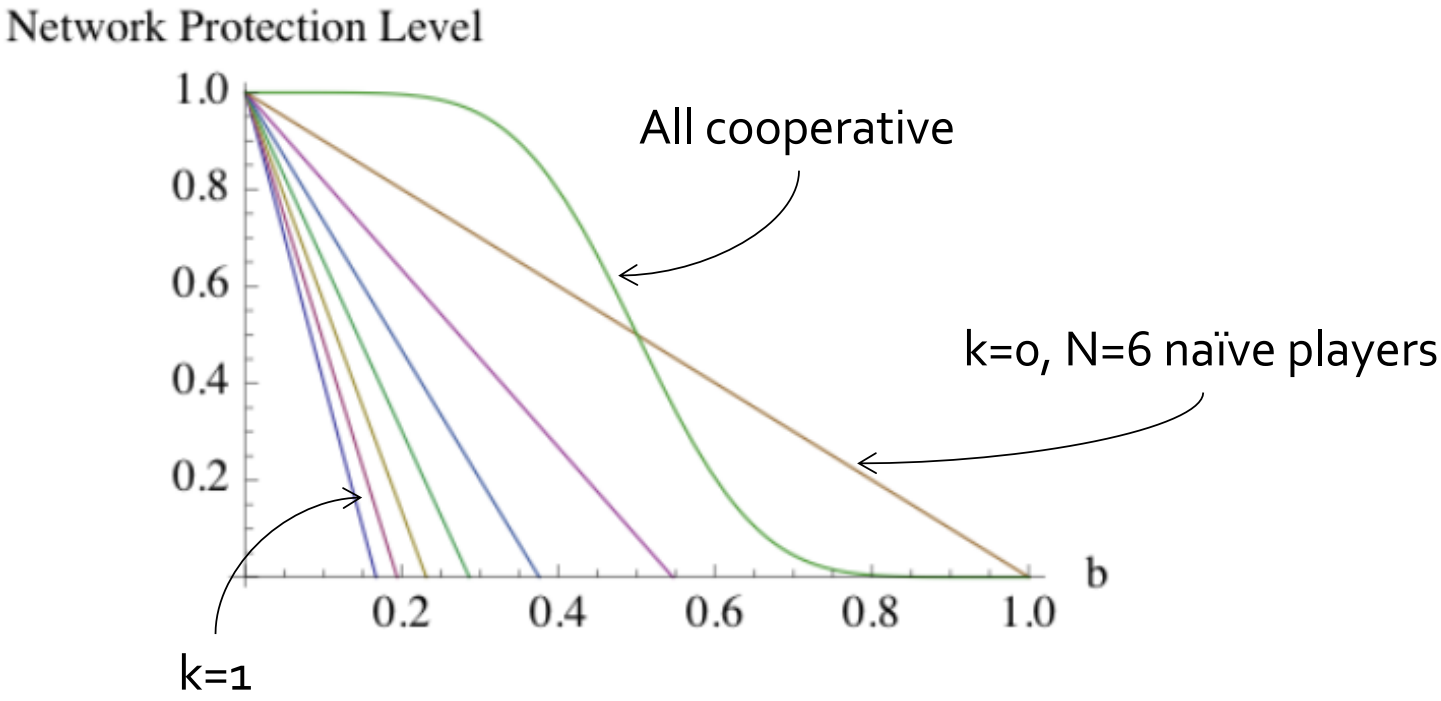

Implications

- (In some contexts), security experts are useful when (and only when) they collaborate.
- When security is divided among independent agencies, it is important to develop mechanisms for facilitating interagency collaboration.
- User education should focus on the collaborative nature of security

Summary

- Incentives of defenders matter
  - Negative externalities
    - Self-insurance actually makes self-protection less appealing
  - Heterogeneity, limited information, limited expertise do not change the overall picture much
  - Security expertise can actually make things worse
    - Expert users understand negative externalities
- But… how close to rational are users in practice?

Part II: The influence of behavioral economics

Understanding near-rationality

Related papers

1. Acquisti and Grossklags. Privacy and Rationality in Individual Decision-Making. *IEEE S&P Magazine 2004*

2. Kahneman and Tversky. Choices, Values and Frames. *American Psychologist 1984*

3. Grossklags and Acquisti. When 25 Cents is too much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. *WEIS 2007*

4. Christin, Egelman, Grossklags, and Vidas. It's all about the Benjamins: An Empirical Study on Incentivizing Users to Ignore Security Advice. *FC 2011*

Perceptions of risk

- Suppose I give you two alternatives:
  - I give you $1 or
  - I flip a fair coin; if it's tails, you get $2; heads, you get nothing
- What do you choose?

Perceptions of risk

- Suppose I give you two alternatives:
  - I give you $10,000,000 or
  - I flip a fair coin; if it's tails, you get $20,000,000; heads, you get nothing
- What do you choose?

Perceptions of risk

- Suppose I give you two alternatives:
  - I give you $10,000,000 or
  - I flip a fair coin; if it's tails, you get $50,000,000; heads, you get nothing
- What do you choose?

Previous findings

- Humans generally **risk-averse** when it comes to gains (prefer fixed payoff even though reward potentially less)
  - Risk aversion generally increases with the amount at stake (relative to the endowment)
- Humans generally **risk-seeking** when it comes to losses
  - Asian flu experiment (Kahneman and Tversky)

    *Two courses of action have been suggested. If program A is adopted**, 400 will die**. If program B is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. Which of the two programs do you favor?*

Free Giveaway!

Name: _____

Address: _____

City: _____ State: _____ Zip: _____

Home Phone: _____

Work Phone: _____

❑ Single ❑ Married

Age: _____ Occupation: _____

Spouses Age: _____ Occupation: _____

Combined Income:

❑ Under $30,000 ❑ Over $30,000 ❑ Over $50,000

DO YOU: ❑ RENT OR ❑ OWN YOUR HOME?

❑ VISA ❑ MASTERCARD ❑ AMERICAN EXPRESS

Would you like info on special events & promotions at Pier 39?

❑ Yes ❑ No

E-mail address: _____

Details of Participation and Eligibility Requirements

- Only one Entry per Family.
- Winner allows the use of his or her name, photo, and statements for future promotional use without further compensation.
- Winner must be 18 or over. I.D. required. Winner must provide all necessary federal and state tax reporting information before receiving prizes.
- Drawing held February 23, 2003. Last date to enter drawing is February 16, 2003.
- Winner need not be present to win. Winner will be notified by phone.
- Drawing will be conducted by a Certified Public Accounting Firm at the corporate office of Grand Pacific Resorts, 5900 Pasteur Ct., #200, Carlsbad, CA 92008. To request winner information, correspondence may be forwarded to Grand Pacific Resorts, Promotions Dept. P.O. Box 4068, Carlsbad, CA 92018.
- All local, state, and federal taxes, fees and licenses are the winner's responsibility. Acceptance of the prizes constitutes a release of Facility Management, it's agents and employees from all responsibility to the winner.
- Odds are based on number of entries received, approximately 1 in 700,000.
- No purchase or attendance is necessary to be entered into the drawing. Entrants may be invited to attend a sales presentation about the Red Wolf Lodge at Squaw Valley.
- Entries become the property of PNR Marketing Inc.
- The annual "Grand Prize" Giveaway consists of any vehicle with a retail value not to exceed $25,000 or a three year lease (value to $25,000) on a luxury car; or any prize (or similar) displayed in a Grand Pacific Resorts Promotion February 25, 2002 - February 23, 2003 (valued up to $15,000), or the winner may choose cash in the amount of $15,000.

What can the individual infer?

- Benefits:
  - Non-monetary benefit (e.g., excitement of participation)
  - Expected monetary benefit:
    - *1/700000 * $15000 = 2 cent*
- Costs:
  - Promotions, unsolicited mailing, sales contacts (cannot exclude further use and consequences)
  - Expected monetary cost:
    - *?*
- *What behavioral variables are missing?*

Psychological biases

- We react differently depending on framing of messages
- We make time-inconsistent decisions
- We seek immediate gratification
- We are susceptible to strong biases with ambiguous and unknown information

Paying people to install malware
[Christin, Egelman, Vidas, Grossklags, FC 2011]

- We paid people to download and run an unknown executable
- Payment was increased every week
  - Log scale
  - $0.01/$0.05/$0.10/$0.50/$1.00
- Mechanical Turk as experimental platform
  - Measured views vs. downloads vs. runs

Experimental environment

- CMU Distributed Computing Project
  - No such project exists
  - All code was hosted on a third-party domain
  - No connection to us or our institutions

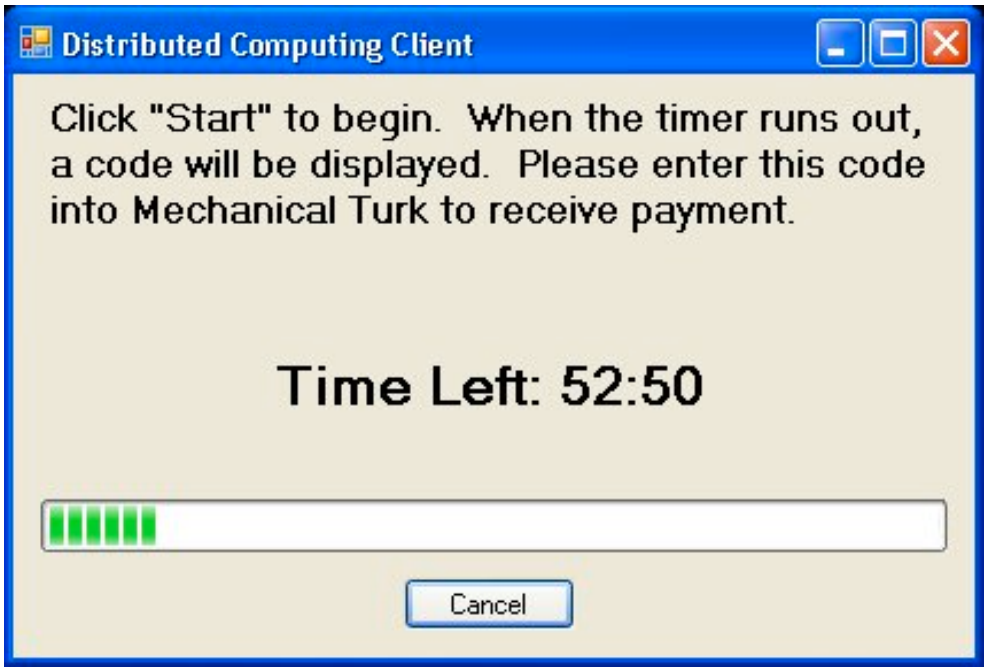

Experimental Environment

- Are current mitigations effective?

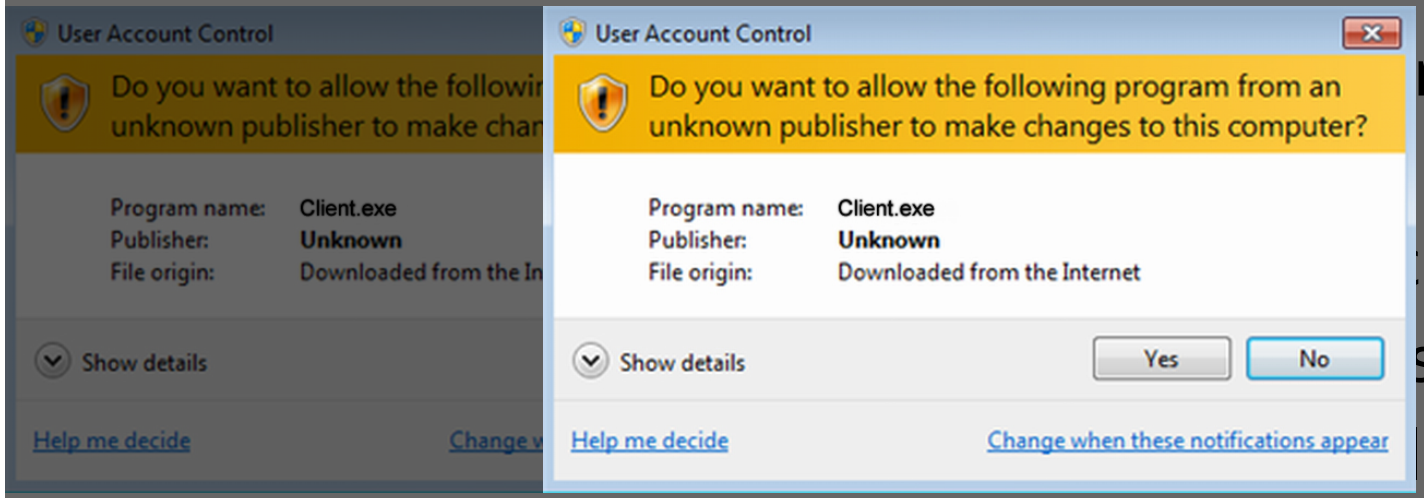

mpt for 50%
data:
ontrol
s version
ist
  - VM detection
- Displayed payment code
- Sent an exit survey

Results

	$0.01		$0.05		$0.10		$0.50		$1.00	
Viewed	291		272		363		823		1,105	
Downloaded	141	49%	135	50%	190	52%	**510**	**62%**	**738**	**67%**
Executed	64	22%	60	22%	73	20%	**294**	**36%**	**474**	**43%**

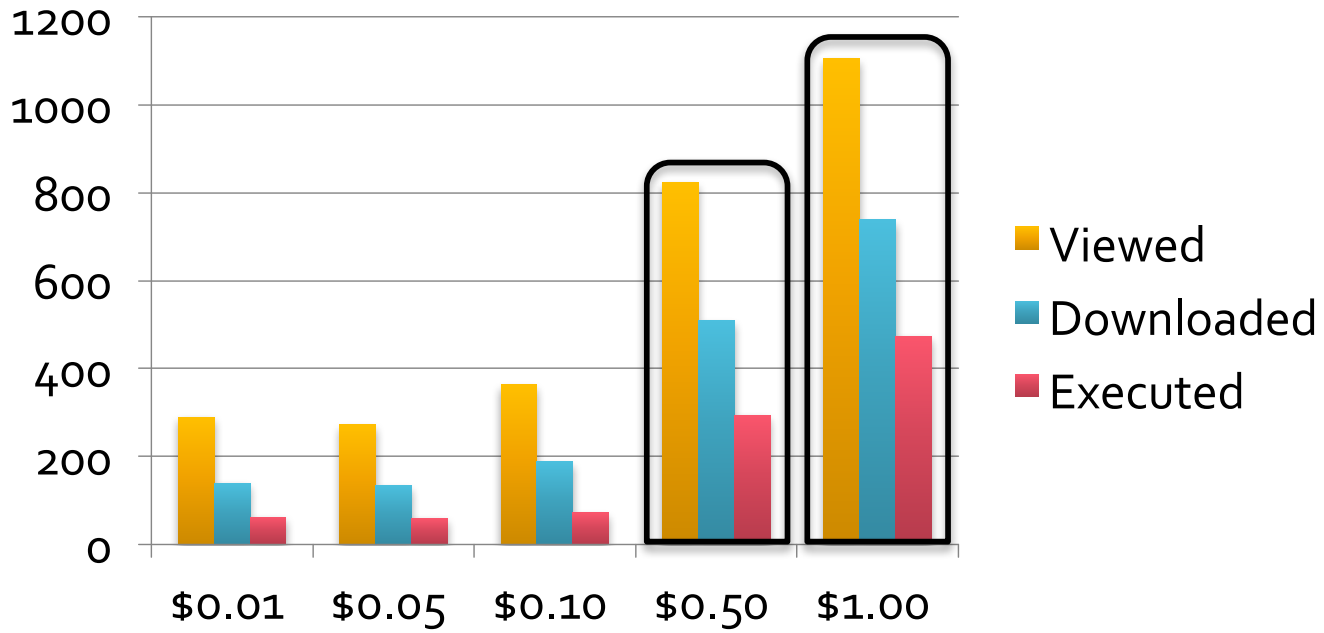

UAC was ineffective

- 501 users had either Vista or Windows 7
- Conditions randomly assigned
  - $X^2_1=0.449$, $p<0.503$

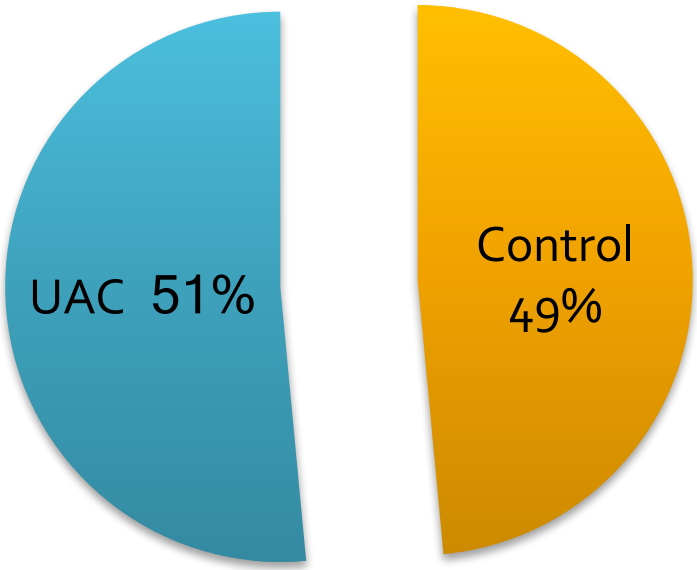

Security behaviors

- 17 participants used a VM (1.8% of 965)
- We categorized 3,110 unique processes
  - 16.4% had malware
  - 79.4% had security software
  - Correlation between malware/security software
    - $\Phi=0.066$, $p<0.039$

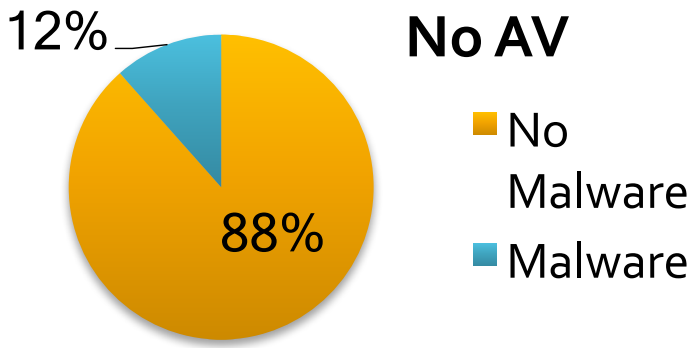

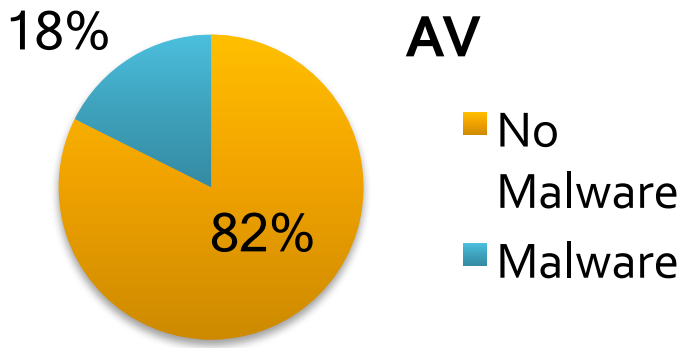

Validating behaviors

- Exit survey for a $0.50 bonus payment
- 513 people responded
  - 40% from India
  - 30% from US/Canada
  - Percentage from the developed world increased with price, 9.4% to 23.4%

Security perceptions

- Danger of running code from MTurk on a 5-point scale
  - $F_{4,508} = 3.165$, $p < 0.014$
  - People who *should* have known better participated once the price was right
- 70% of participants knew it was dangerous to download unknown programs
  - All of them did so anyway

What is rational?

- Peltzman effect
  - Availability of seatbelts leads to more risky driving
- Same effect observed here
  - Installation of security software does not limit risky behaviors, far from it!

What did we learn?

- Modeling rational choice extremely valuable
  - Important conclusions about market processes and behavior of economic agents
  - Identification of incentive misalignment
- Behavioral biases can and should be tested through experimentation
  - Humans do not react “rationally” to catastrophic events
  - Systematic (e.g., hyperbolic discounting, risk-aversion/seeking behavior) and non-systematic (e.g., Peltzman effect) biases impact utility functions
- Models must include aspects of limited rationality
  - E.g., near-rational agents

Part III: Measurements

Measuring and modeling adversaries' ecosystems

Question

- How can we model attacker behavior?
- Attackers usually not keen on being interviewed
- Modeling based on utility assumptions needs to be grounded in empirical evidence

- …however…
- Online attackers leave **lots** of data for us to analyze

Relevant papers (case studies)

- Online sale of prescription drugs
  1. Leontiadis, Moore and Christin. Measuring and analyzing search-redirection attacks in the illicit online prescription trade. *USENIX Security 2011*
  2. Leontiadis, Moore and Christin. A nearly four-year longitudinal study of search-engine poisoning. *ACM CCS 2014*
  3. Leontiadis, Moore and Christin. Pick Your Poison: Pricing and Inventories at Unlicensed Online Pharmacies. *ACM EC 2013*
- Online anonymous markets
  1. Christin. Traveling the Silk Road: A measurement study of a large anonymous online marketplace. *WWW'13*
  2. Soska and Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. *USENIX Security 2015*

Case study: Online sale of drugs

- One of the best known illicit online trades
  - Who hasn't received email spam for prescription drugs?

- Potentially most dangerous form of online crime
  - Wrong dosage can kill: cf. Ryan Haight
- Complex supply chain that can tell us a lot about the online criminal ecosystem

Supply chain: high-level overview

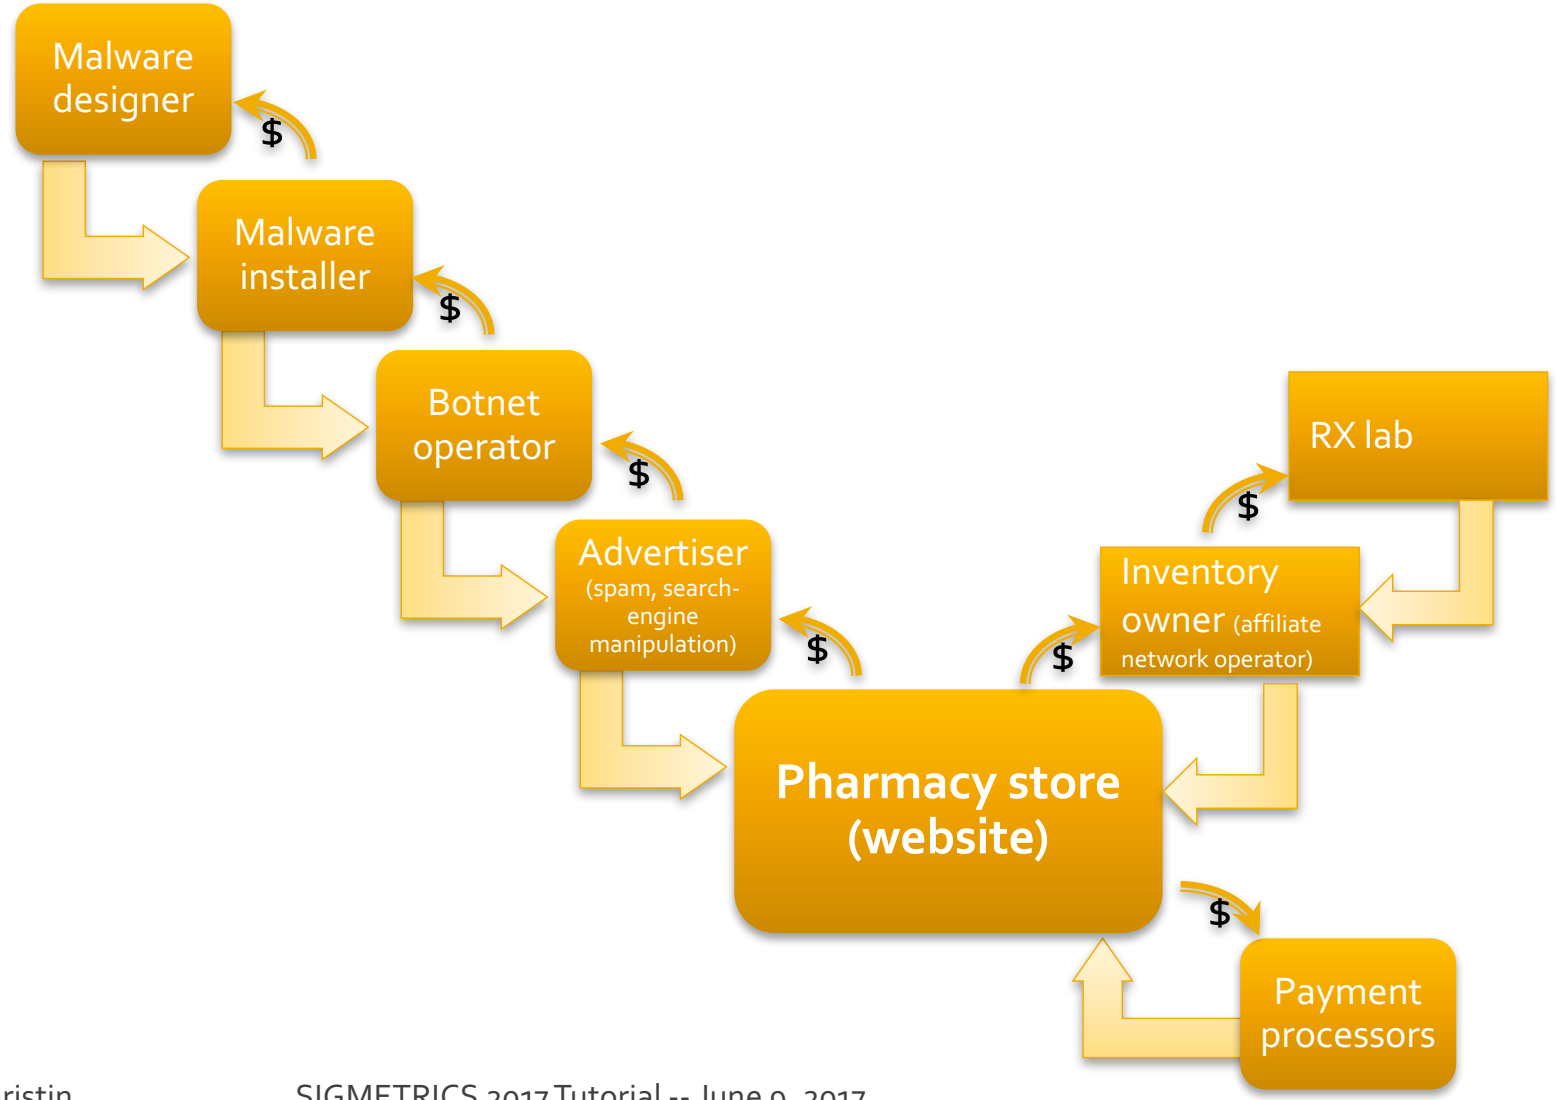

Supply chain: high-level overview

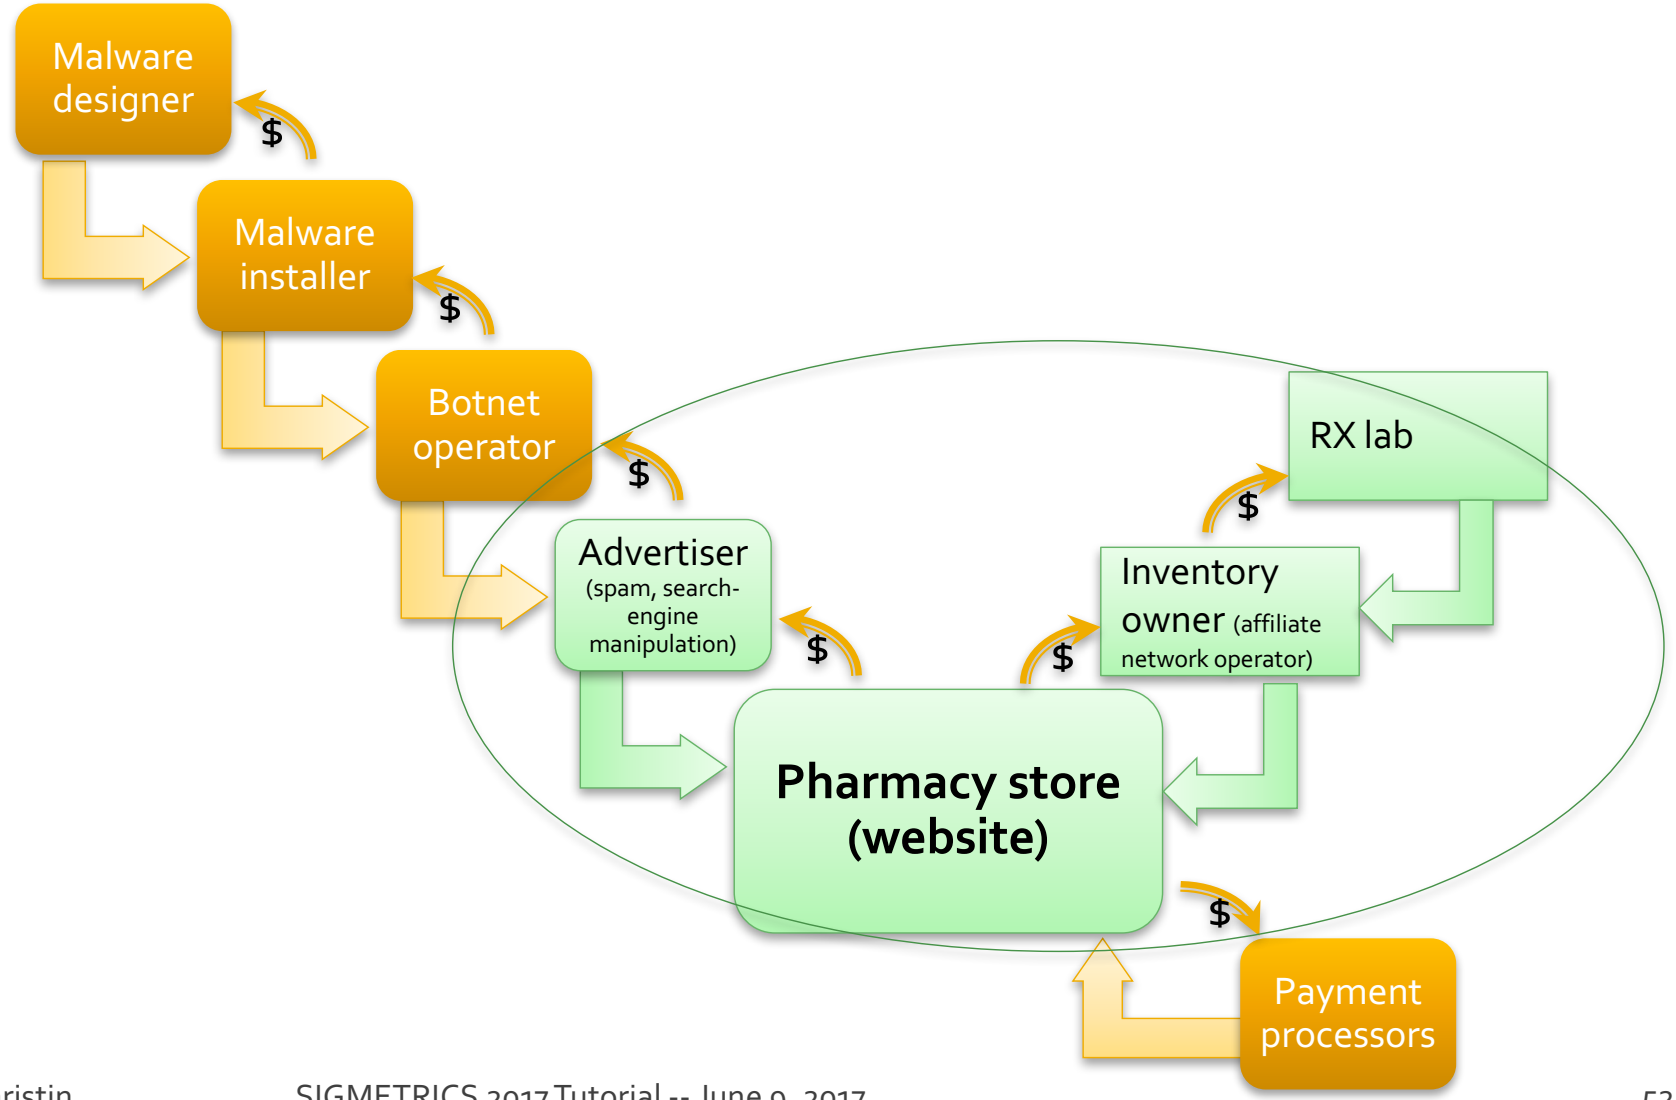

Advertising unlicensed drugs

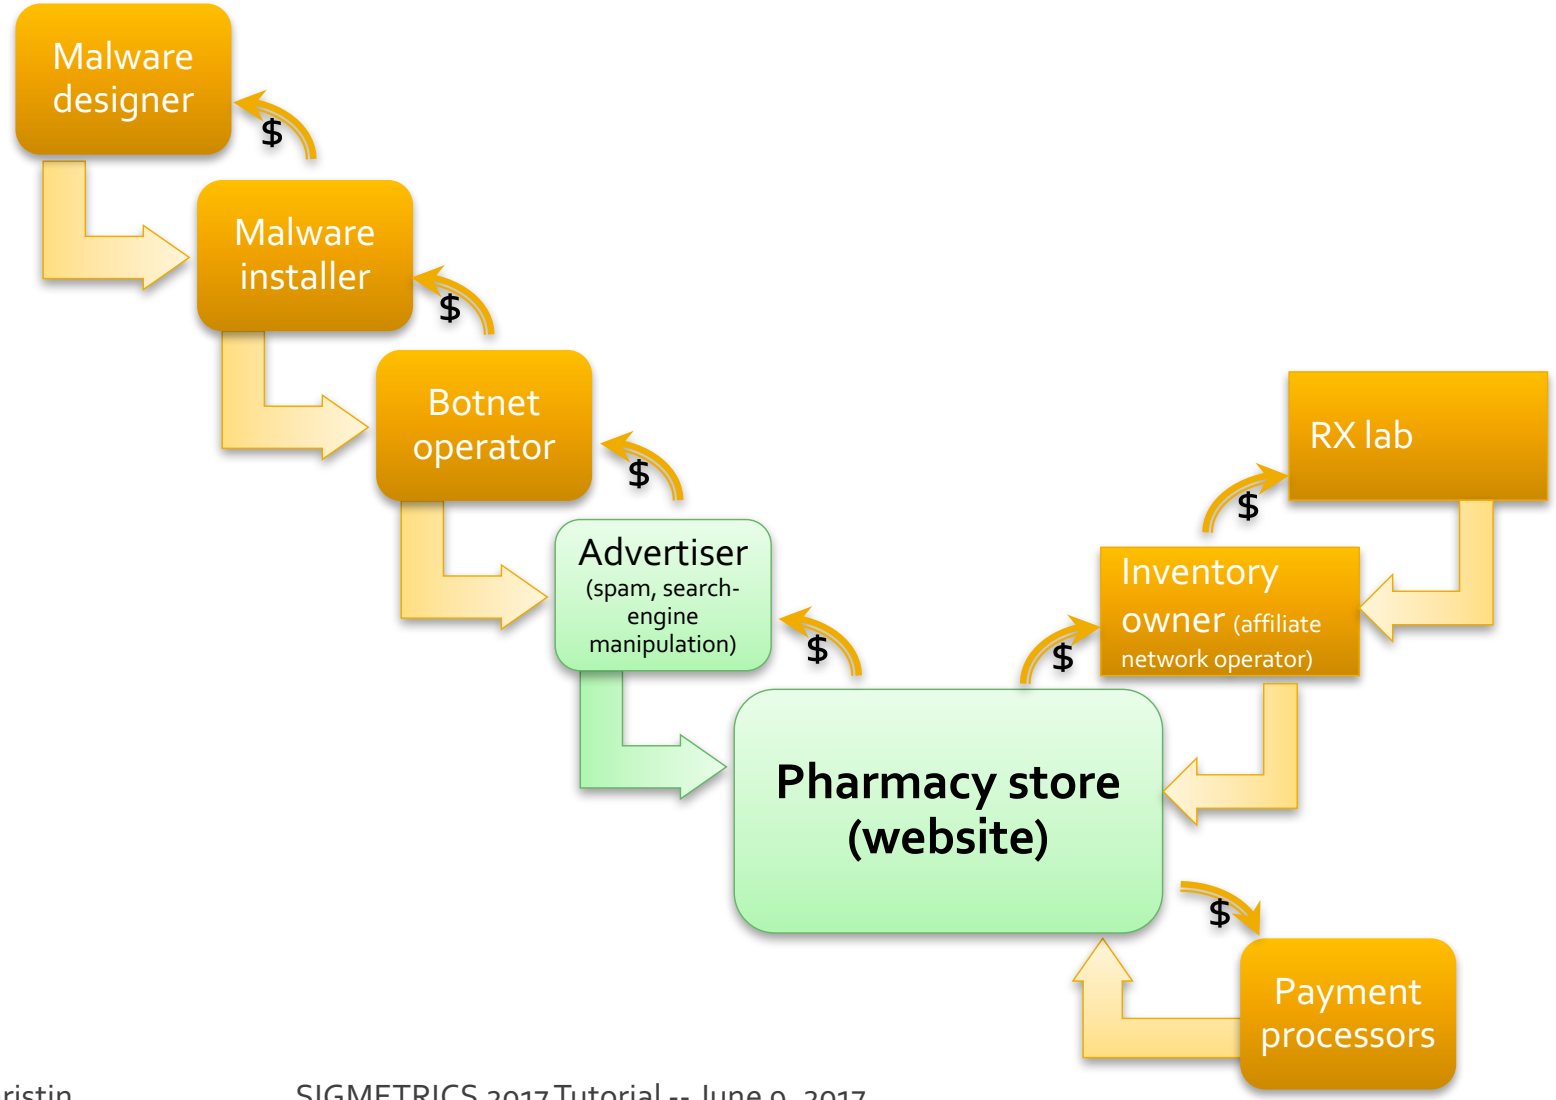

Evolution of advertising of illicit products

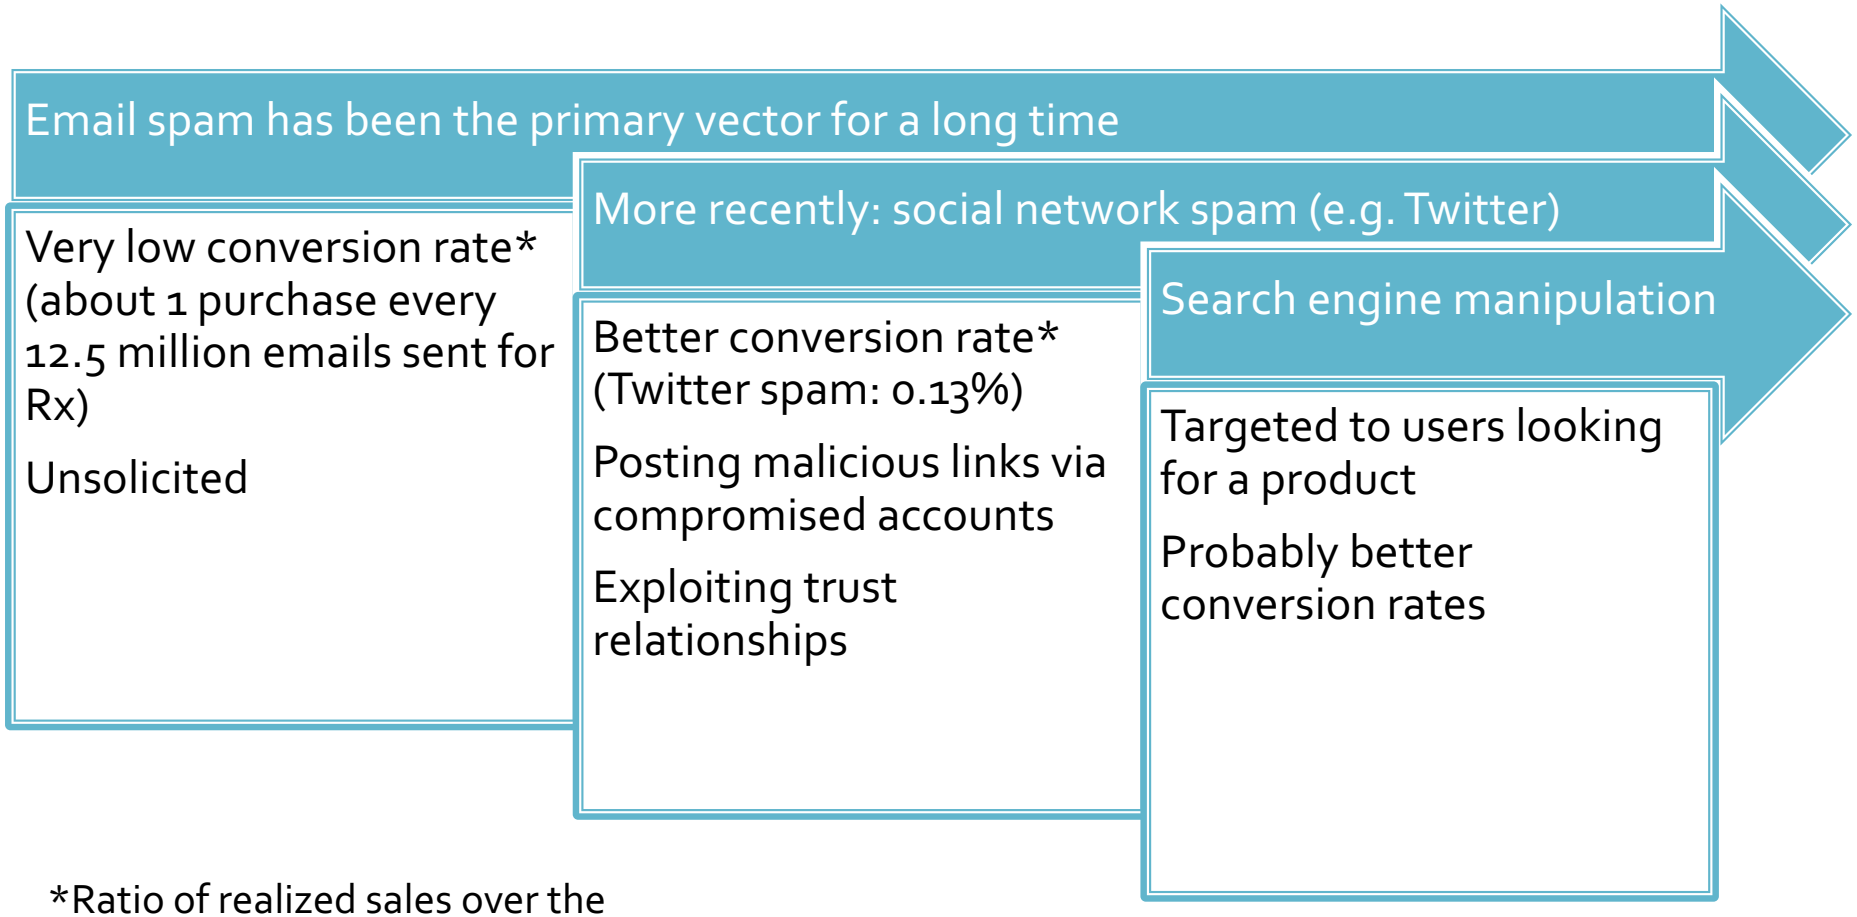

*Ratio of realized sales over the number of emails/clicks

Search-redirection

[LMC, USENIX Security 2011]

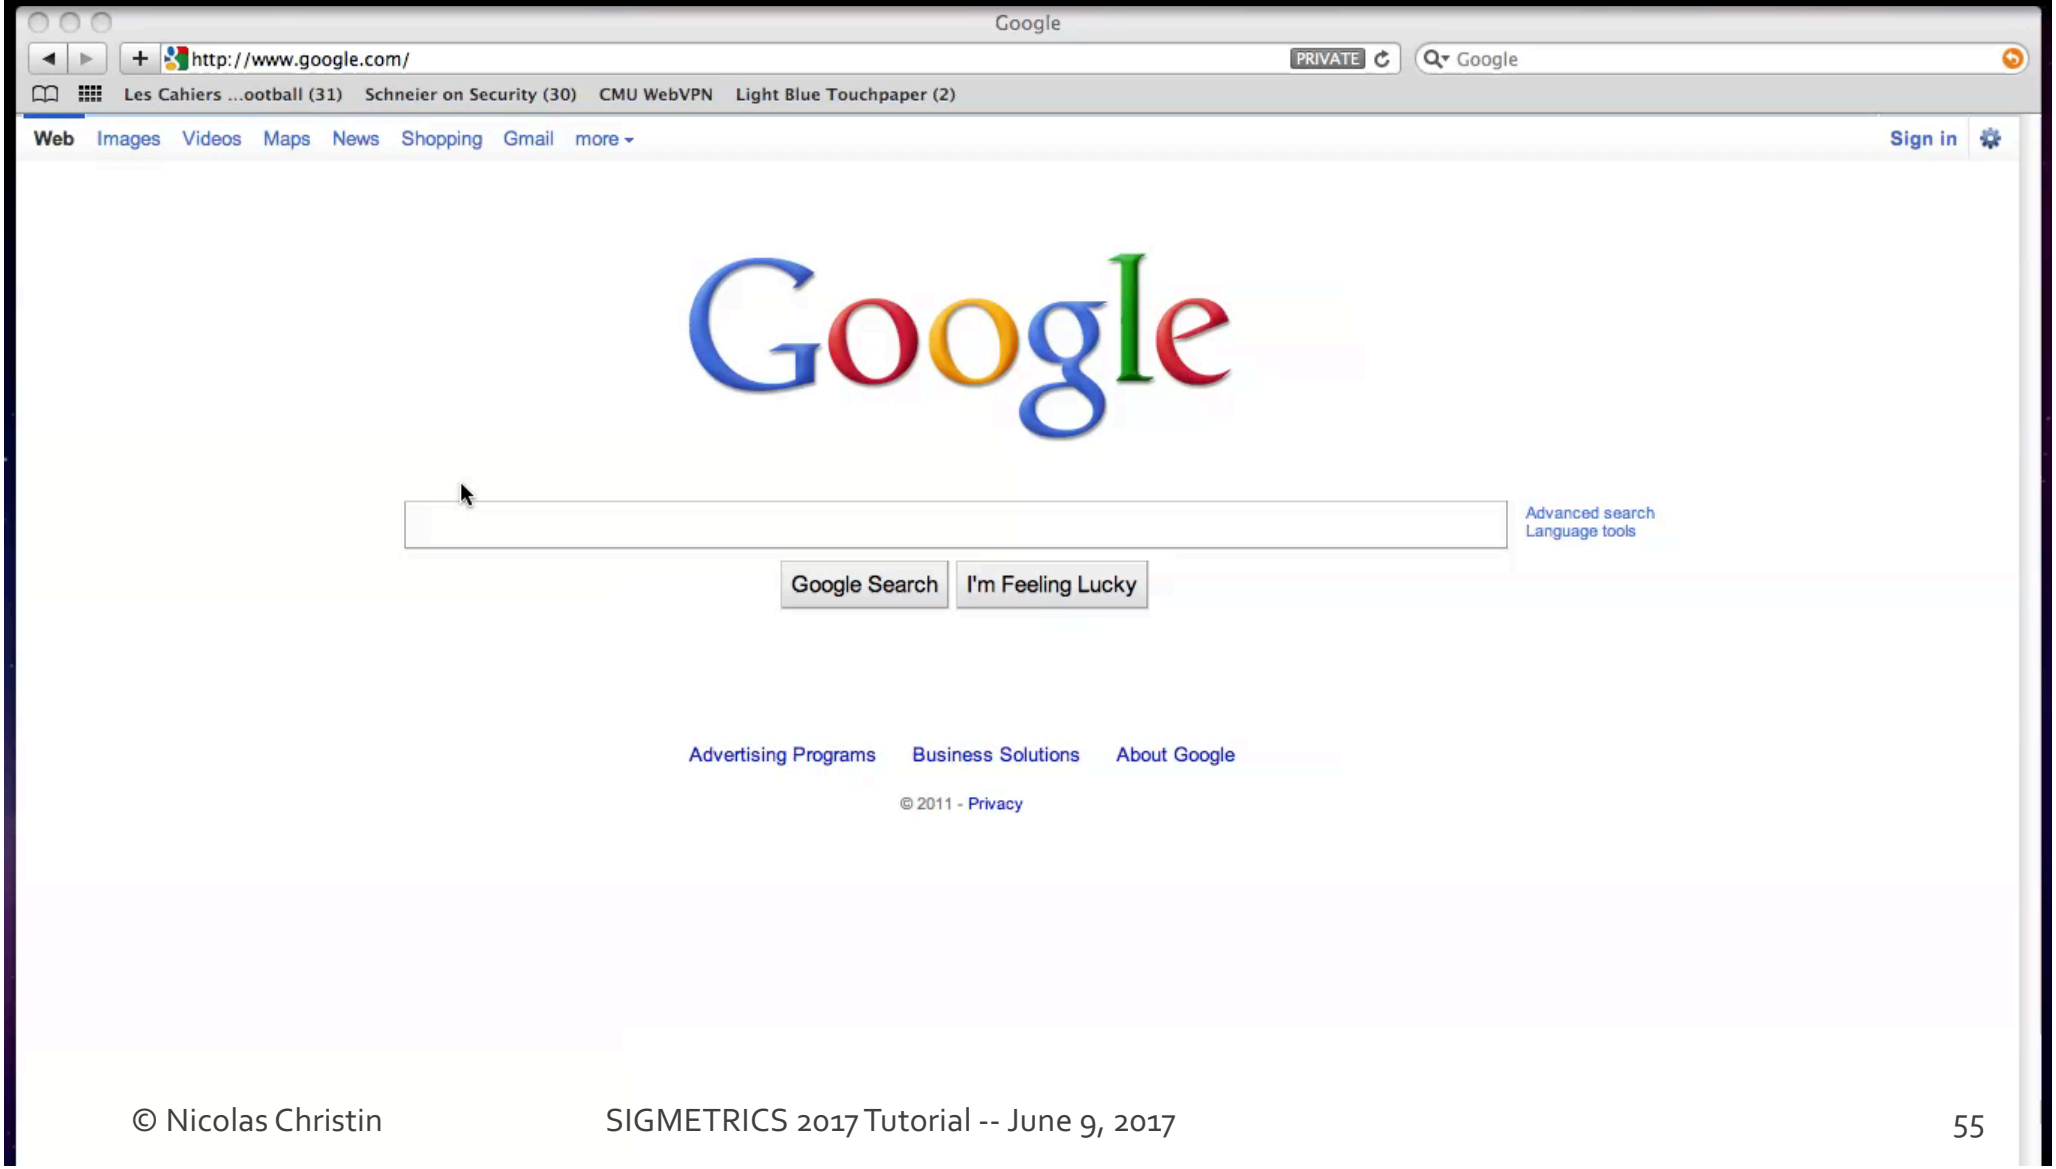

Attack modus operandi: Malware logic

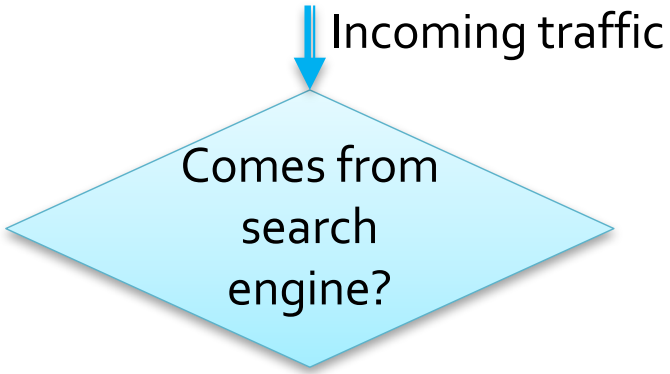

Attack modus operandi: Malware logic

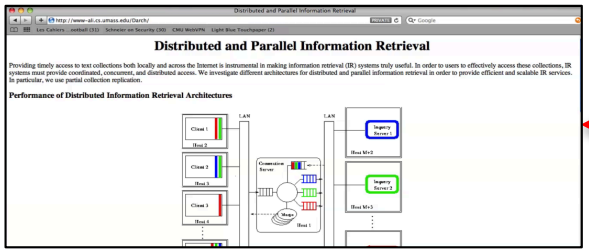

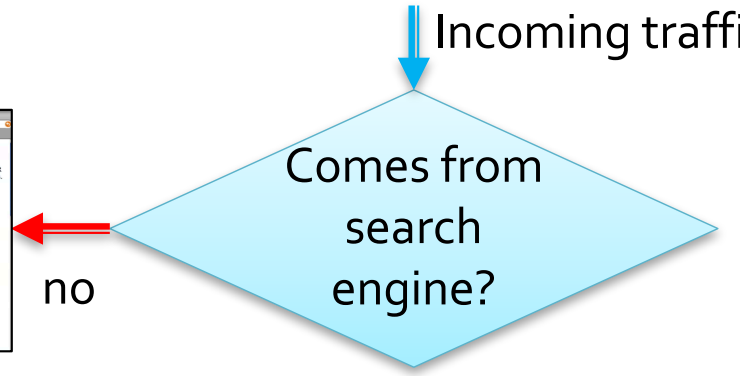

Attack modus operandi: Malware logic

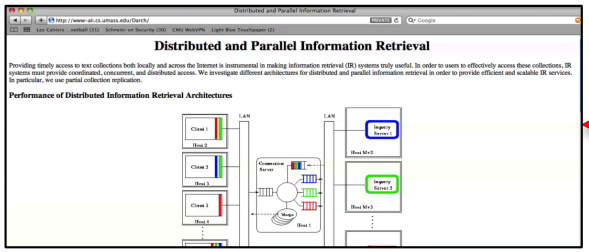

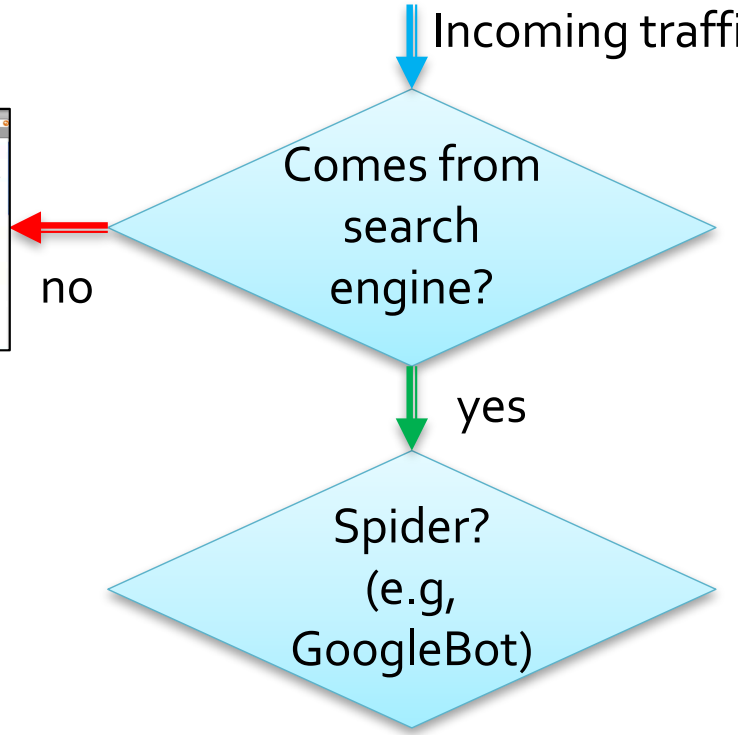

Attack modus operandi: Malware logic

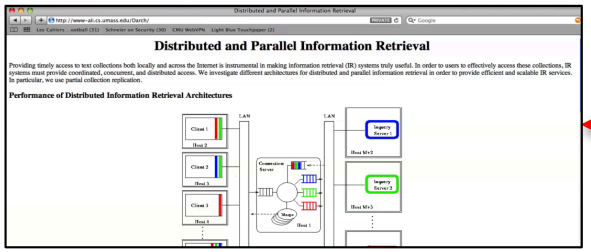

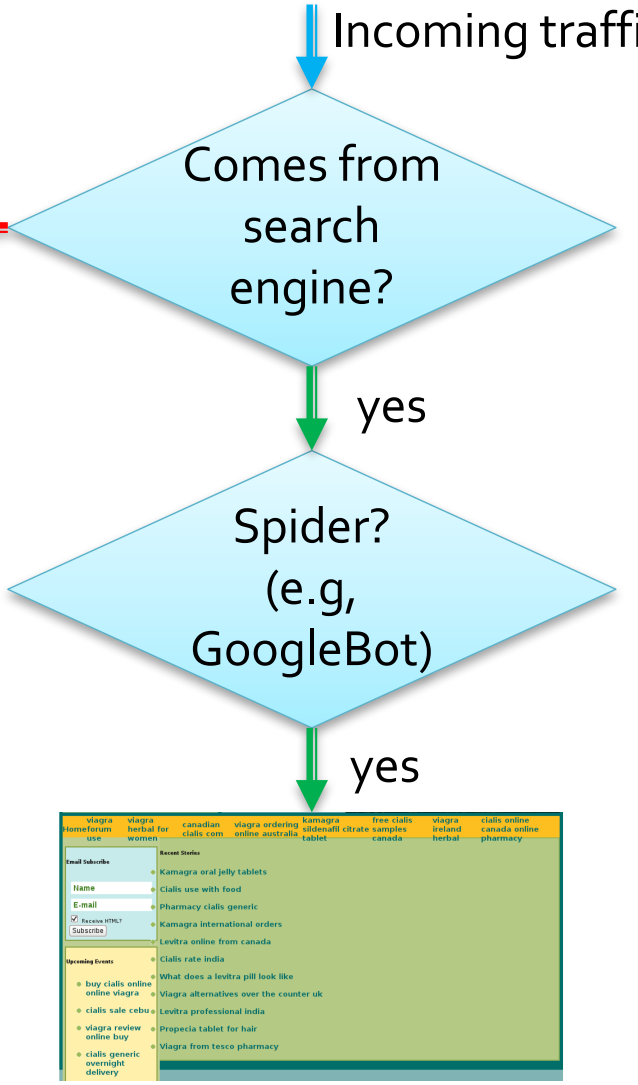

Attack modus operandi: Malware logic

Incoming traffic

Attack modus operandi: Malware logic

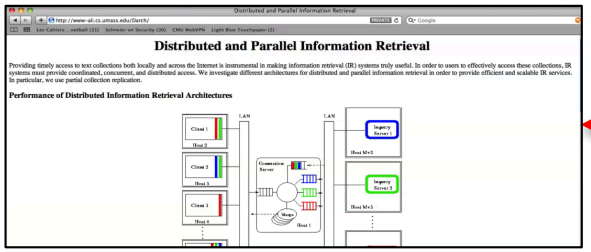

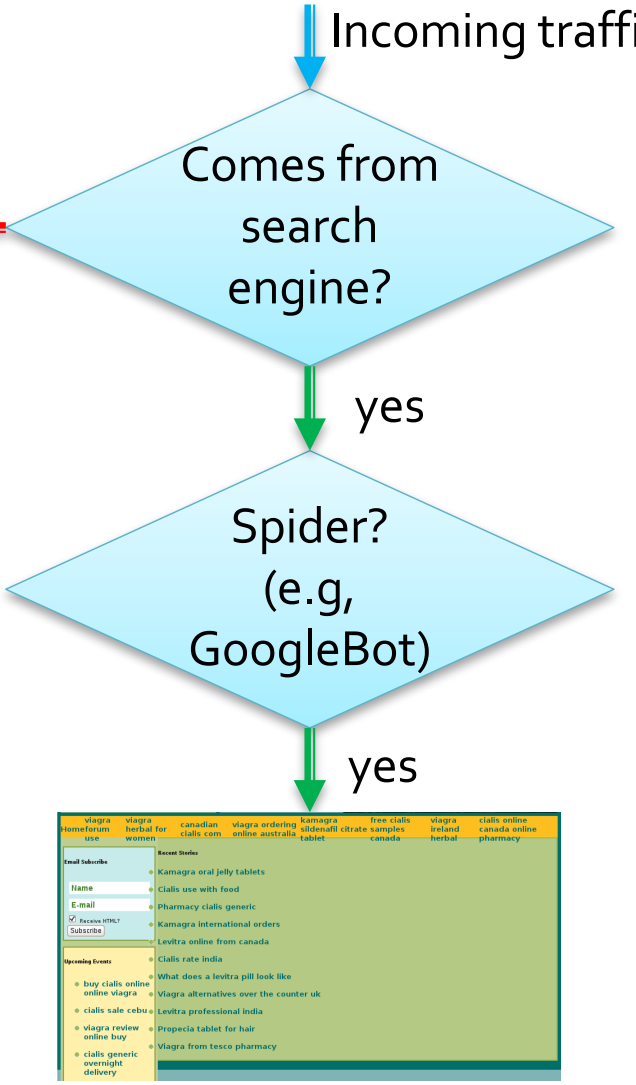

Attack modus operandi: Malware logic

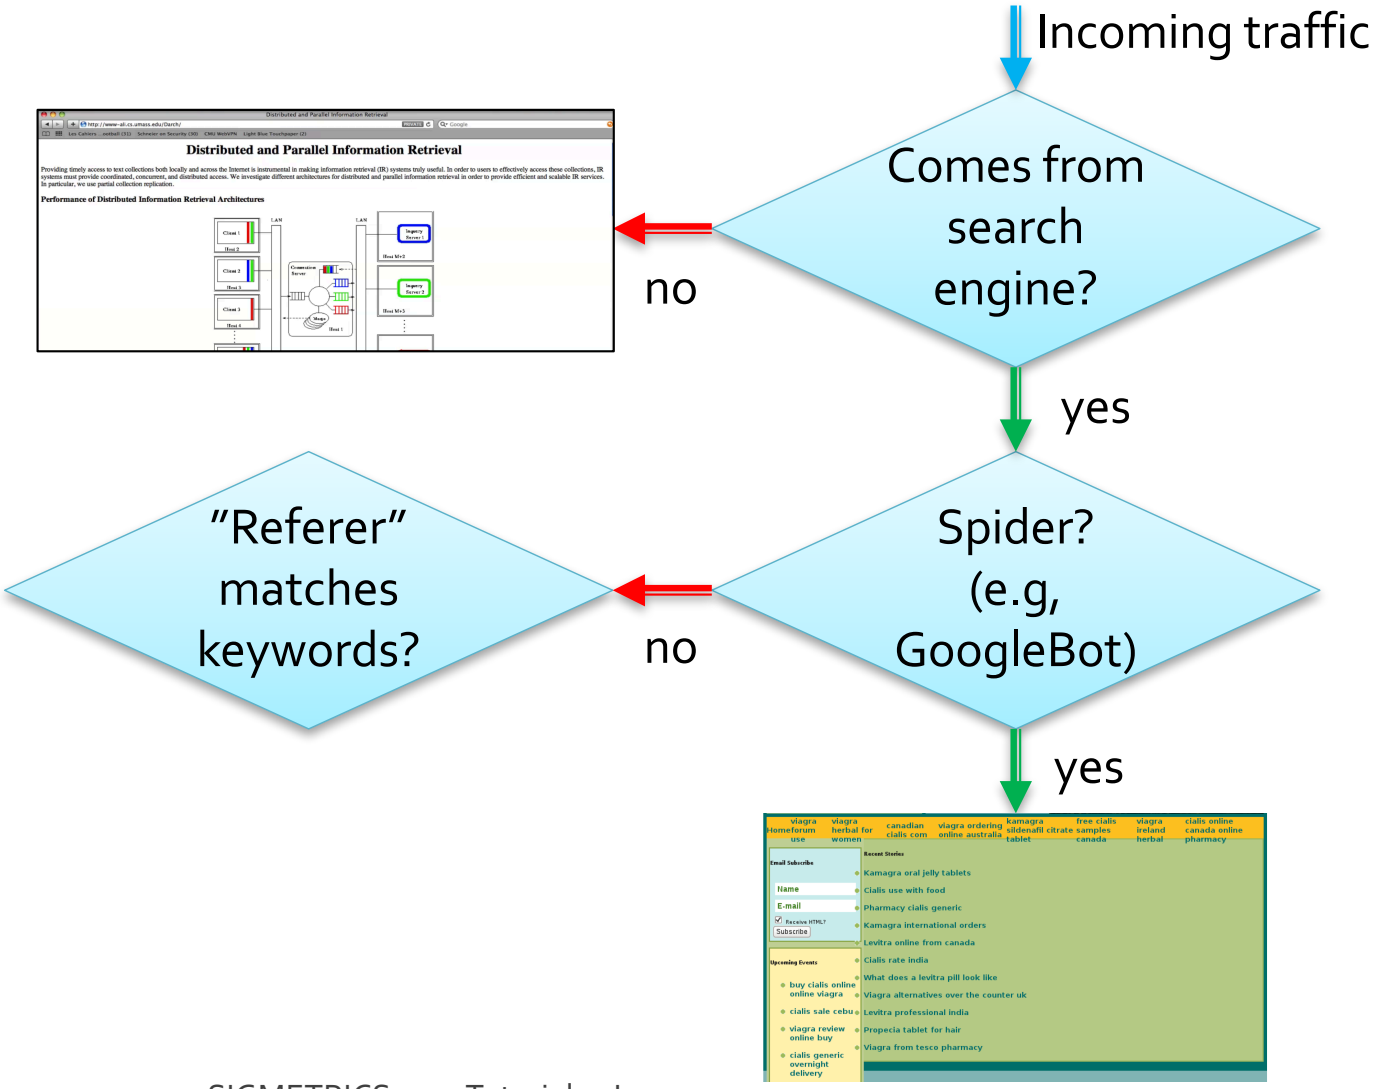

Attack modus operandi: Malware logic

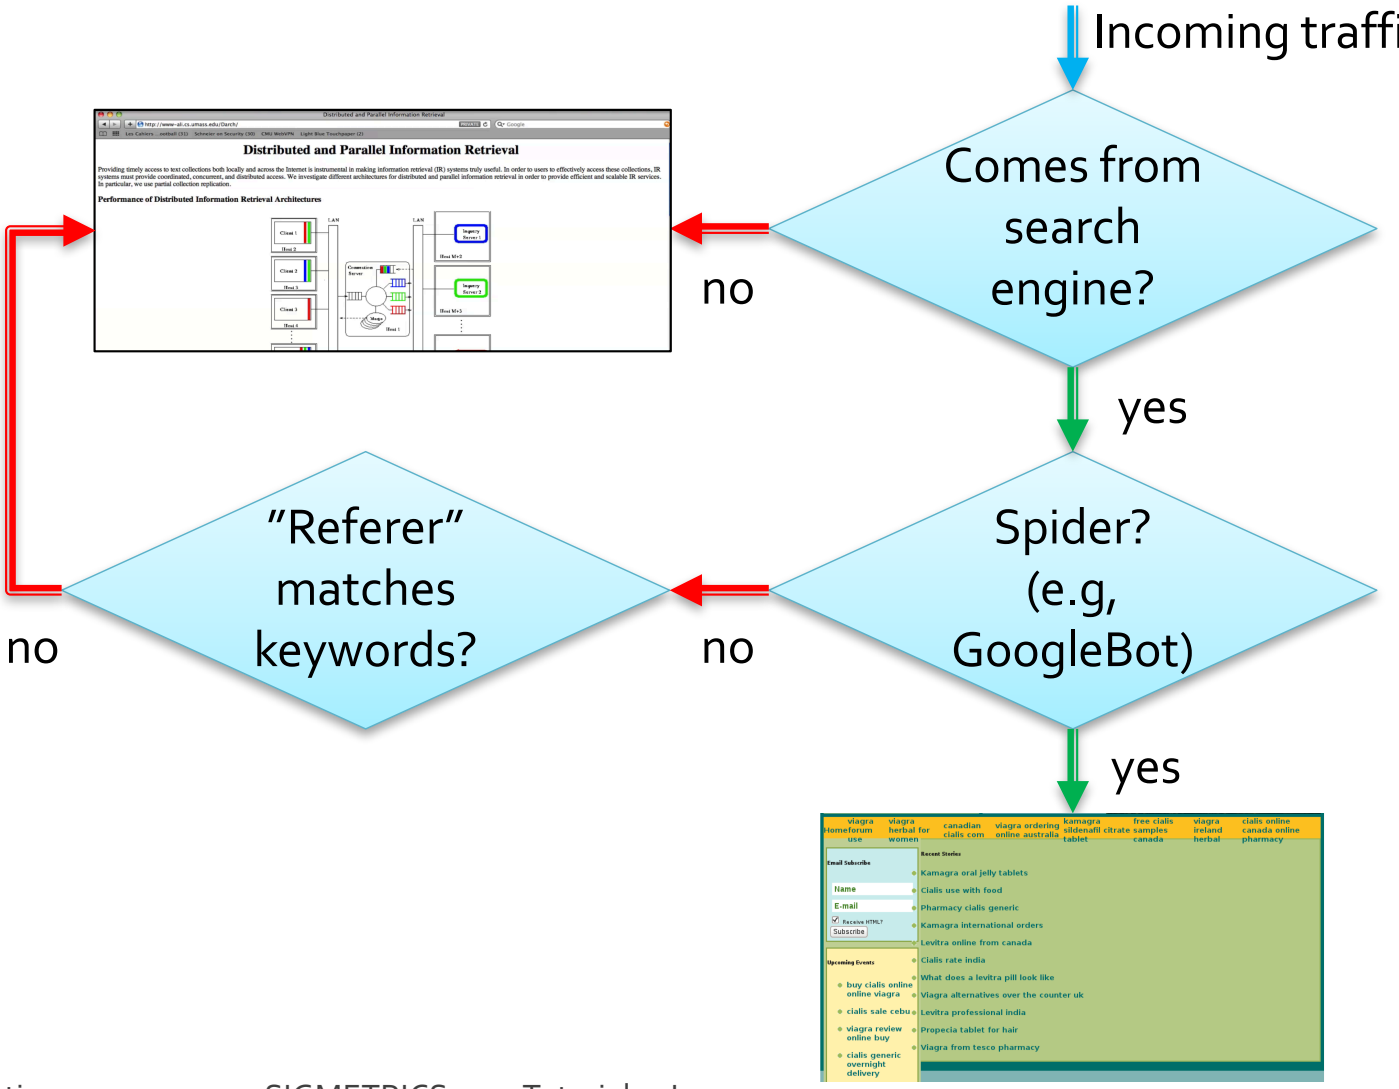

Attack modus operandi: Malware logic

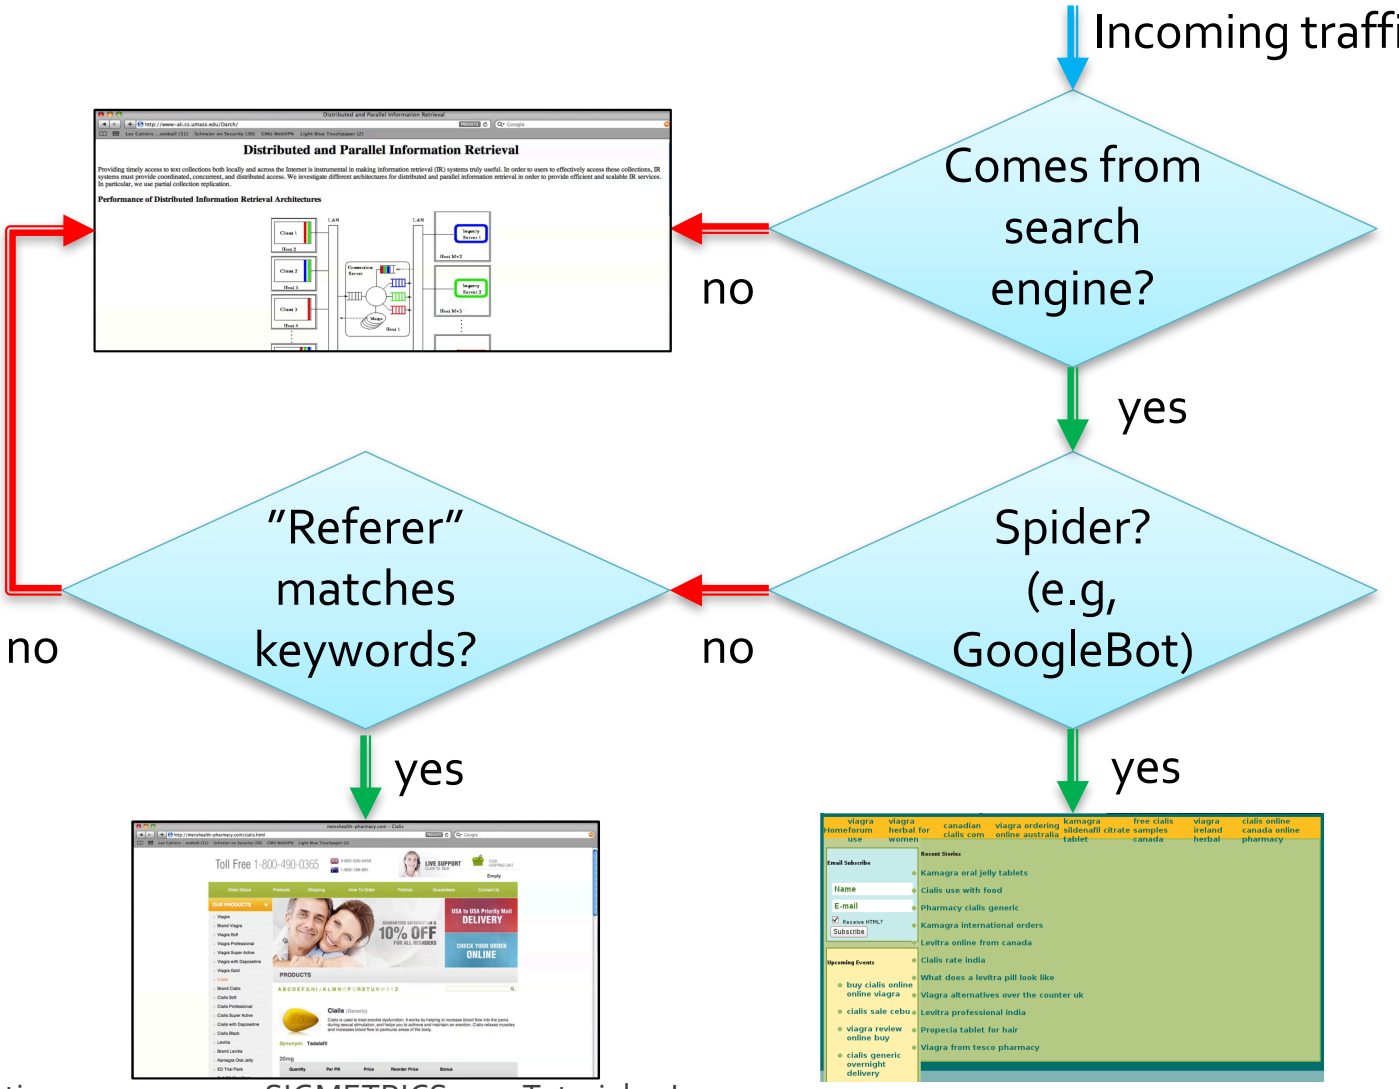

Attack modus operandi: Redirection chains

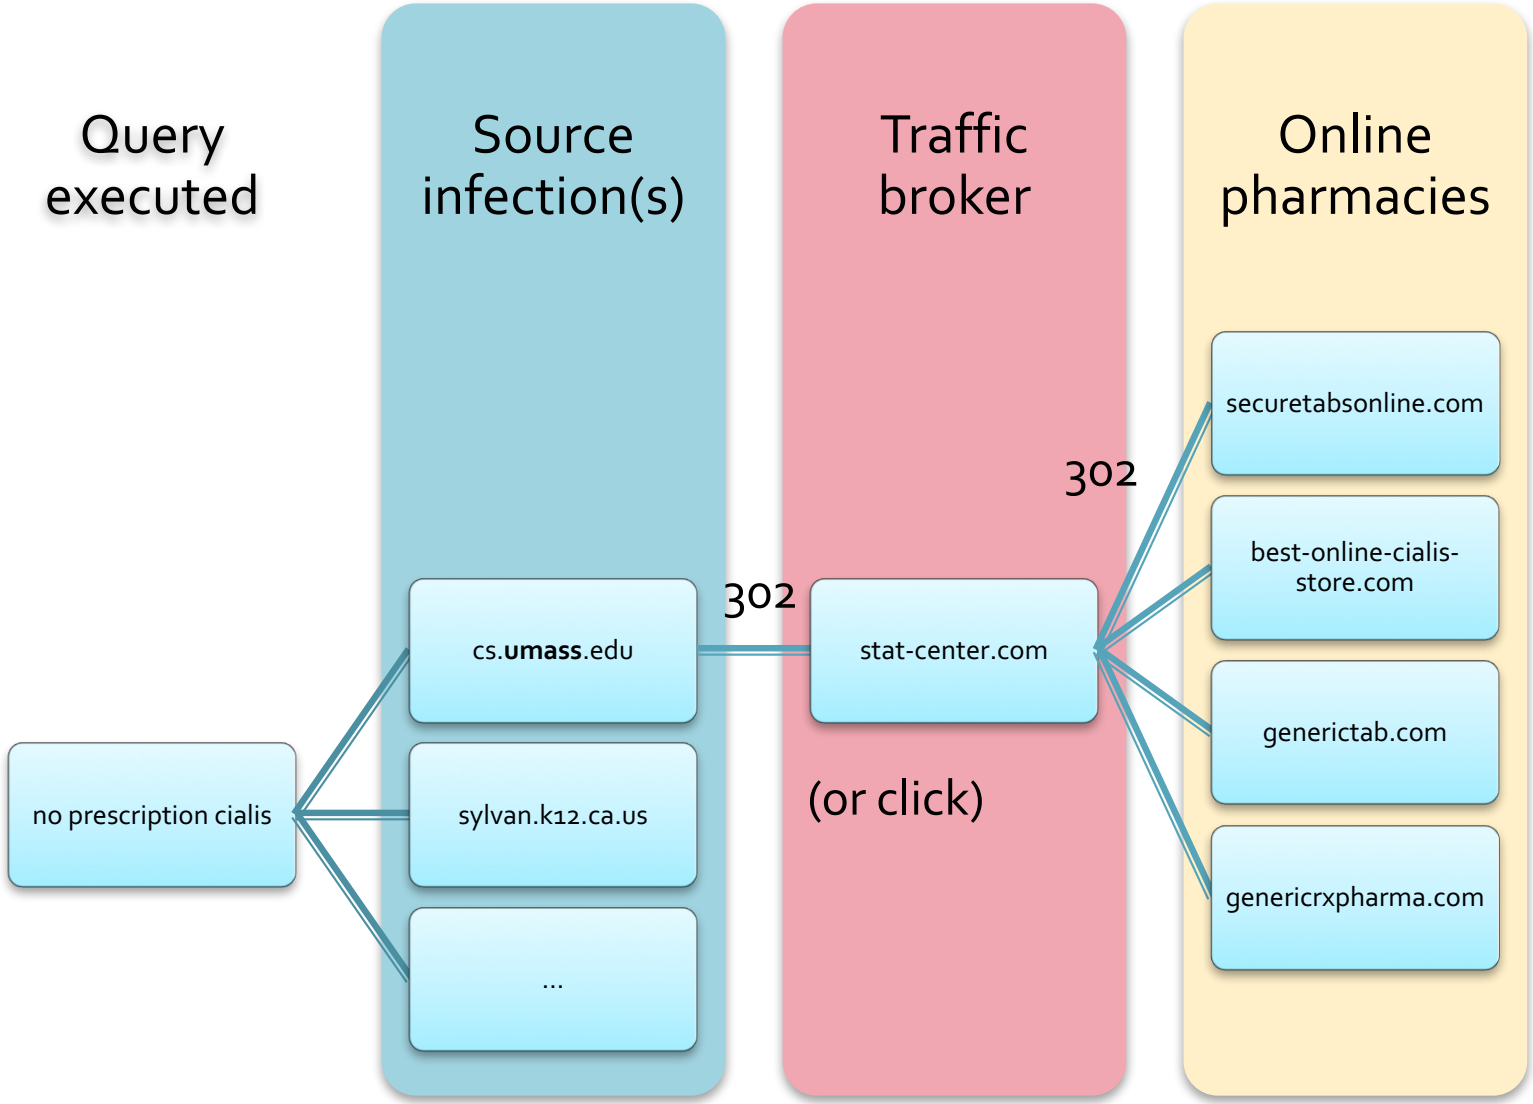

Questions

- How has this attack evolved?
  - Volume and impact: Does this even matter?
  - Techniques
- Why has the attack evolved?
  - Effectiveness of the defenses?
- Can this be thwarted?
  - Policy interventions vs. technical defenses

Data collection process

- Daily collection from 4/12/2010 through 9/16/2013

- Limit due to Google Search API
- Store all results for later processing
- Will also examine position information

- A search result defines the website that a user will be redirected to when clicking on the link
- If the browser is redirected instead to a different website (domain), the result is infected.

- Follow each result identified as infected from previous step
- Follow all redirections that might occur
- Record all the redirection information

Datasets collected

[LMC USENIX Security 2011, CCS 2014]

- Dataset 1
  - Aggregate results only
  - Rank of the results unknown
  - Mapping query-results unknown
- Dataset 2
  - Same as Dataset 1, but ranking information known
  - Mapping query-result doesn't include rank
- Dataset 3
  - All information is captured
  - … but new Google API (slightly) limits results returned

Dataset	1	2	3
Period	4/12/10-11/15/10	11/15/10-10/8/11	10/8/11-9/16/13
Search results/query	64	64	16/32
Total results	260,824	3,609,675	1,530,099
Unique URLs	150,955	189,023	122,382
Unique domains	25,182	36,557	30,881

This is 3.5 years worth of data!

Some of the 218 queries used

vicodin no prescription
cheap valium non prescription
buy ativan online injecting pills
buy xanax valium online florida
order vicodin si levitra online
buy xanax valium online florida
color of adipex pills safest place to buy online
vicodin without prescription
generic cialis free sample
cheap tadalafil
20 mg ambien overdose
prozac side effects
ambien buy online
alprazolam online without prescription buy cheap

Type	Count	Percent
Malicious (Black)	26	22%
Benign (White)	75	34%
Ambiguous (Gray)	117	54%
Total	218	100%

Long-term evolution

[LMC CCS 2014]

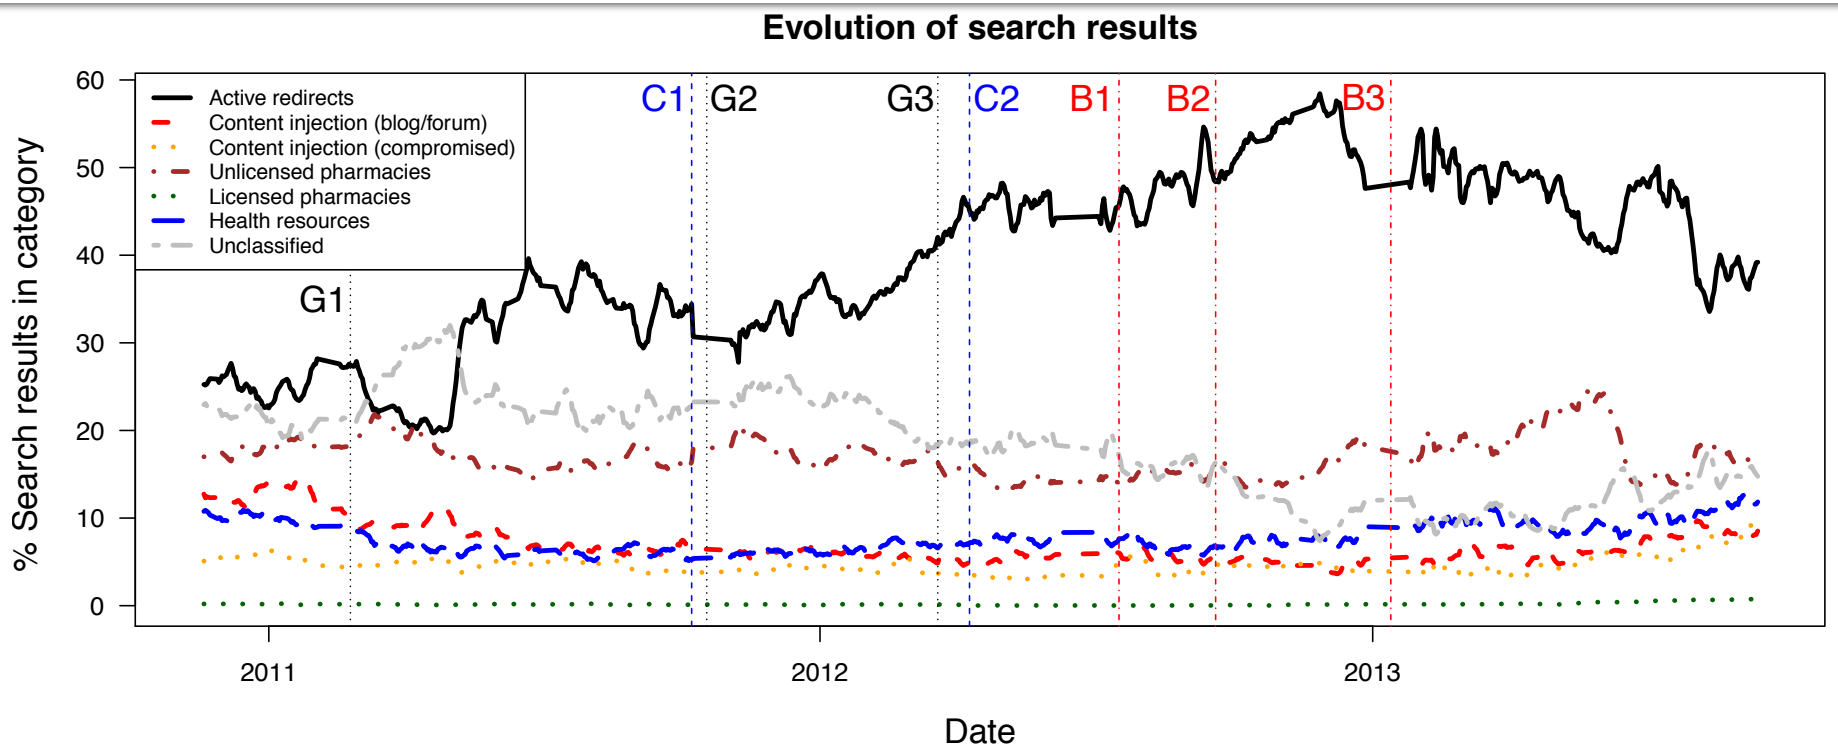

G1: Google changes search ranking algorithm
G2: Google starts removing query info from "Referer" field
G3: Google is done deploying Referer modifictations
B1, B2, B3 : Firefox, Safari, Chrome switch to HTTPS-only search
(C1,C2: major changes to our collection infrastructure)

Uncovering relationships in search results

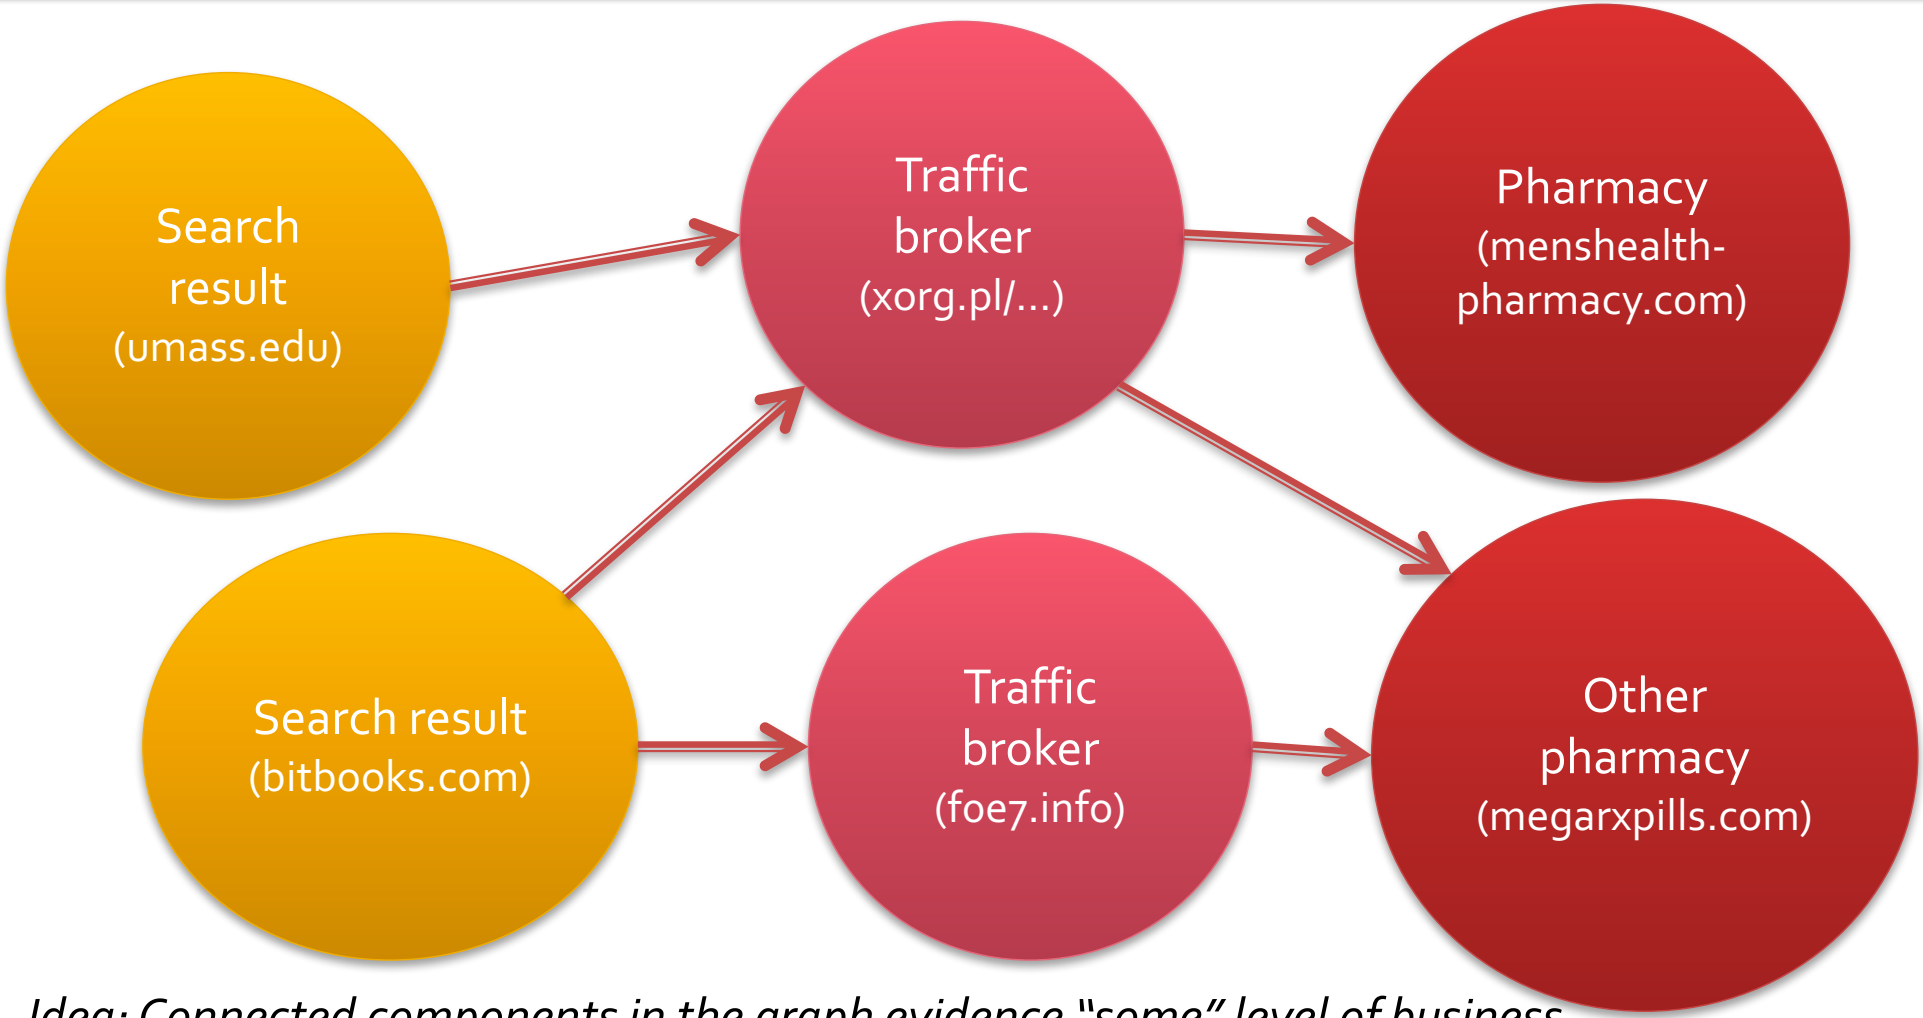

Idea: Connected components in the graph evidence "some" level of business relationships between the nodes they connect

Connected components

- 34 connected components
- One connected component contains
  - 96% of all infected domains
  - 90% of all redirection domains
  - 92% of all pharmacies

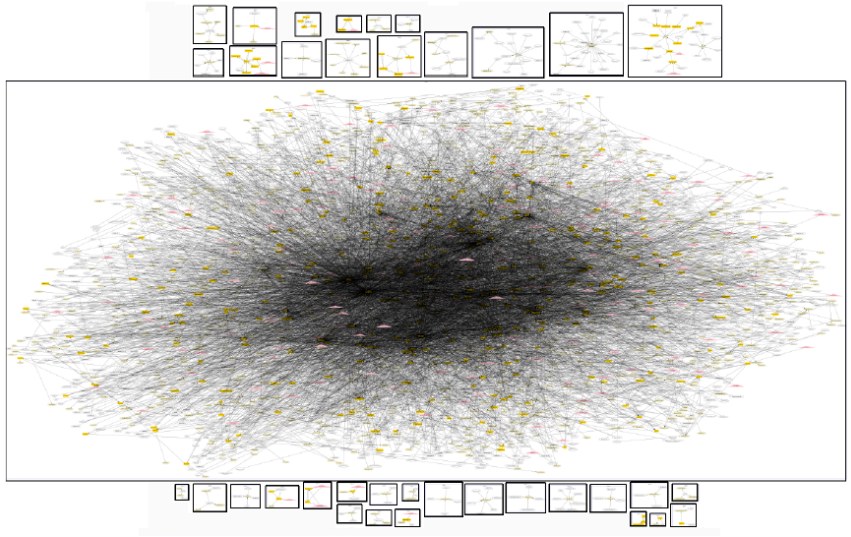

- Is one person responsible for all of this?!

Connected components

- 34 connected components
- One connected component contains
  - 96% of all infected domains
  - 90% of all redirection domains
  - 92% of all pharmacies

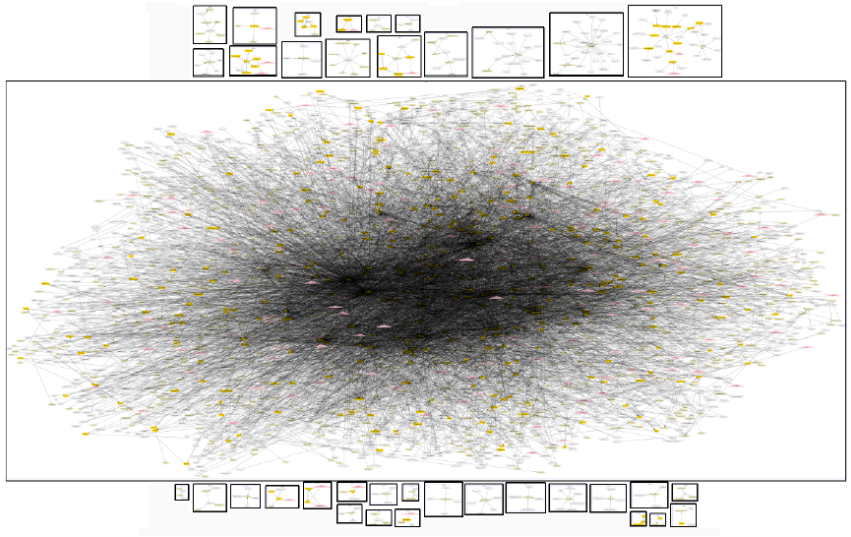

- Is one person responsible for all of this?!
  - **NO!**
  - Some advertisers work for several different affiliate networks
  - Certain domains are (ab)used by multiple advertisers

Identifying the main players

- Run (spinglass) clustering algorithm in big connected component
- Each cluster represented by different color
- Evidence of separate organized groups/campaigns more loosely connected to each other
  - About 10-12 large groups

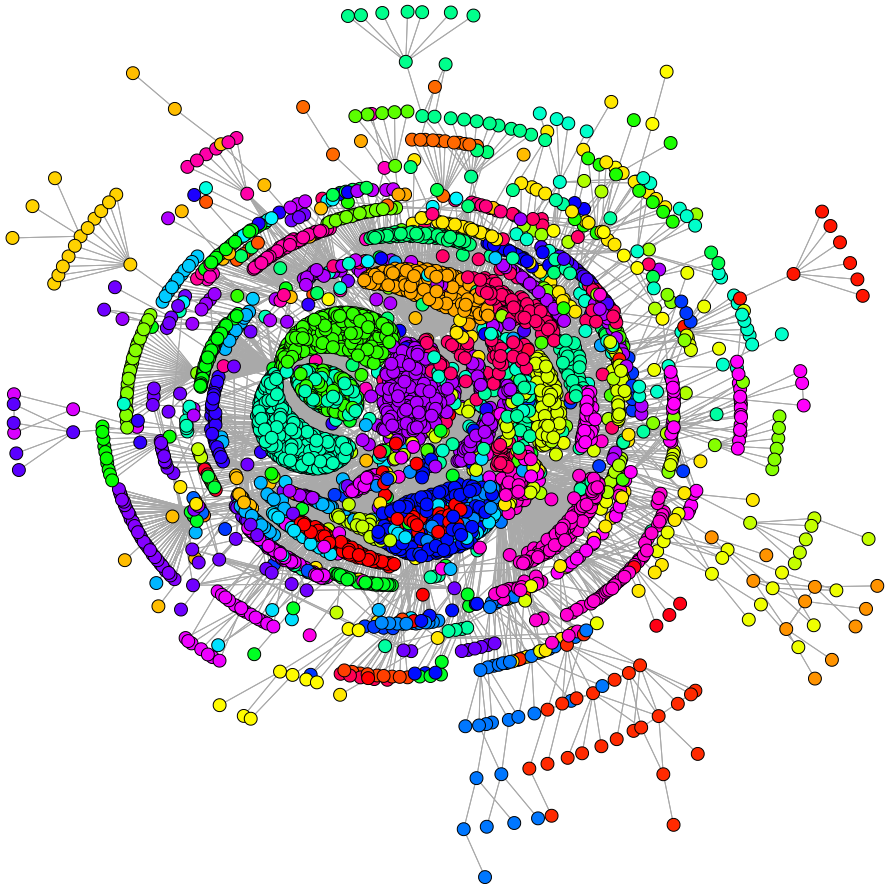

Identifying the main players

- Clusters may denote
  - A set of hosts compromised at the same time
  - A set of hosts all compromised by the same people
- Interesting network properties
  - All major redirectors hosted on 11 ASes

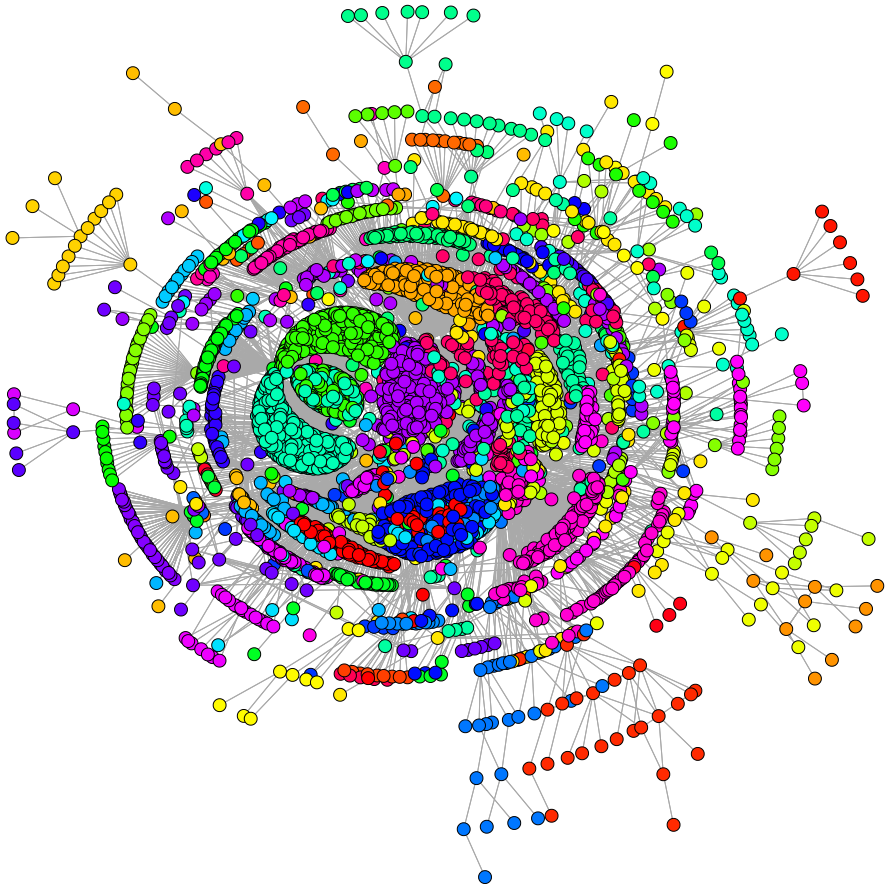

Possible interventions

Valuable infected sites (e.g., umass.edu)

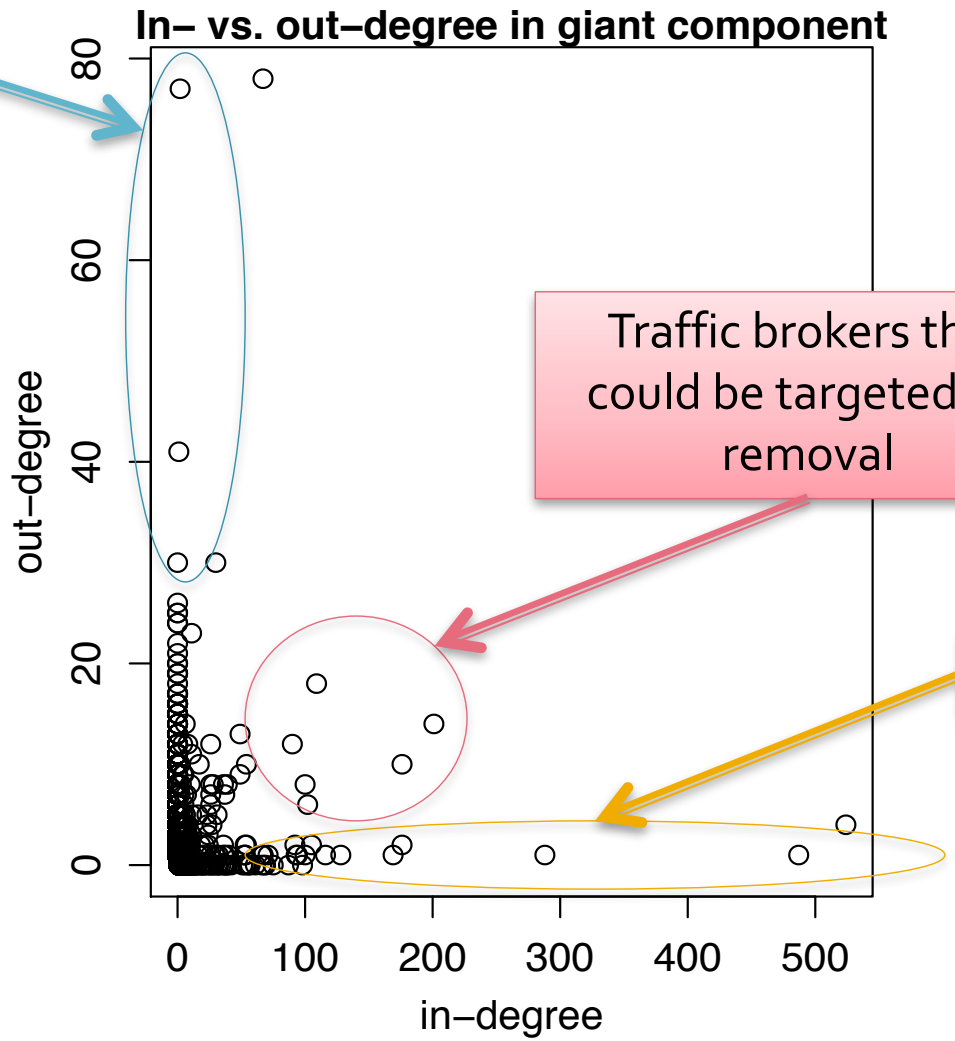

Traffic brokers that could be targeted for removal

Pharmacies

(*) Results here are for 2010-11

Illicit advertising infrastructure

- Traffic brokers are disproportionately hosted on **very few** networks

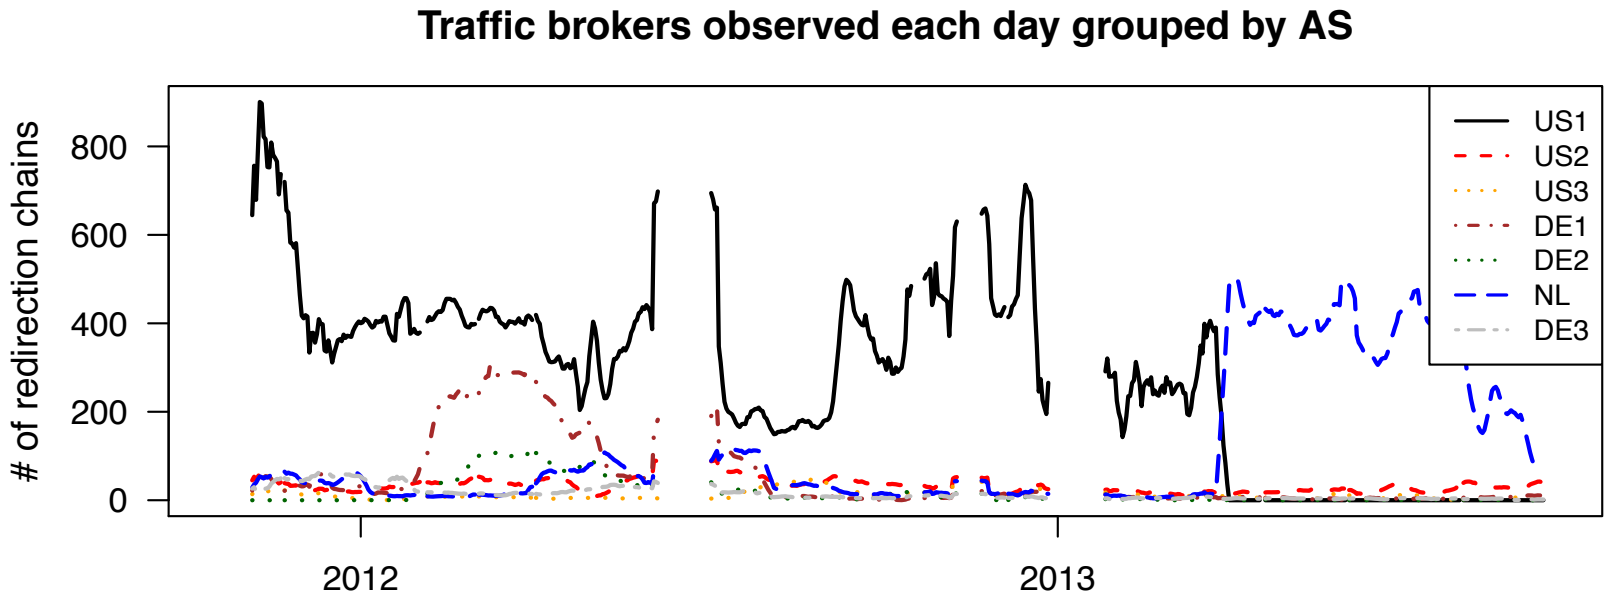

Procuring unlicensed drugs

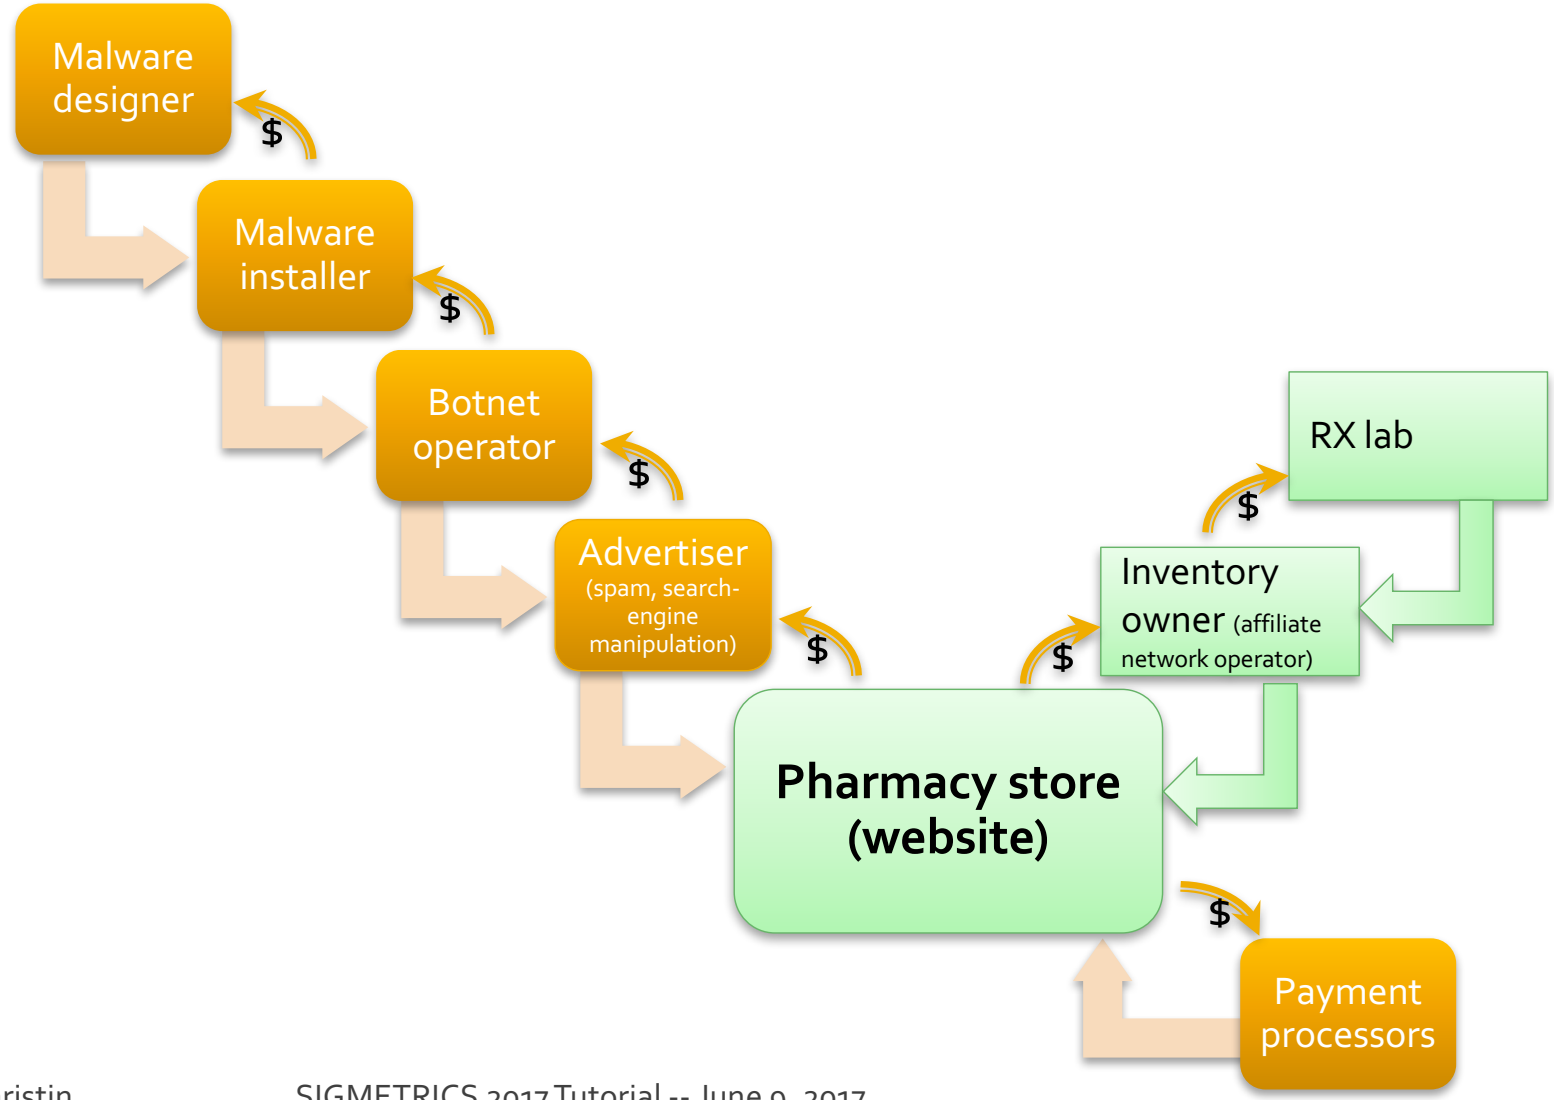

Inventory analysis [LMC, 2013]

Data collected

- Scraped for prices and inventories:
  - 265 unlicensed pharmacies (doing search-redirection attacks) collected between April 3, 2012 and October 16, 2012
  - 265 "blacklisted" pharmacies
    - Randomly sampled out of a corpus of 9000+ NABP "not recommended" pharmacies
    - No overlap with the unlicensed pharmacy corpus

Scraping

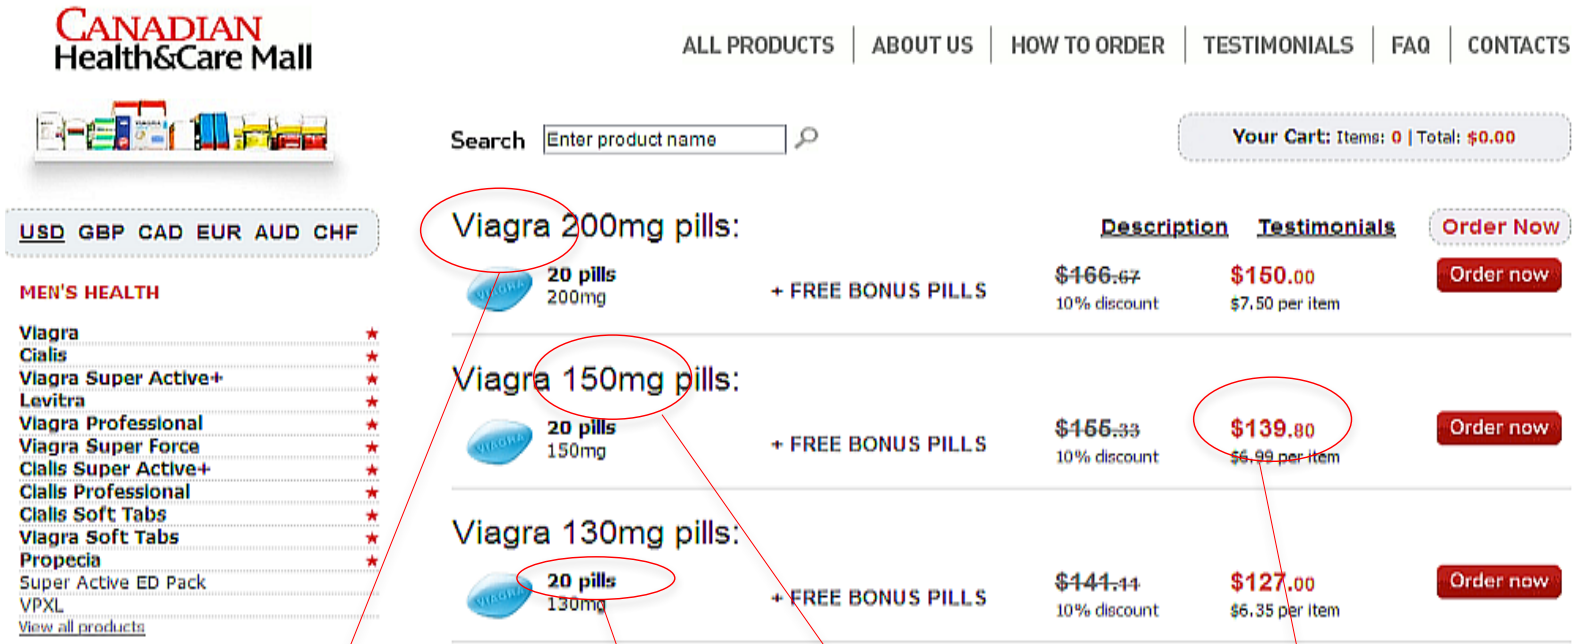

Drug name/product (subsequently mapped to active ingredient/condition)

Quantity

Dosage

Price

Total = 1,451,587 distinct (drug, active ingredient, dosage, unit) tuples collected
1,661 different drug names

Identifying common suppliers: Inventory overlap

- How much overlap is there between distinct unlicensed pharmacies' inventories?
- Jaccard distance:

$$J_\delta(A,B) = 1 - \frac{|A \cap B|}{|A \cup B|}$$

Inventory of pharmacy $\mathscr{A}$

Inventory of pharmacy $\mathscr{B}$

  - Identical inventories $\Leftrightarrow J_\delta(A,B) = 0$
  - No overlap at all $\Leftrightarrow J_\delta(A,B) = 1$

Clustering inventories

- Inventory A and Inventory B belong to the same cluster iff $J_\delta(A,B) < t$
  - t is an arbitrary threshold, $0<t<1$
- Distance between two groups of inventories X, Y:
  - Minimum linkage: $J_\delta(X,Y) = \min\{J_\delta(x,y) : x \in X, y \in Y\}$
  - Maximum linkage: $J_\delta(X,Y) = \max\{J_\delta(x,y) : x \in X, y \in Y\}$
  - Average linkage: $J_\delta(X,Y) = \frac{1}{|X| \cdot |Y|} \sum_{x \in X} \sum_{y \in Y} J_\delta(x,y)$

Clustering inventories: Average linkage, *t=0.31*

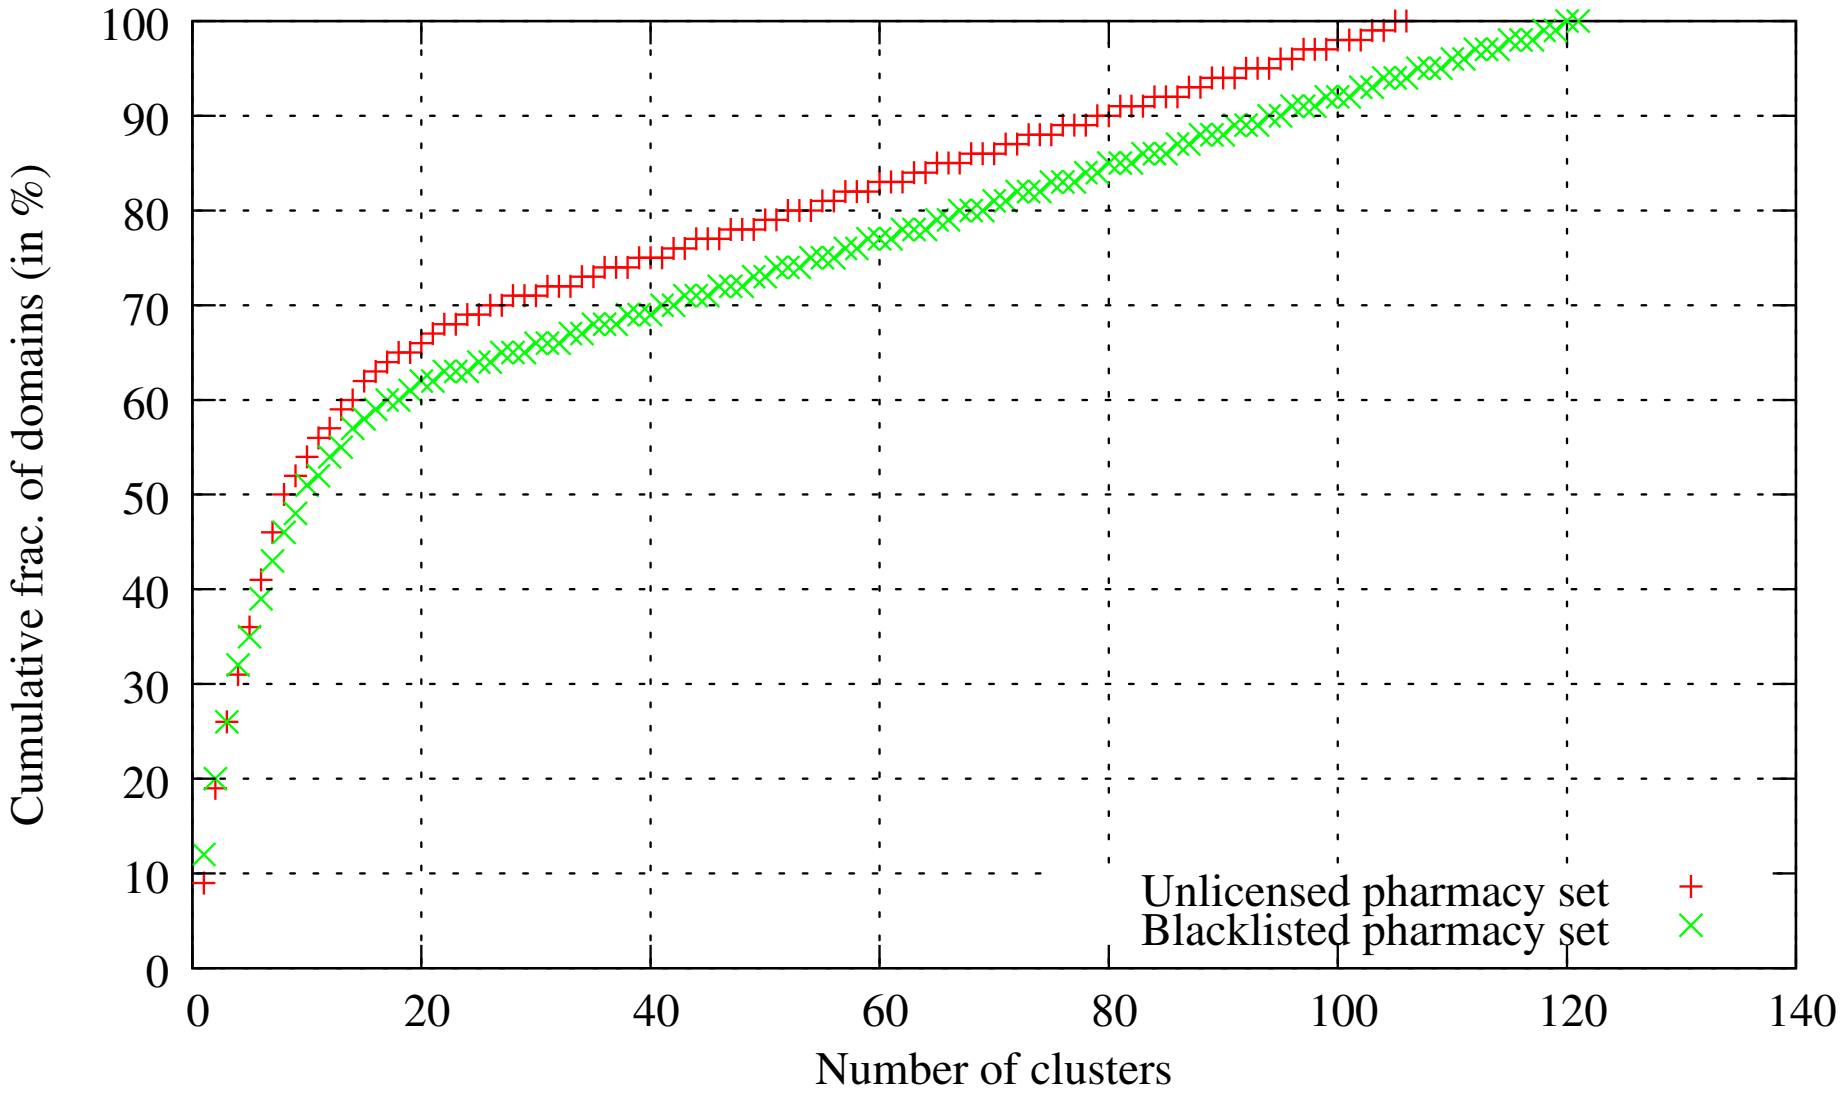

Clustering inventories: Average linkage, *t=0.31*

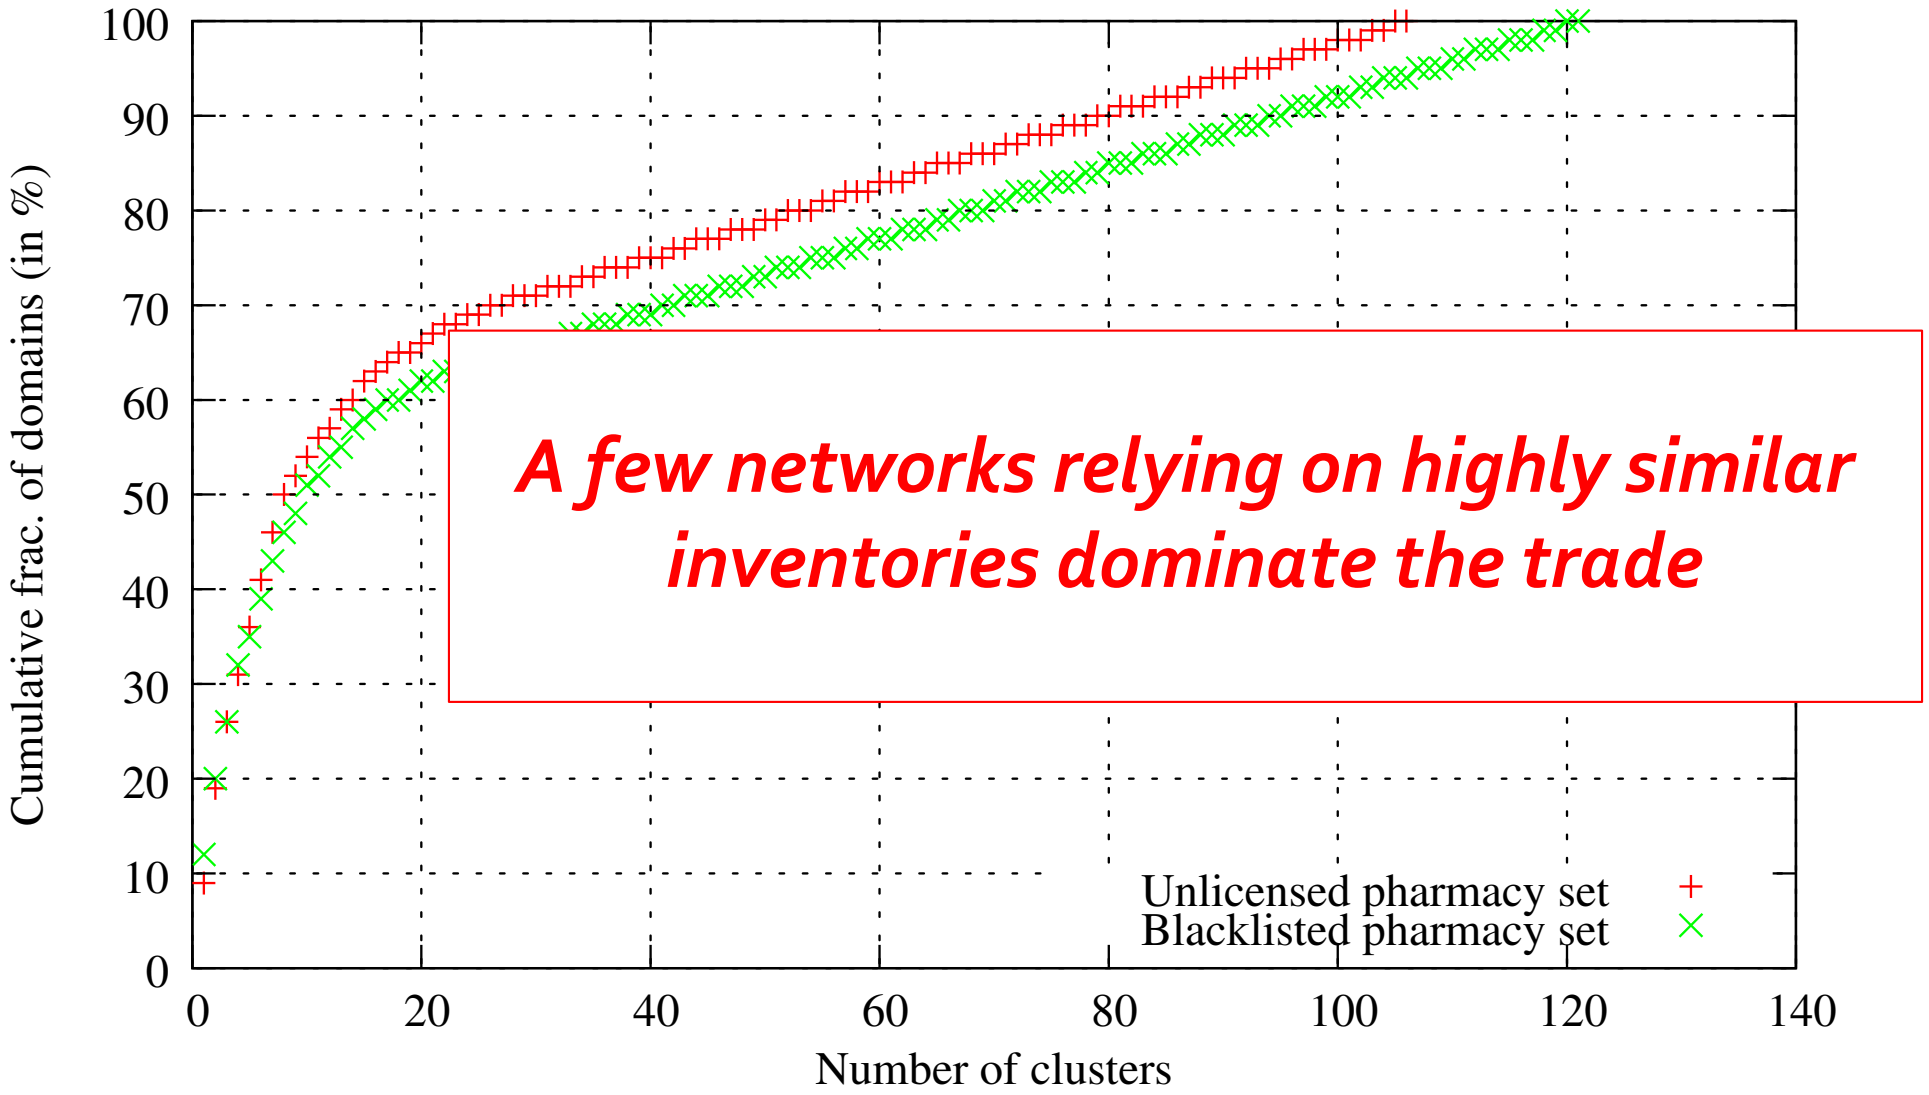

A few networks relying on highly similar inventories dominate the trade

Clustering inventories: Average linkage, *t=0.31*

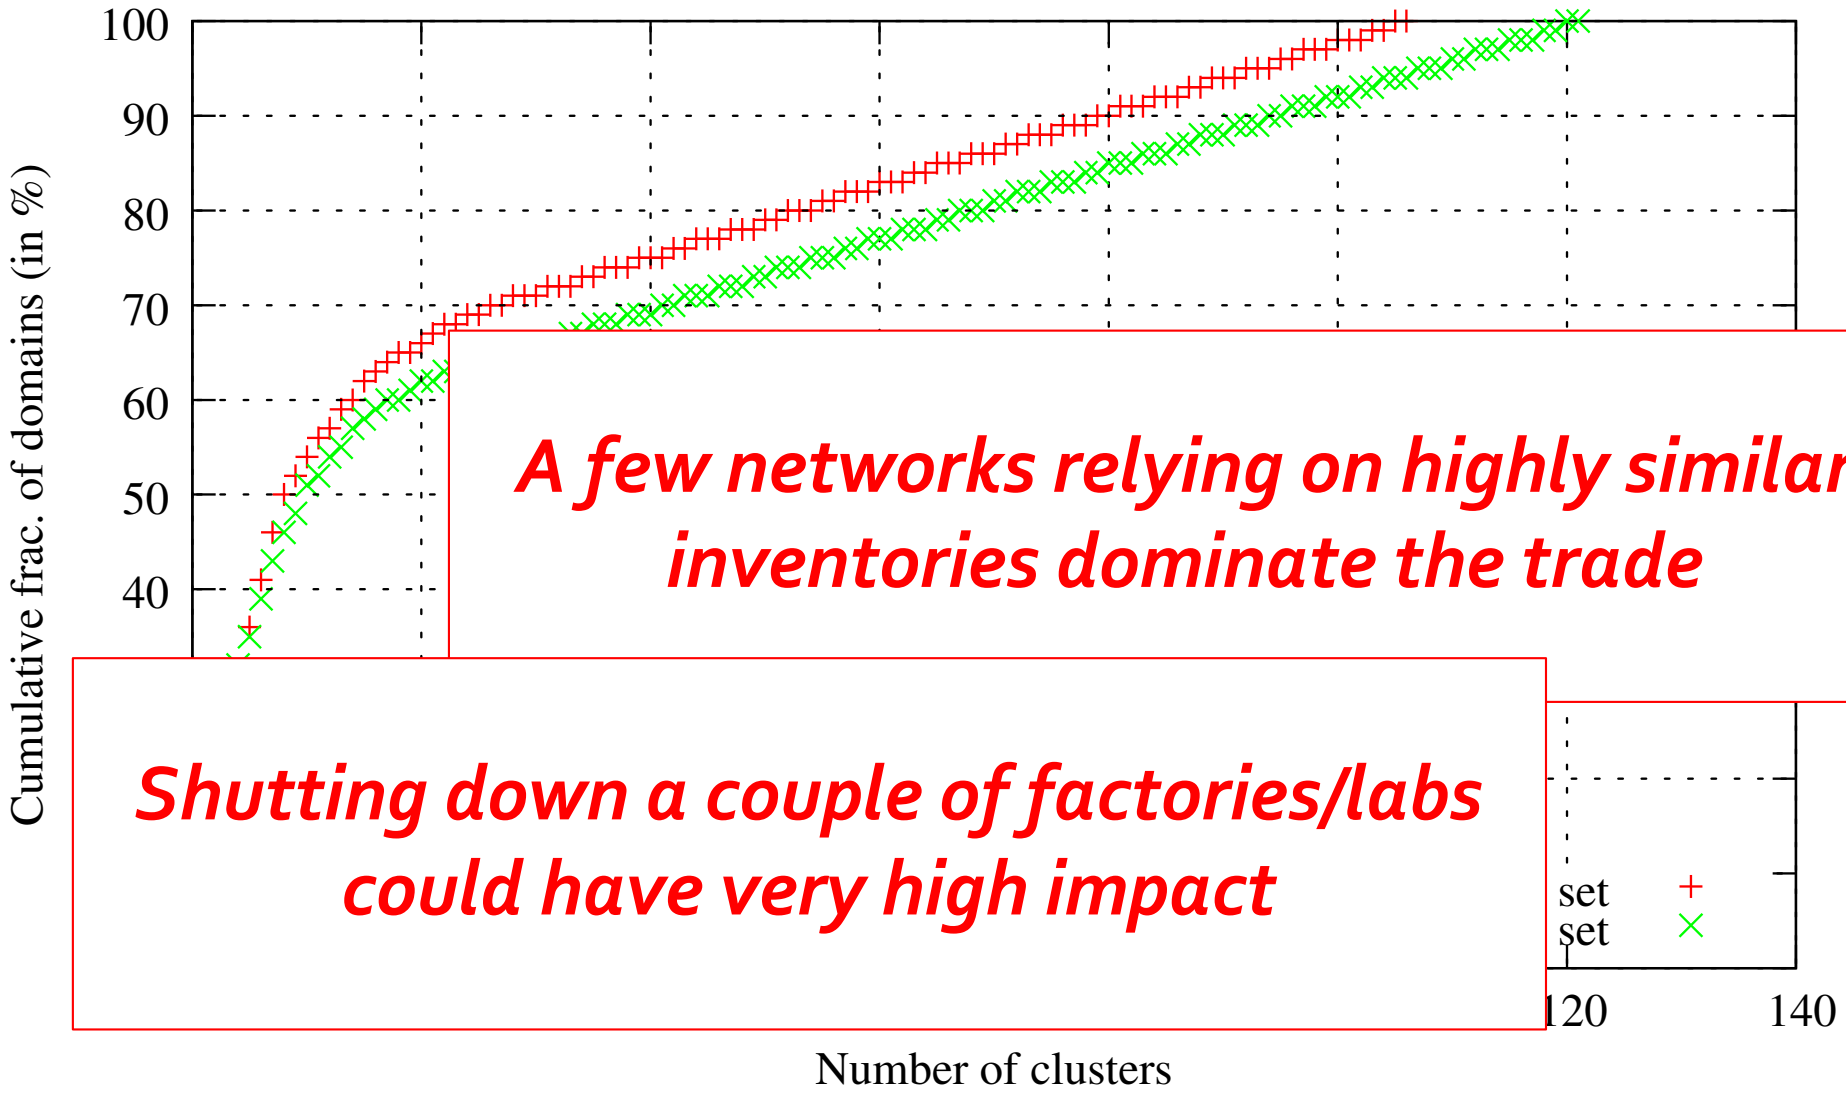

A few networks relying on highly similar inventories dominate the trade

Shutting down a couple of factories/labs could have very high impact

Intervening, or: what does the analysis tell us

- High concentration in traffic brokers
  - Orders of magnitude less numerous than pharmacies and infected hosts
  - Mostly hosted on same networks
  - Structure hasn't changed much over four years
  - Opportunities for takedowns seem ripe
    - Jurisdiction issues?
- High concentration is suppliers (labs)
  - Of strong interest to manufacturers…

Evolution of illicit Internet commerce

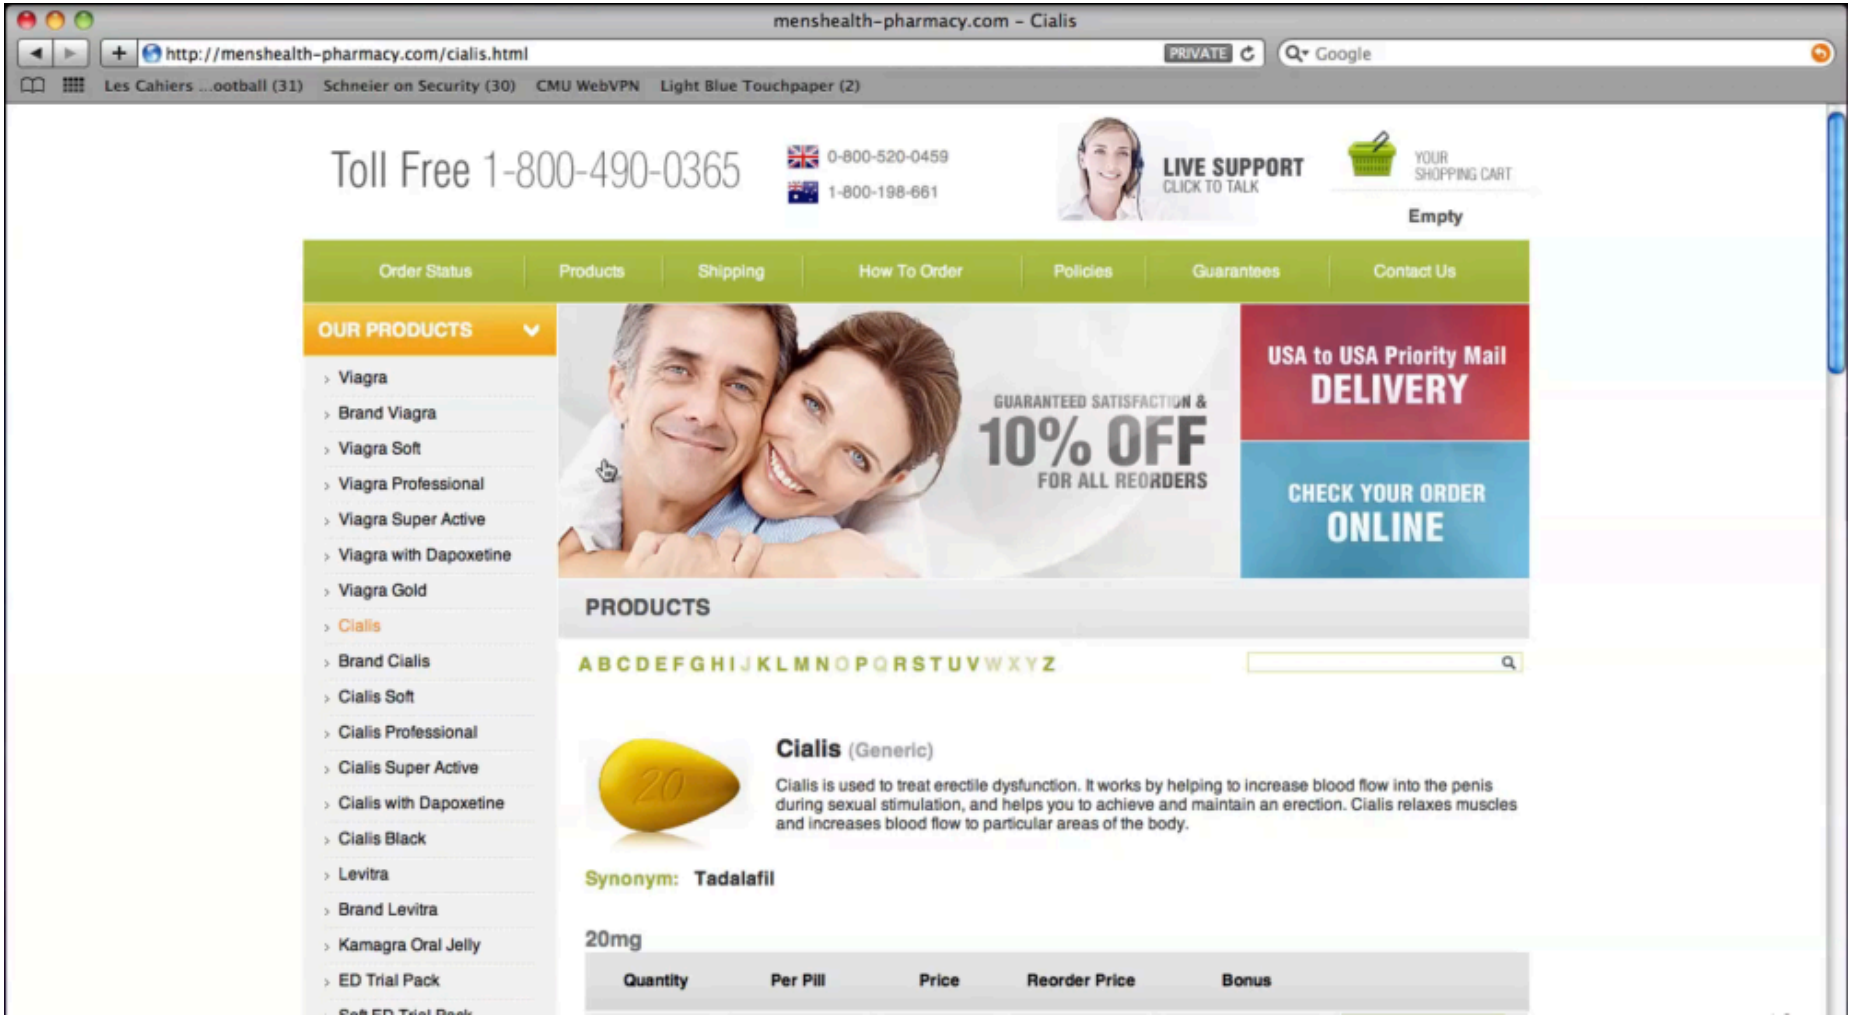

Evolution of illicit Internet commerce

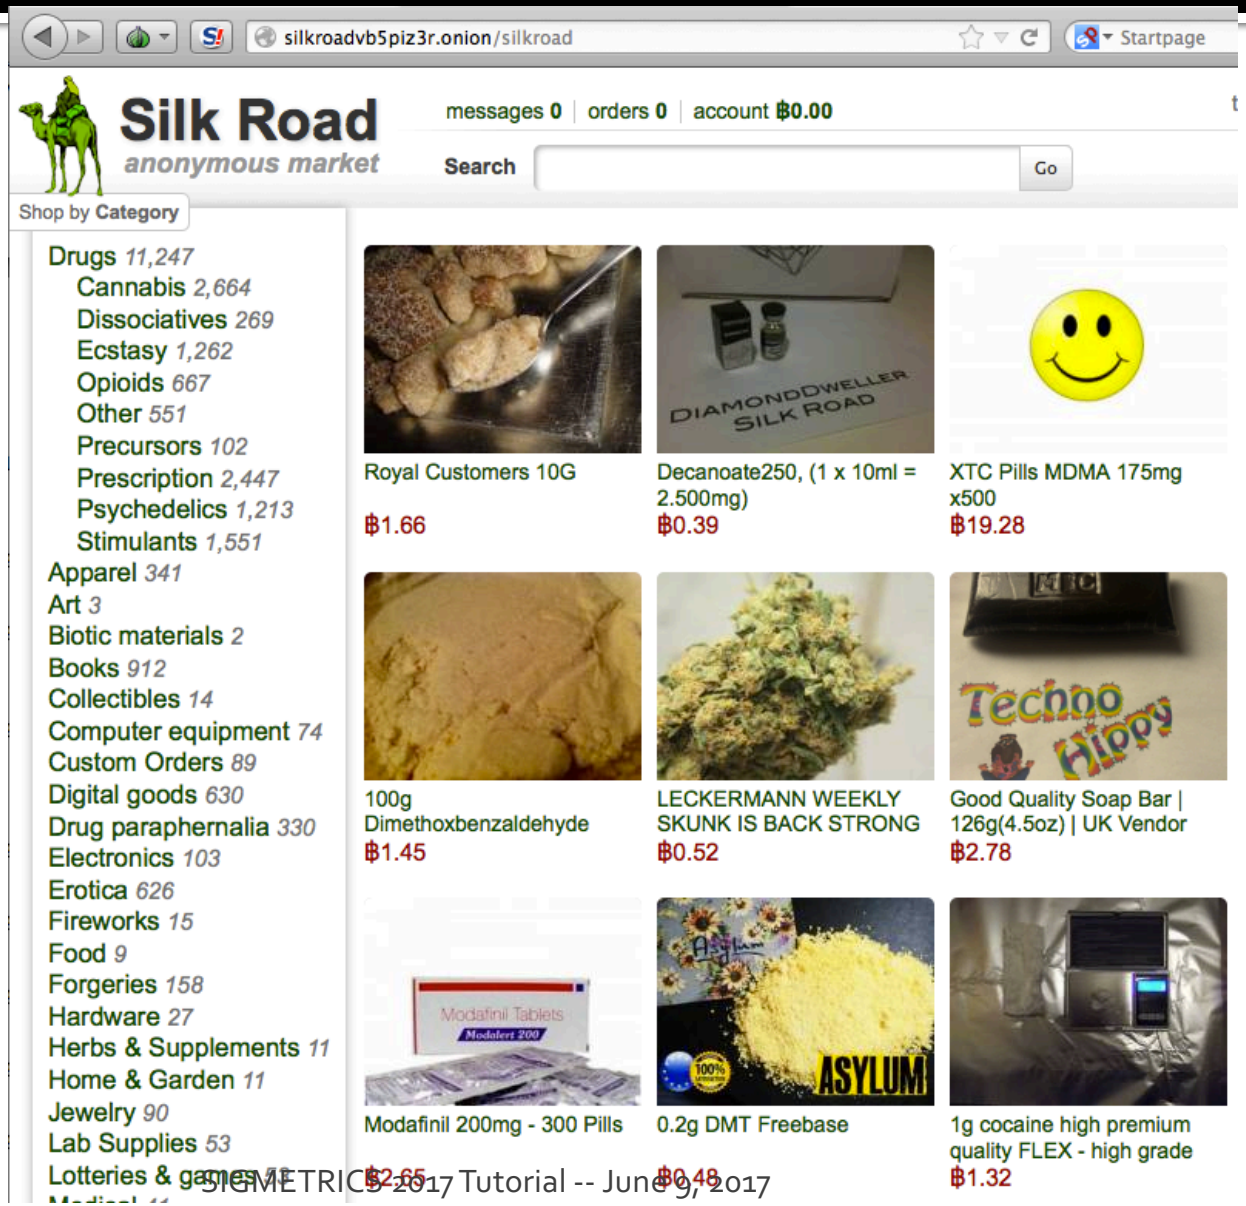

Case study: Online anonymous marketplaces

- Amazon.com of illegal goods
  - Drugs, CC's & Fake IDs, Weapons, etc.
  - No child pornography
- Safety
- Convenience
- Variety
- Accountability
- Competition

Online anonymous marketplace technology

- Hidden Website (Tor Hidden Service, I2P)
  - Customers
    - No cost of creation
    - No information needed
  - Vendors
    - Vendor bonds required
    - Often invite only
    - Public feedback history

- Payments (Bitcoin)
  - Marketplaces often act as escrow agent
  - Escrow sometimes acts as a mixing service

- Encrypted Messages(PGP)

Questions

- How much is being sold?
- What is being sold?
- How many vendors are relevant?
- What are potentially successful interventions?

Typical listing page

Books

Hacking for beginners

Seller:
[redacted] (98)

Price:
฿0.12

Ships from: undeclared
Ships to: Worldwide

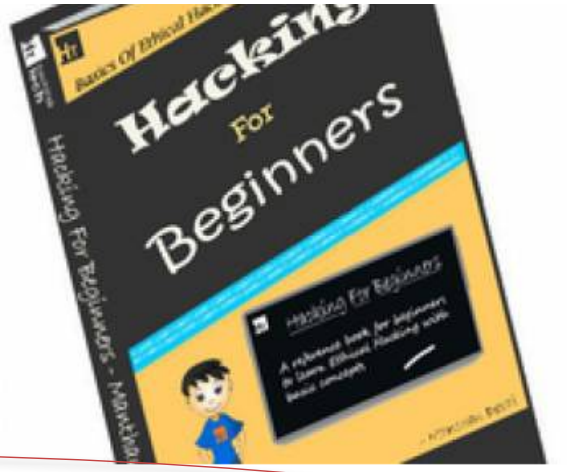

Description:
Hacking For Beginners is a reference book for beginners to learn ethical hacking for free and from basic level to clear all the fundamental concepts of ethical hacking.the book has been prepared by Hacking Tech (www.hackingtech.co.tv) website for the users benefit.so enjoy the book and site...

add to cart

Recent feedback

rating	feedback	freshness
5 of 5	Fast delivery	3 days
5 of 5	Thanks!	4 days
5 of 5	Leave feedback here	9 days
5 of 5	Leave feedback here	9 days
5 of 5	5 of 5	10 days

Feedback is often **mandatory!**

➔Acceptable proxy for sales volume

Measurements

- **Started collection in November 2011**
- **As of August 2015, we had collected**
  - 35 marketplaces
  - 1,908 scrapes total – 3.2 TB
  - 27 – 331,691 pages per scrape
- **Still collecting…**

Data completeness

- **How complete is the data?**
  - Unreliable dynamic marketplaces that take days to scrape
  - Empirical observations – lower bound
- **Idea:** Estimate population via mark and recapture
  - Schnabel estimator allows multiple recapture

Mark and recapture

Population Size = 24

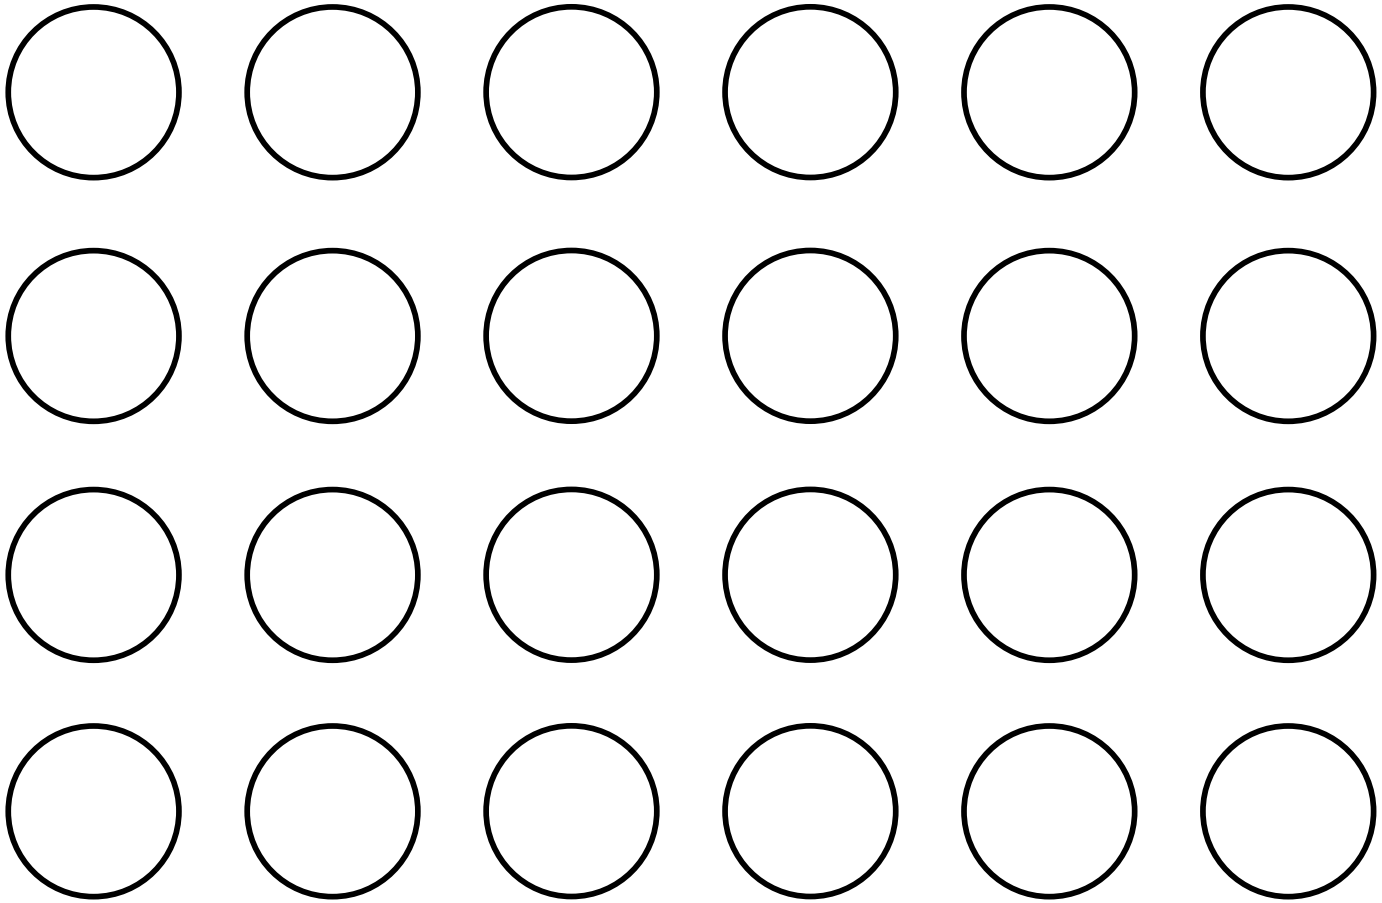

Mark and recapture

Sample Size = 10

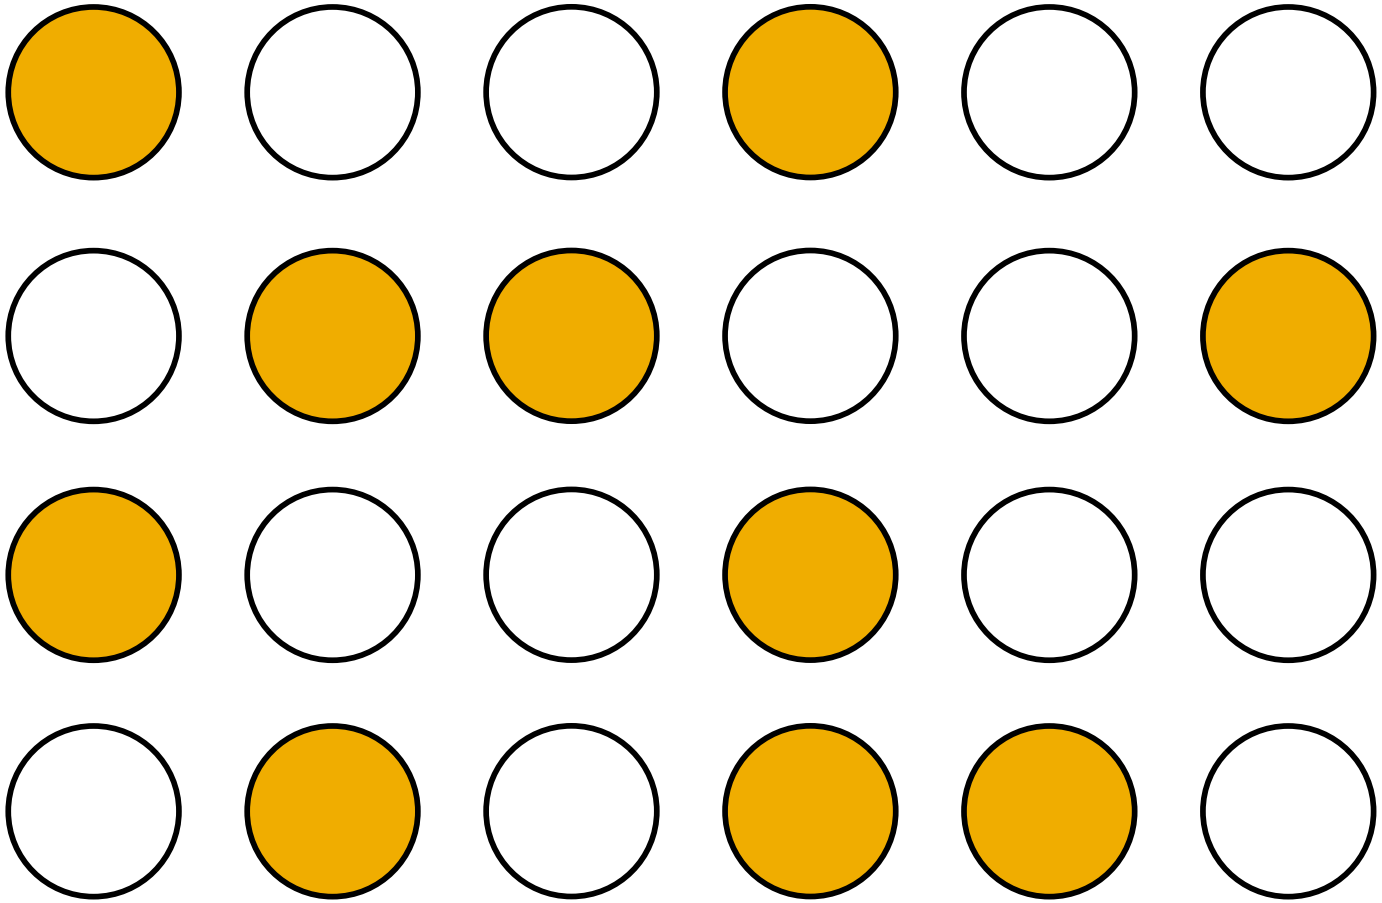

Mark and recapture

Sample Size = 13

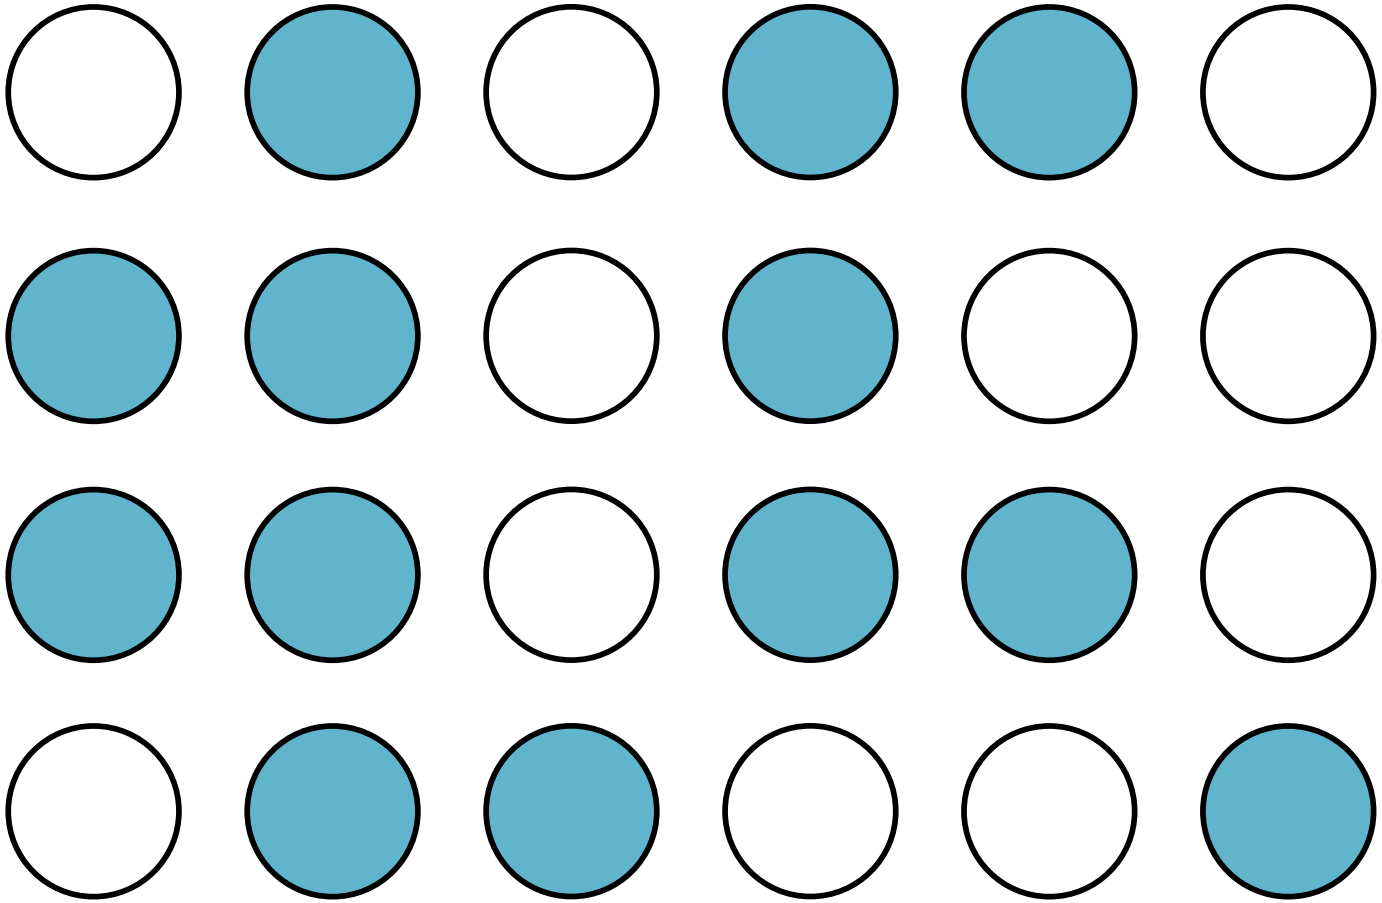

Mark and recapture

Overlap = 5, Population Estimate = 26

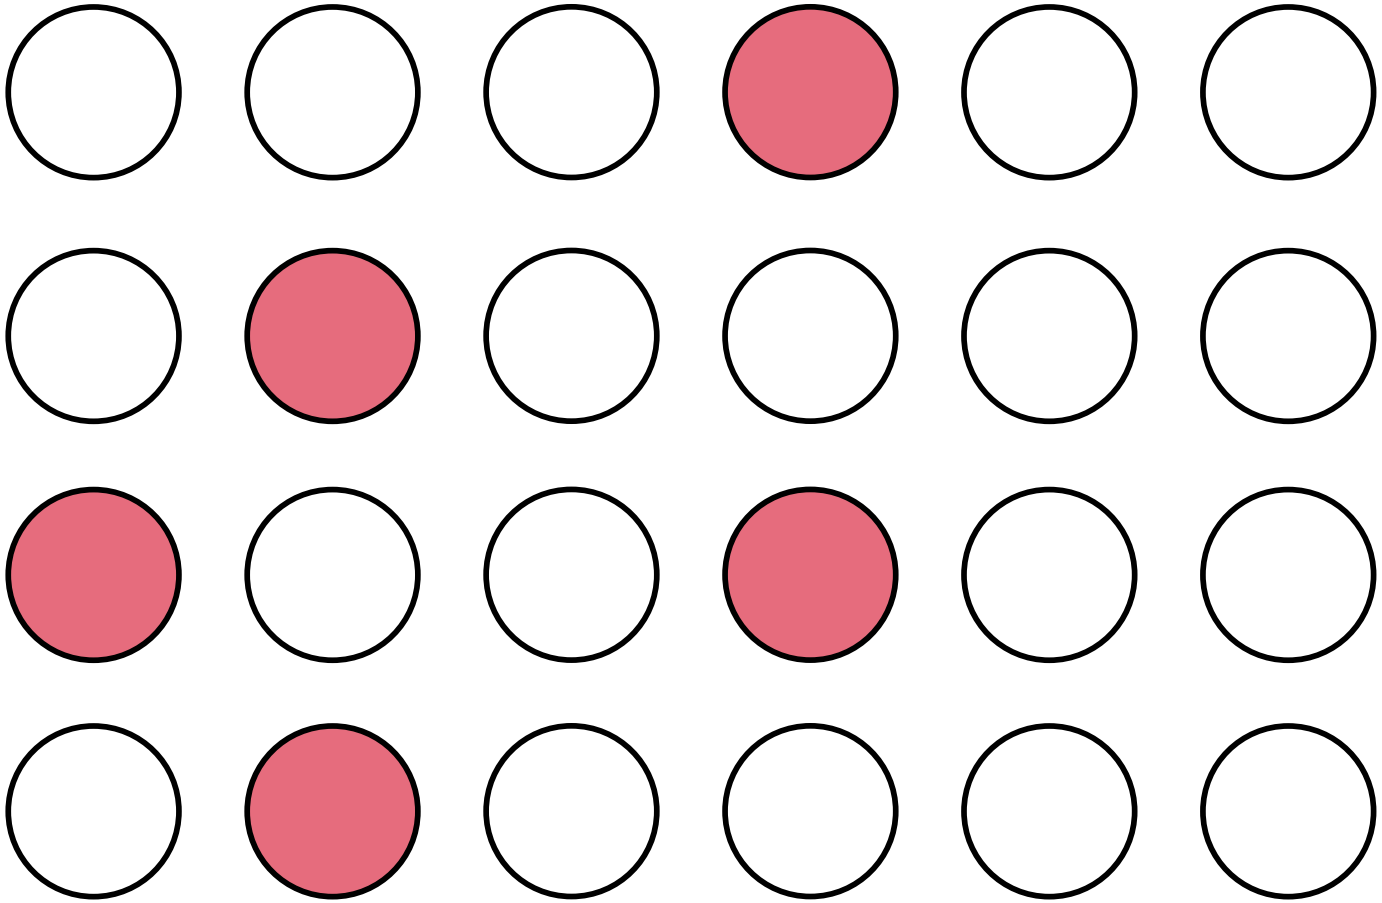

Data completeness

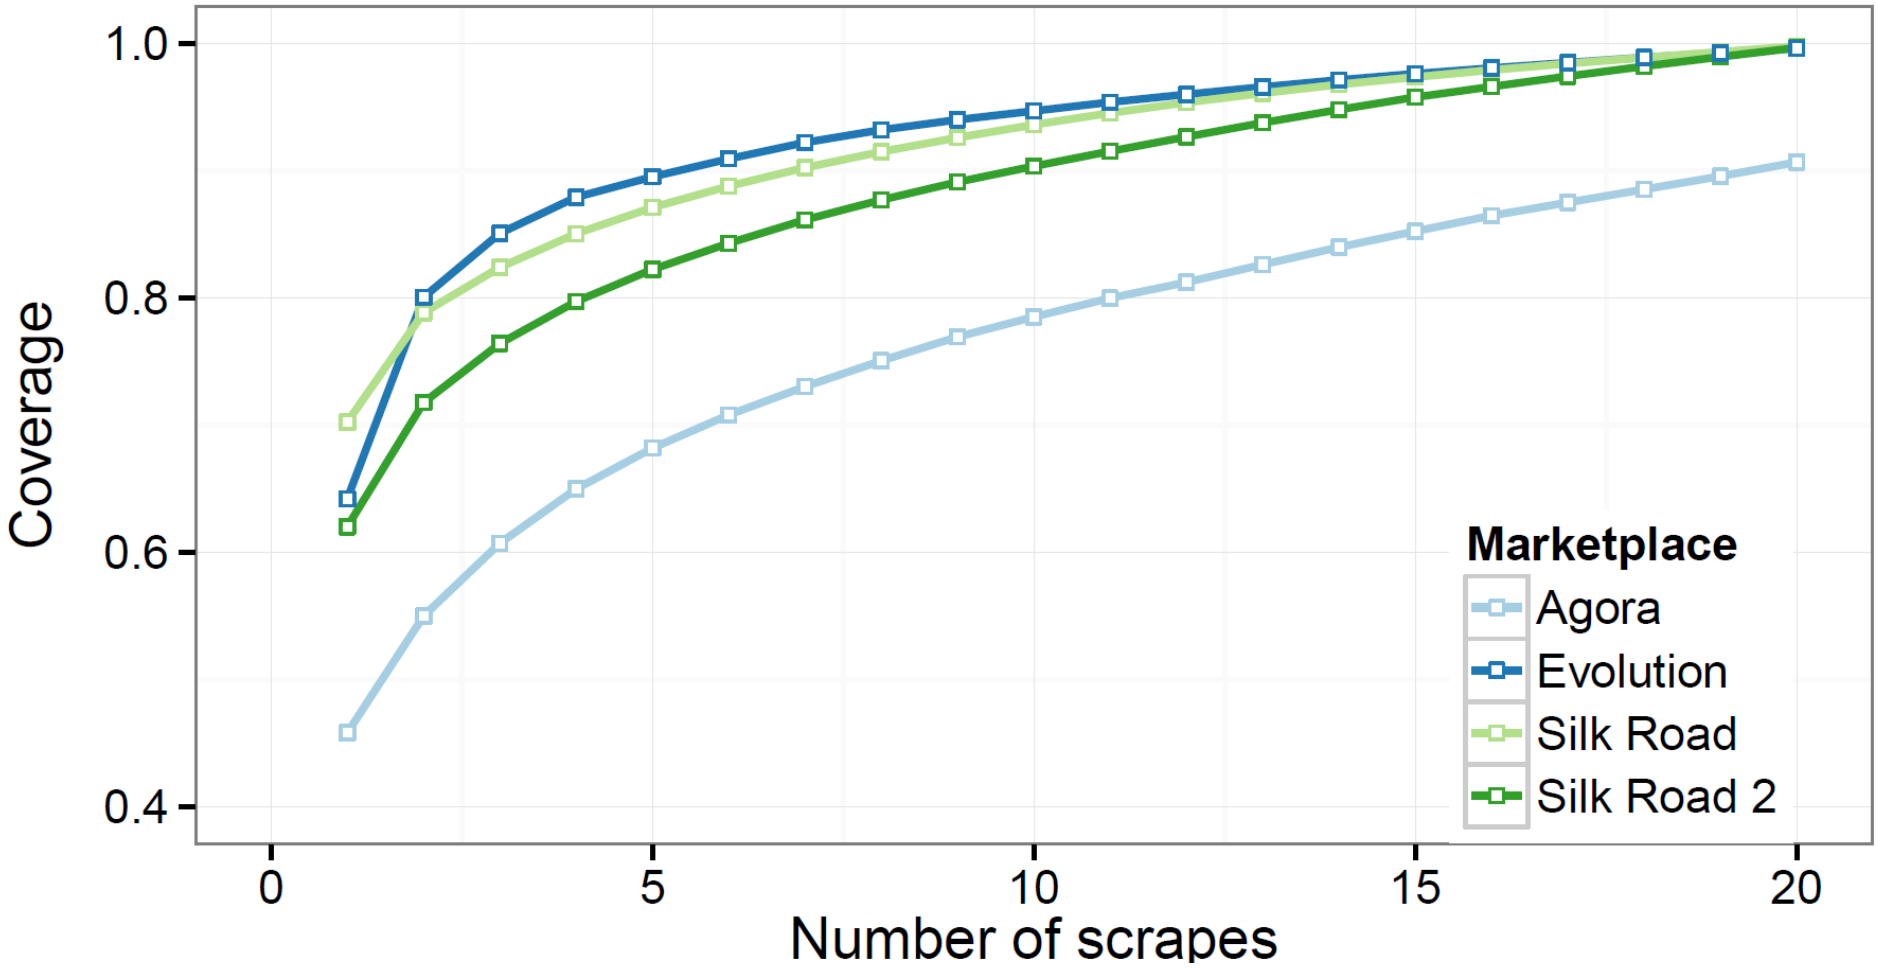

Analysis

- **Assumption:** Each feedback corresponds to precisely one transaction
  - Anonymity requires strictly enforced feedback system to establish reputation
  - Possible on many marketplaces to purchase several quantities of item and leave one feedback, conservative estimate

Analysis challenges

- "Holding prices"
  - Came up with automated statistical filtering of outliers

$0.02 -> $1,000.00

$1,100.00 -> $1,000,000.00

Analysis challenges

- Misleading product categories

  - Define sixteen categories
  - Designed special purpose classifier to infer which category each listing belongs to
    - Extract from tf-idf

Sales volumes

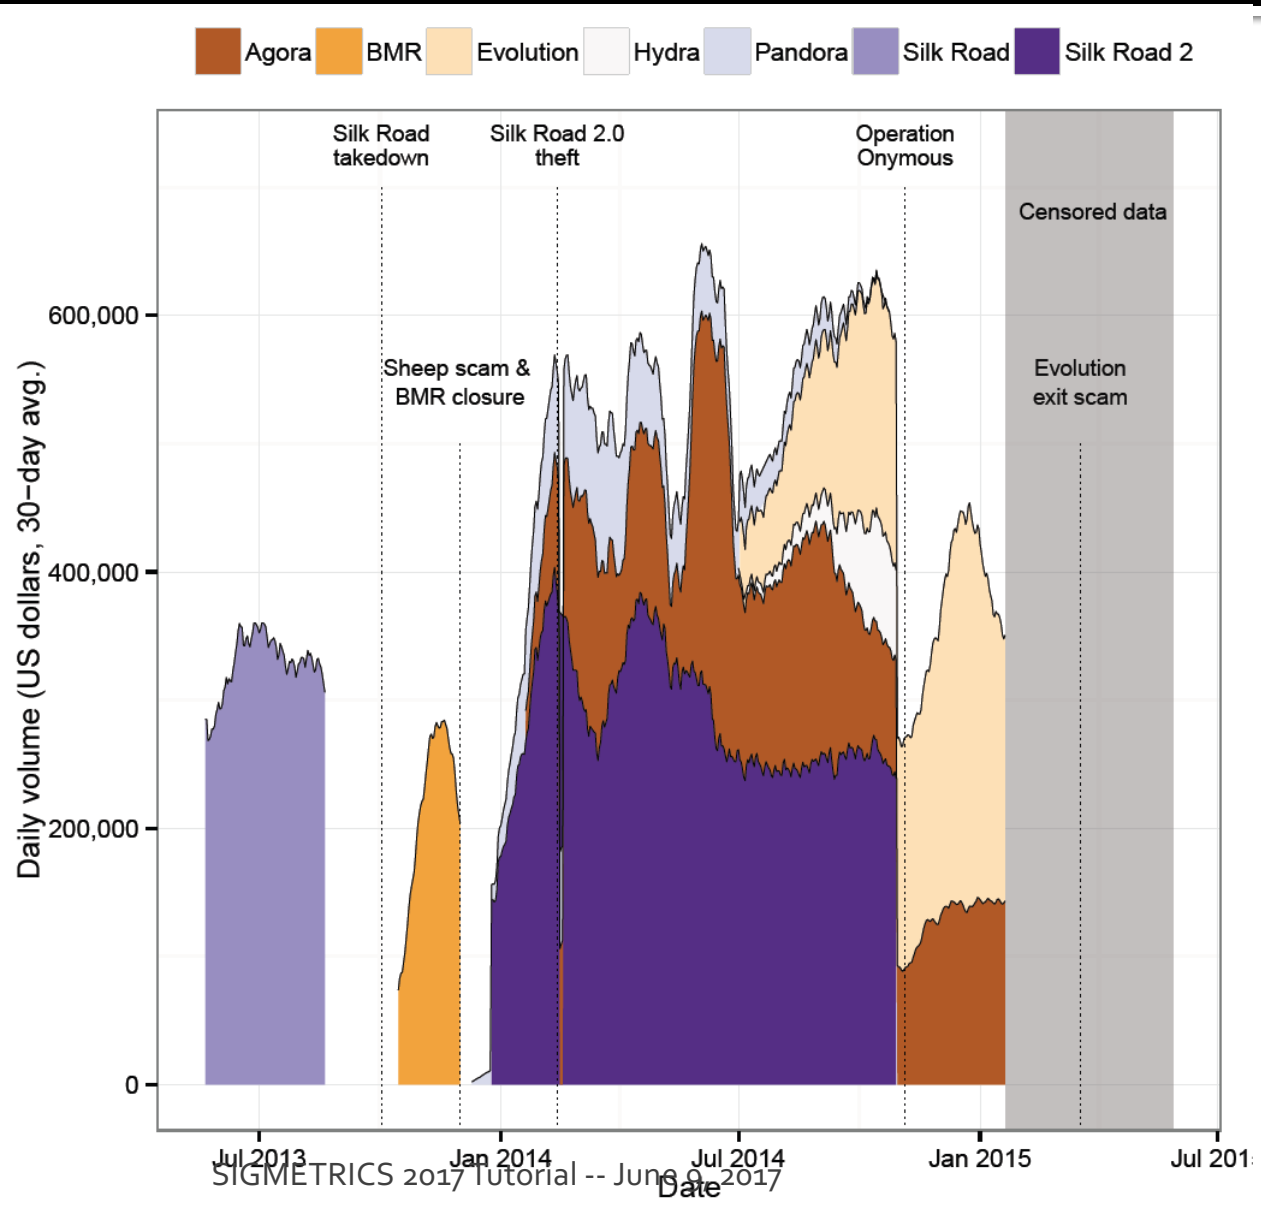

Sales volumes

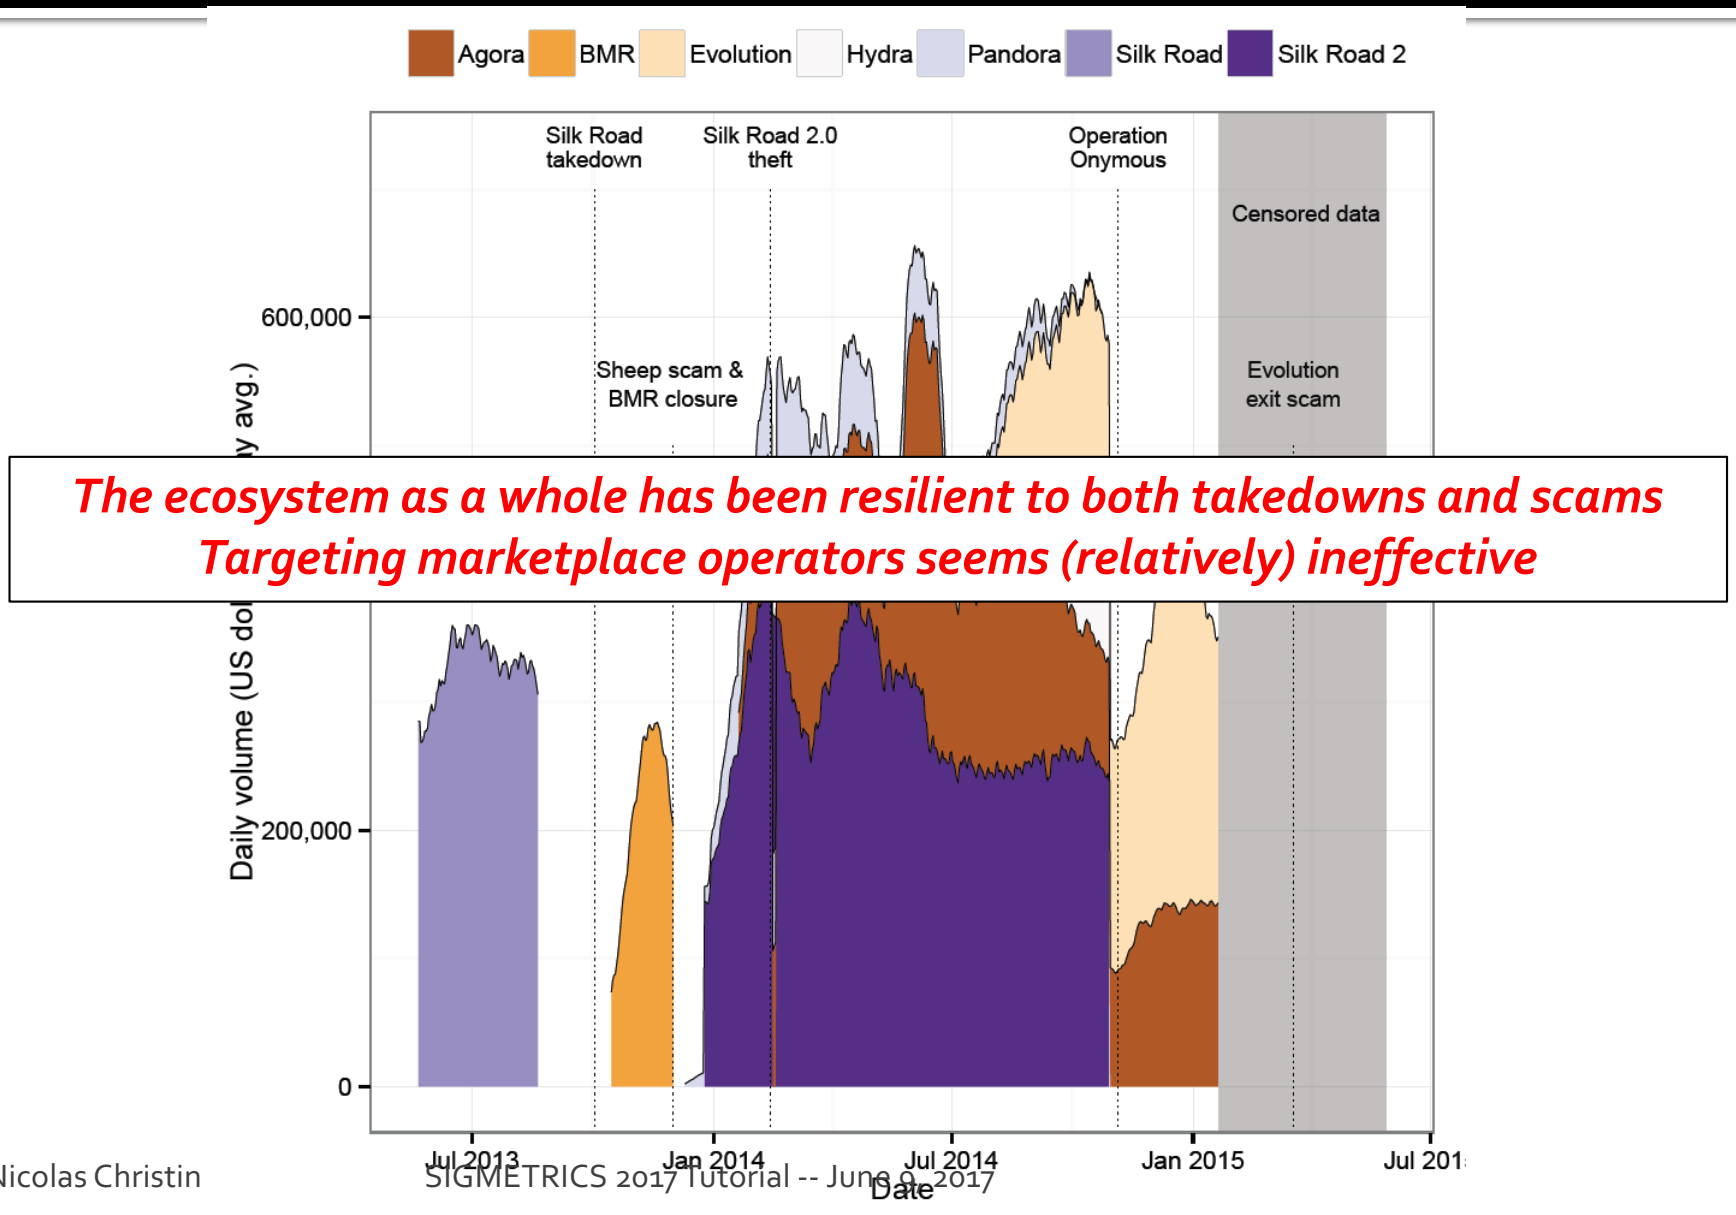

The ecosystem as a whole has been resilient to both takedowns and scams
Targeting marketplace operators seems (relatively) ineffective

Item sales per category

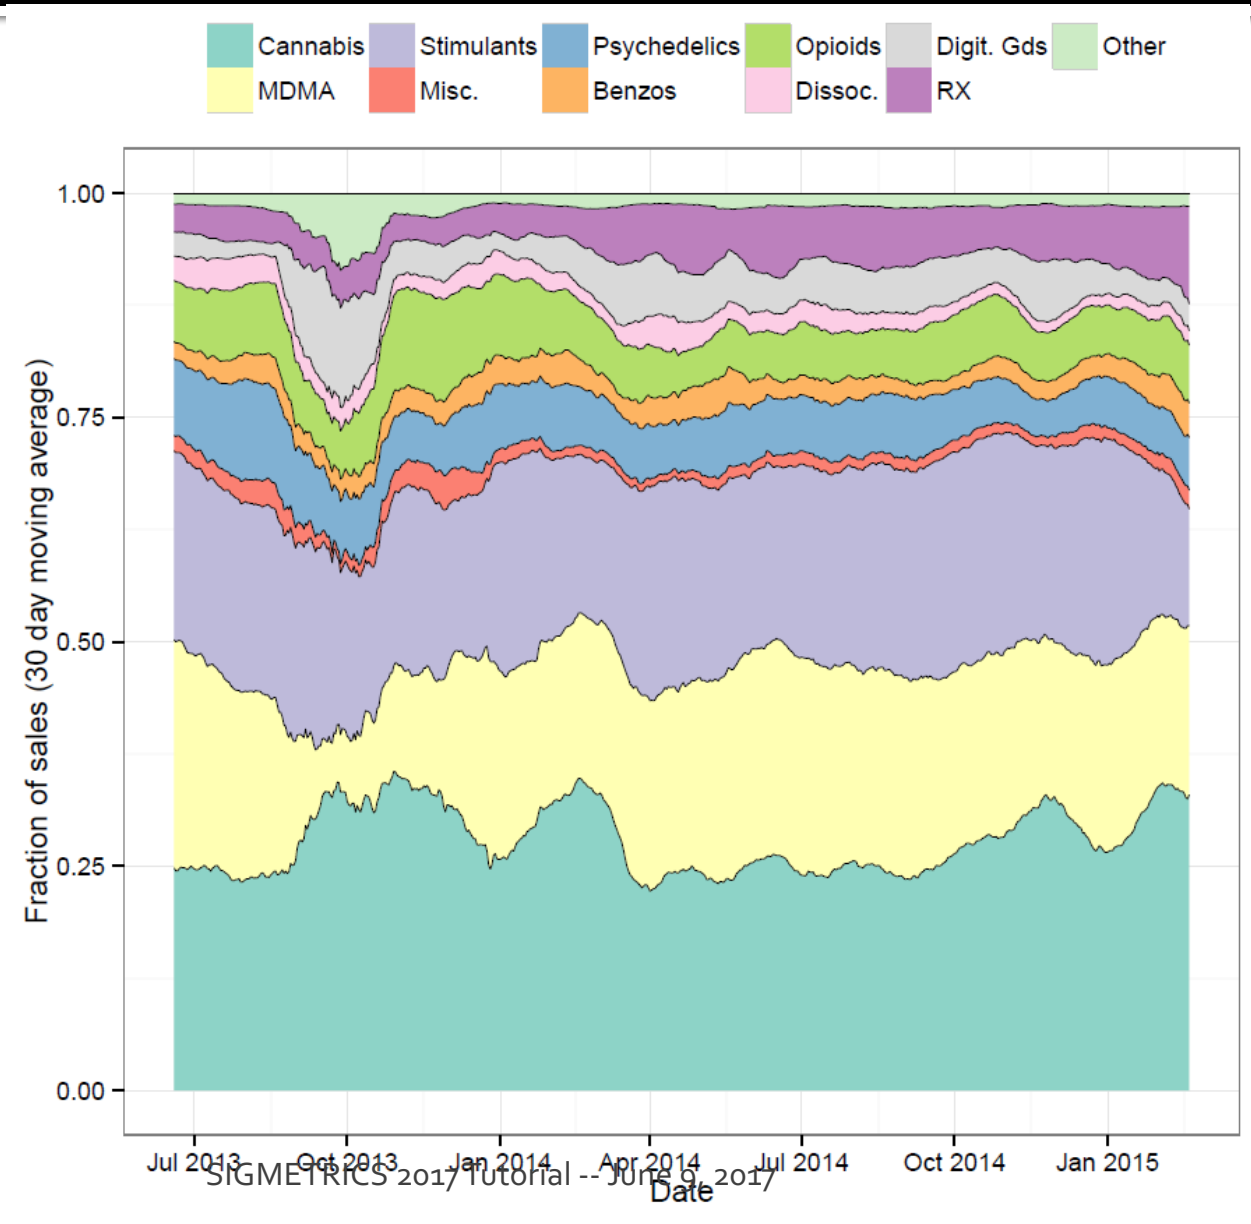

Vendor volumes

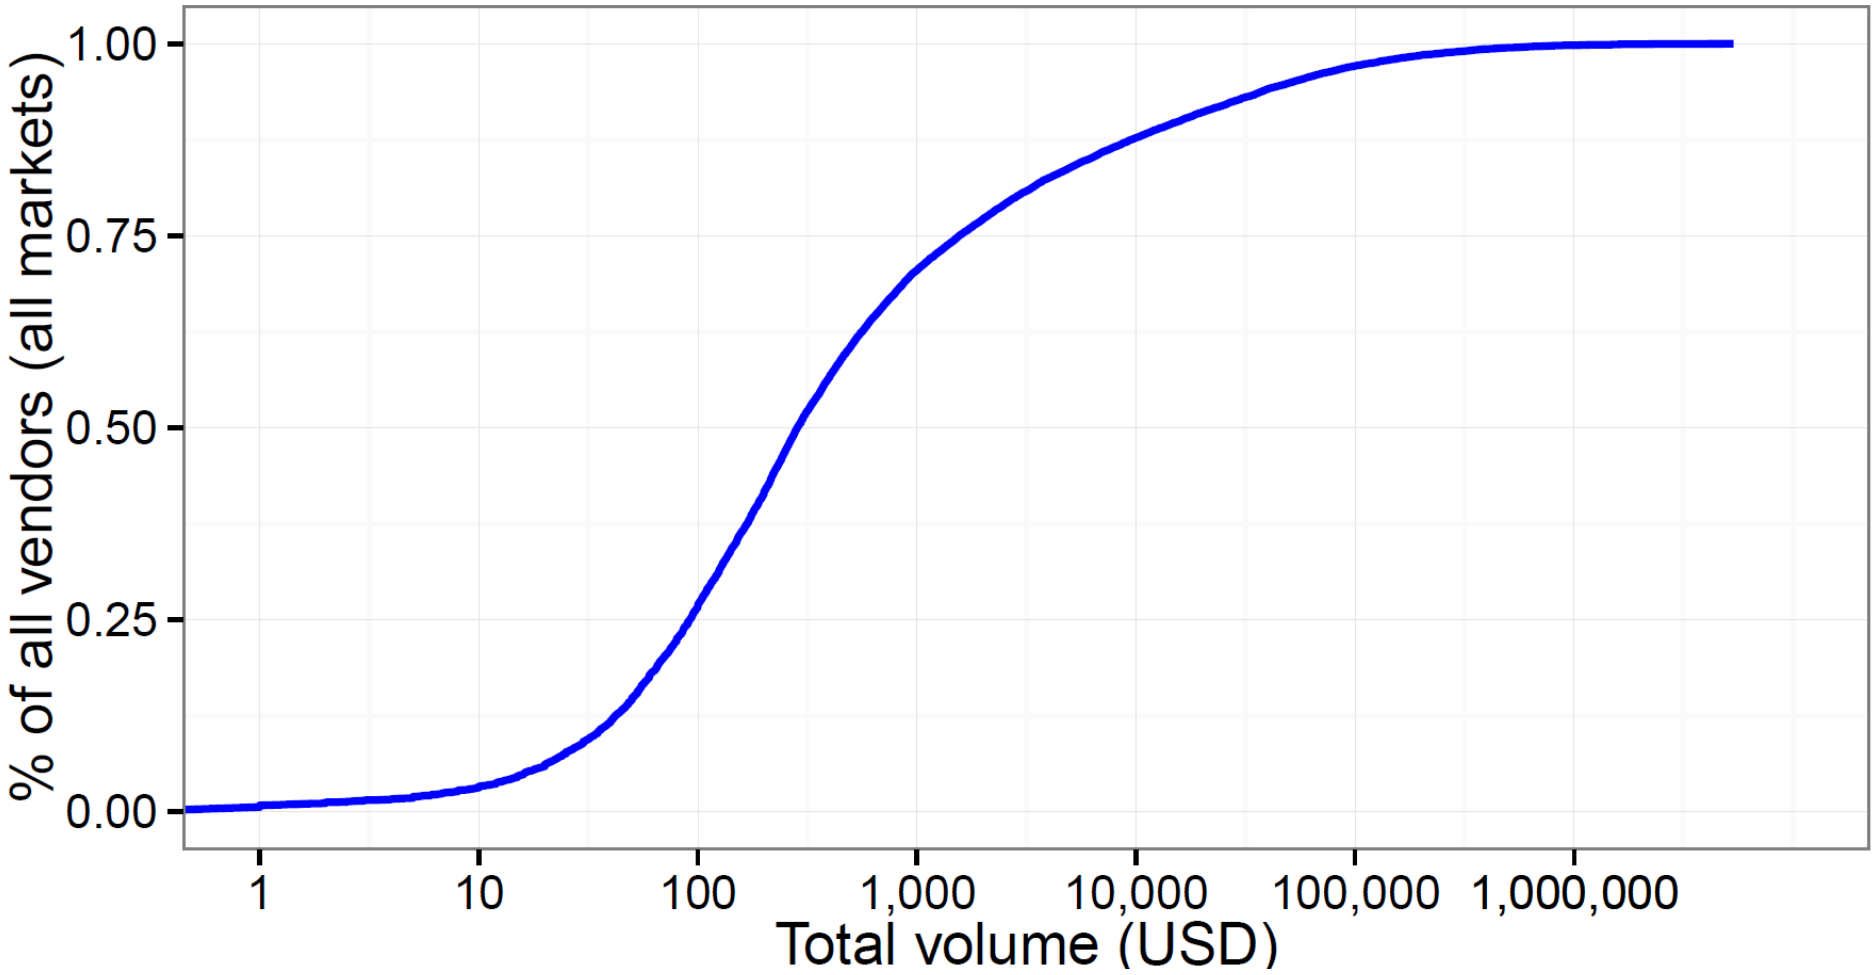

Vendor volumes

Attacker measurements summary

- Collect and analyze data to understand attacker ecosystem and develop better defenses
  - Science of measurement required to understand
  - Emergence of concentrations
    - Traffic brokers & production labs in pharma, large sellers in narcotics…
    - **Driven by economic properties**
    - Possible intervention points
- Ongoing/future work
  - Using our data to build descriptive (mathematical) models of interactions that can then be used to predict future behavior

Final take-aways

Putting it together

- Users (defenders) are near-rational
  - Observed behaviors deviate from predicted models
  - … but not in a random fashion at all
  - Need to incorporate behavioral biases in our models
  - Conjecture: Distance to rationality increases with individual nature of player
    - I.e., institutional actors are likely to be much more rational than individuals
- Increasing amount of data available allows for pretty good modeling of attacker ecosystems
  - Existence of "concentration points"
    - Advertisers, Labs, Major vendors on anonymous marketplaces… → Zipf Law is everywhere
- Attackers closer than defenders to "perfect" rationality
  - Quick reaction to intervention mechanisms can be observed
    - Responses are very rational!
- Furthermore (not in this talk)…
  - Attacks use incentive misalignment **and** behavioral biases

Toward a security economics research agenda

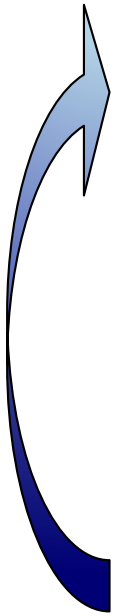

1. Mathematical analysis
   - Game-theoretic predictions, selfishness vs. altruism
   - Impact of various parameters
2. Experimental research
   - Controlled lab experiments
   - Behavioral modeling
3. Field data measurement
   - Acquisition of attacker data
   - Acquisition of investment patterns
4. Testing intervention mechanisms
   - Incentives, legal…
   - Mostly through simulation?

Nicolas Christin
nicolasc@cmu.edu / @nc2y
https://www.andrew.cmu.edu/user/nicolasc