Forecasting Cybersecurity Incidents and Its Role in Designing Incentive Mechanisms

Mingyan Liu

Joint work with

Yang Liu, Armin Sarabi, Parinaz Naghizadeh, Michael Bailey

Motivation

Increasingly frequent and high-impact data breaches

- Target, JP Morgan Chase, Home Depot, Anthem, to name a few
- Increasing social and economic impact of such cyber incidents
  - 95% increase in average cost of from 2010 to 2014

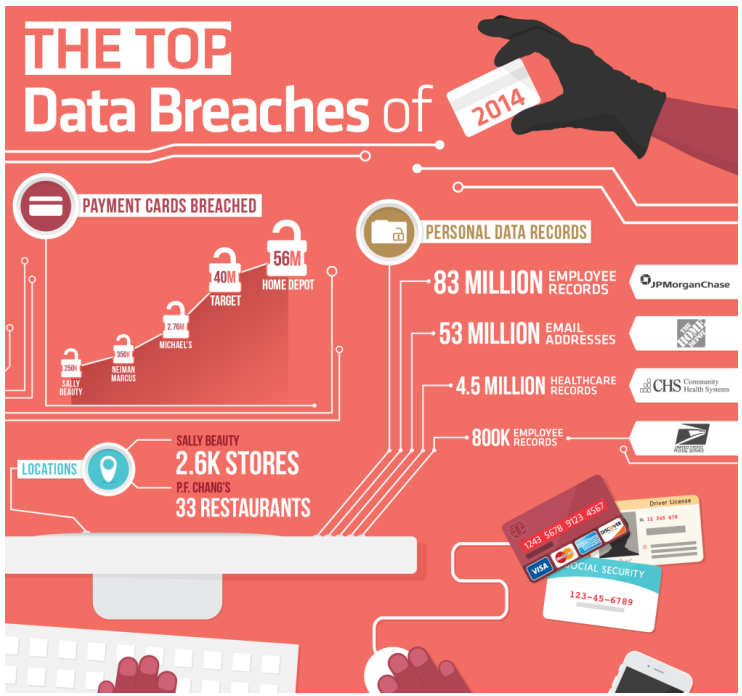

Current approaches

- Heavily *detection* based
- Even when successful, maybe or too late by the time a breach is detected
- Damage control post breach
- Clear need for more *proactive* measures pre breach

Current approaches

- Heavily *detection* based
- Even when successful, maybe or too late by the time a breach is detected
- Damage control post breach
- Clear need for more *proactive* measures pre breach

Detection

- analogous to diagnosing a patient who may already be ill (e.g., via biopsy).
- [Qian et al. NDSS14, Wang et al. USENIX Sec14]

Prediction

- predicting whether a presently healthy person may become ill based on a variety of relevant factors.
- [Soska & Christin, USENIX Sec14]

Objective

Understand the extent to which one can *forecast* incidents at an organizational level.

Objective

Understand the extent to which one can *forecast* incidents at an organizational level.

Desirable features:

- *Scalability:* we rely solely on *externally* observed data.

Objective

Understand the extent to which one can *forecast* incidents at an organizational level.

Desirable features:

- *Scalability:* we rely solely on *externally* observed data.
- *Robustness:* data will be noisy, incomplete, not all of which is under our control.

Objective

Understand the extent to which one can *forecast* incidents at an organizational level.

Desirable features:

- *Scalability:* we rely solely on *externally* observed data.
- *Robustness:* data will be noisy, incomplete, not all of which is under our control.

Key steps:

- Tap into a *diverse* set of data that captures different aspects of a network's security posture: source, type (*explicit* vs. *latent*).
- Follow a supervised learning framework.

Takeaway from this talk

If you are interested in cybersecurity

Takeaway from this talk

If you are interested in cybersecurity

- This is the right time to apply data analytics to make new contributions.
  - An abundance of data; need domain expertise to make sense of.
  - Good analysis can inform policy design, opening up new areas.

Takeaway from this talk

If you are interested in cybersecurity

- This is the right time to apply data analytics to make new contributions.
  - An abundance of data; need domain expertise to make sense of.
  - Good analysis can inform policy design, opening up new areas.

If you are only interested in the underlying methodology

Takeaway from this talk

If you are interested in cybersecurity

- This is the right time to apply data analytics to make new contributions.
  - An abundance of data; need domain expertise to make sense of.
  - Good analysis can inform policy design, opening up new areas.

If you are only interested in the underlying methodology

- This is a good case study to highlight some of the real challenges in applying machine learning techniques.
  - Data is rarely readily available: they are misaligned, grossly incomplete, with various unknown errors/biases.
  - But if you do come out the other end, the results can be very rewarding; you might even get ideas on how to further the methodology.

Outline of the talk

- **Data and Preliminaries**
  - Data sources
  - Pre-processing
- Forecasting method and results
  - Feature extraction
  - Construction of the classifier
  - Prediction performance
- Fine-grained prediction
- Risk assessment as a form of "public monitoring"

Security posture data

Malicious Activity Data: a set of 11 reputation blacklists (RBLs)

- Daily collections of IPs seen engaged in some malicious activity.
- Three malicious activity types: spam, phishing, scan.

Security posture data

Malicious Activity Data: a set of 11 reputation blacklists (RBLs)

- Daily collections of IPs seen engaged in some malicious activity.
- Three malicious activity types: spam, phishing, scan.

Mismanagement symptoms

- Deviation from known best practices; indicators of lack of policy or expertise:
  - Misconfigured- HTTPS cert, DNS (resolver+source port), mail server, BGP.

Cyber incident Data

Three incident datasets

- Hackmageddon
- Web Hacking Incidents Database (WHID)
- VERIS Community Database (VCDB)

Incident type	SQLi	Hijacking	Defacement	DDoS
Hackmageddon	38	9	97	59
WHID	12	5	16	45
Incident type	Crimeware	Cyber Esp.	Web app.	Else
VCDB	59	16	368	213

Datasets at a glance

Category	Collection period	Datasets
Mismanagement symptoms	Feb'13 - Jul'13	Open Recursive Resolvers, DNS Source Port, BGP misconfiguration, Untrusted HTTPS, Open SMTP Mail Relays
Malicious activities	May'13 - Dec'14	CBL, SBL, SpamCop, UCEPROTECT, WPBL, SURBL, PhishTank, hpHosts, Darknet scanners list, Dshield, OpenBL
Incident reports	Aug'13 - Dec'14	VERIS Community Database, Hackmageddon, Web Hacking Incidents

- Mismanagement and malicious activities used to extract features.
- Incident reports used to generate labels for training and testing.

Data pre-processing

Conservative processing of incident reports:

- Remove irrelevant or ambiguous cases, e.g., robbery at liquor store, "something happened", etc.

Data pre-processing

Conservative processing of incident reports:

- Remove irrelevant or ambiguous cases, e.g., robbery at liquor store, "something happened", etc.

Challenge in data alignment, both in time and in space:

- Security posture records information at the host IP-address level.
- Cyber incident reports associated with an organization.
- Alignment non-trivial: address reallocation, hosting services, etc.

Data pre-processing

Conservative processing of incident reports:

- Remove irrelevant or ambiguous cases, e.g., robbery at liquor store, "something happened", etc.

Challenge in data alignment, both in time and in space:

- Security posture records information at the host IP-address level.
- Cyber incident reports associated with an organization.
- Alignment non-trivial: address reallocation, hosting services, etc.

A mapping process:

- Using maintainer/owner IDs from RIR databases.
- 4.4 million prefixes listed under 2.6 million owner IDs.
- Sample IP from organization + search in maintainer table.
- Other alternatives with different granularity.

Outline of the talk

- Data and Preliminaries
  - Data sources
  - Pre-processing
- **Forecasting method and results**
  - Feature extraction
  - Construction of the classifier
  - Prediction performance
- Fine-grained prediction
- Risk assessment as a form of "public monitoring"

Primary features: raw data

Mismanagement symptoms (5).

- Five symptoms; each measured as a fraction
- Predictive power of these symptoms.

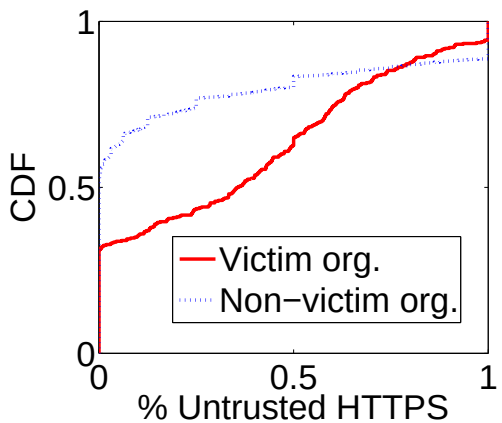

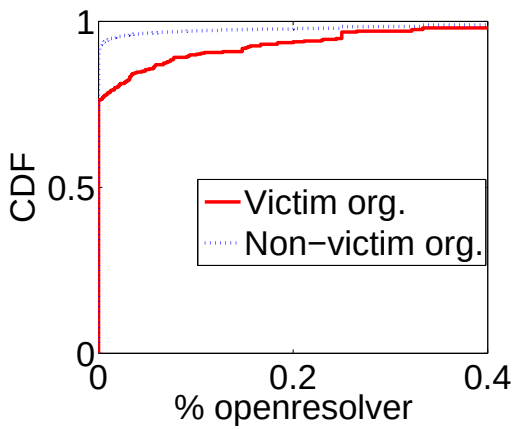

Malicious activity time series (60×3).

- Three time series over a period: spam, phishing, scan.
- Recent 60 v.s. Recent 14.

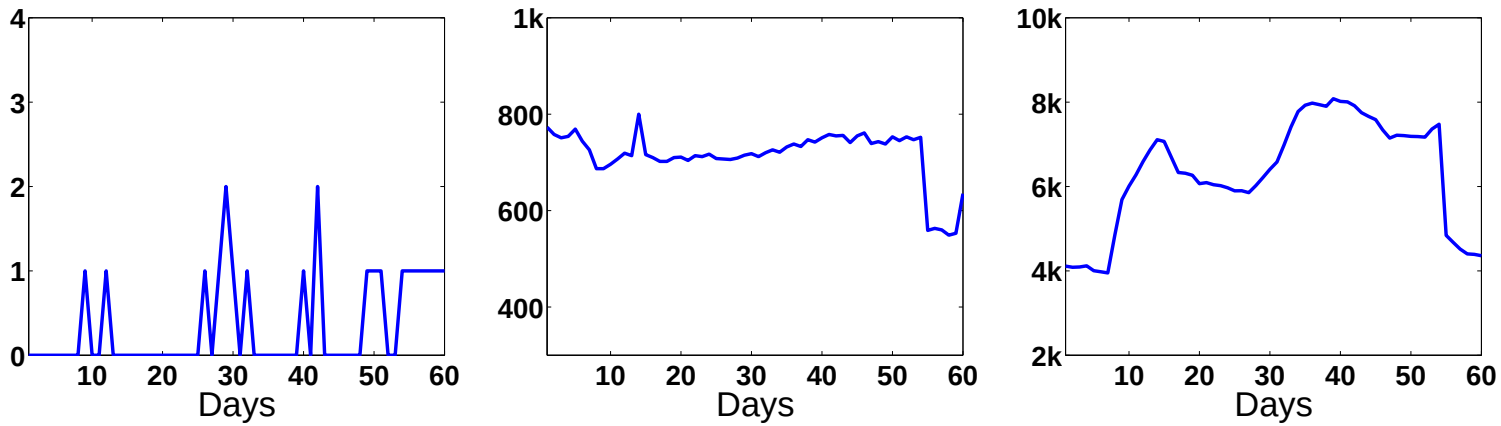

Malicious activity time series (60×3).

- Three time series over a period: spam, phishing, scan.
- Recent 60 v.s. Recent 14.

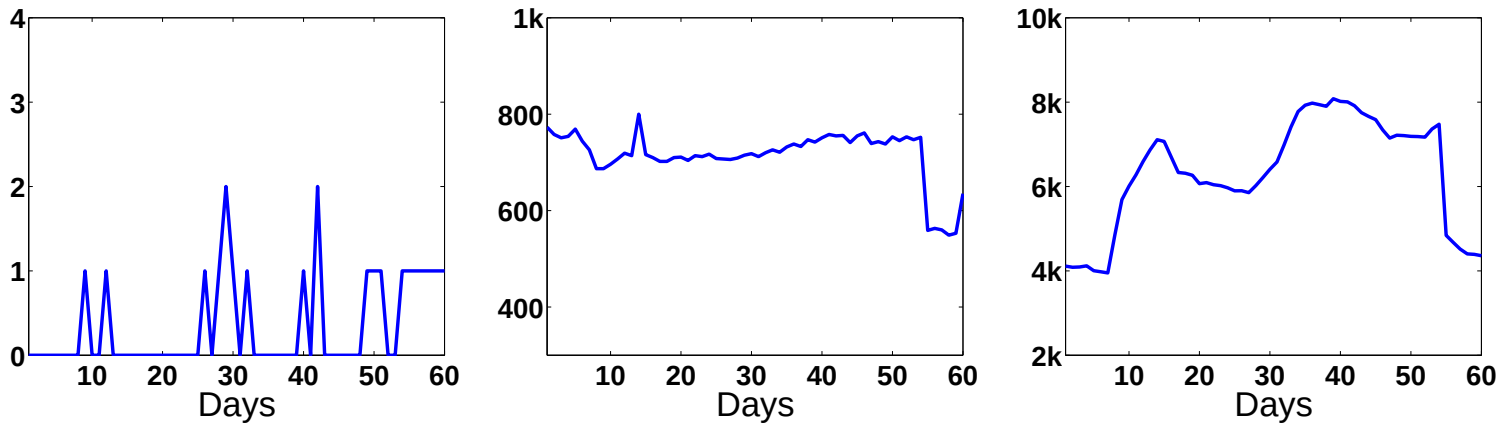

Size: number of IPs in an aggregation unit (1)

- To some extent captures the likelihood of an organization becoming a target of/reporting attacks.

Secondary features

Quantization and second order statistics

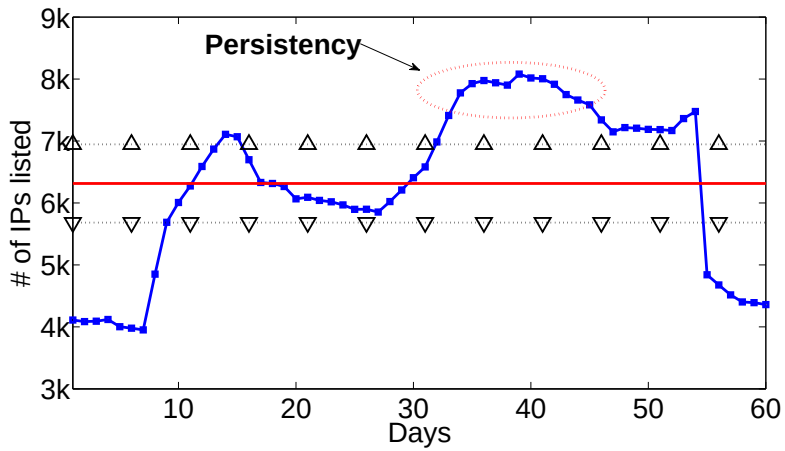

- Measure security efforts and responsiveness.
- In each quantized region, measure average magnitude, average duration, and frequency.

A look at their predictive power (using data from Nov-Dec'13):

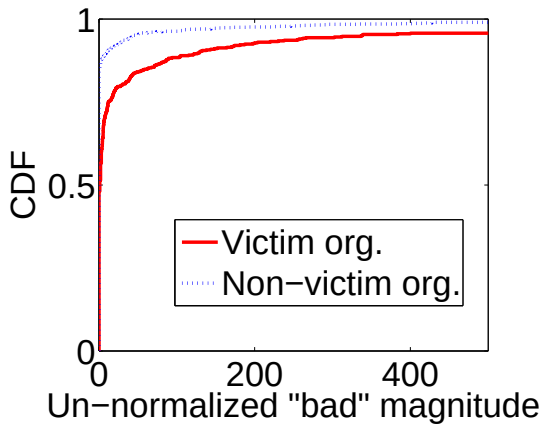

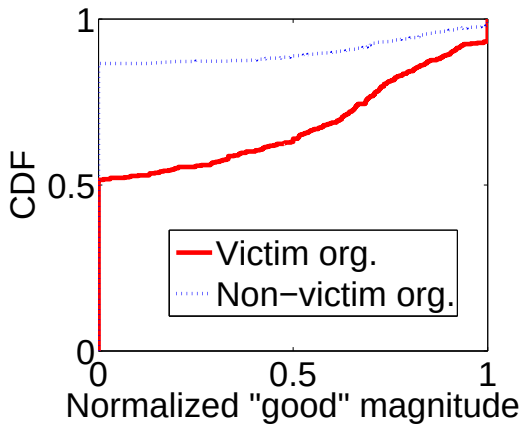

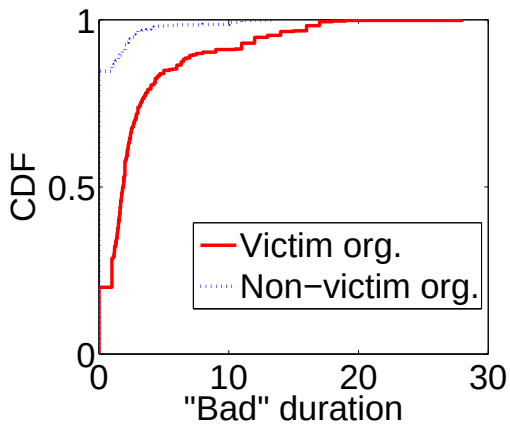

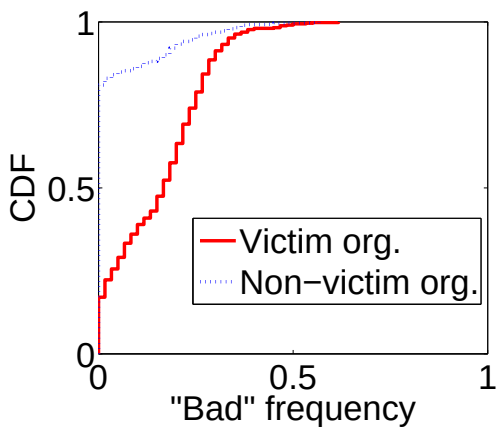

Approach at a glance

Feature extraction

- 258 features extracted from the datasets: Primary + Secondary features.

Approach at a glance

Feature extraction

- 258 features extracted from the datasets: Primary + Secondary features.

Label generation

- 1,000+ incident reports from the three incident sets

Approach at a glance

Feature extraction

- 258 features extracted from the datasets: Primary + Secondary features.

Label generation

- 1,000+ incident reports from the three incident sets

Classifier training and testing

- Random Forest (RF) classifier trained with features and labels.

Training subjects

A subset of victim organizations, or incident group.

- Training-testing ratio, e.g., **70-30** or **50-50** split .
- Split strictly according to time: use *past* to predict *future*.

	Hackmageddon	VCDB	WHID
Training	Oct 13 – Dec 13	Aug 13 – Dec 13	Jan 14 – Mar 14
Testing	Jan 14 – Feb 14	Jan 14 – Dec 14	Apr 14 – Nov 14

Training subjects

A subset of victim organizations, or incident group.

- Training-testing ratio, e.g., **70-30** or **50-50** split .
- Split strictly according to time: use *past* to predict *future*.

	Hackmageddon	VCDB	WHID
Training	Oct 13 – Dec 13	Aug 13 – Dec 13	Jan 14 – Mar 14
Testing	Jan 14 – Feb 14	Jan 14 – Dec 14	Apr 14 – Nov 14

A random subset of non-victims, or non-incident group.

- Random sub-sampling necessary to avoid imbalance; procedure is repeated over different random subsets.

Prediction procedure

Long term prediction

Short term prediction

Prediction procedure

Long term prediction

Short term prediction

Prediction procedure

Long term prediction

Short term prediction

Prediction performance

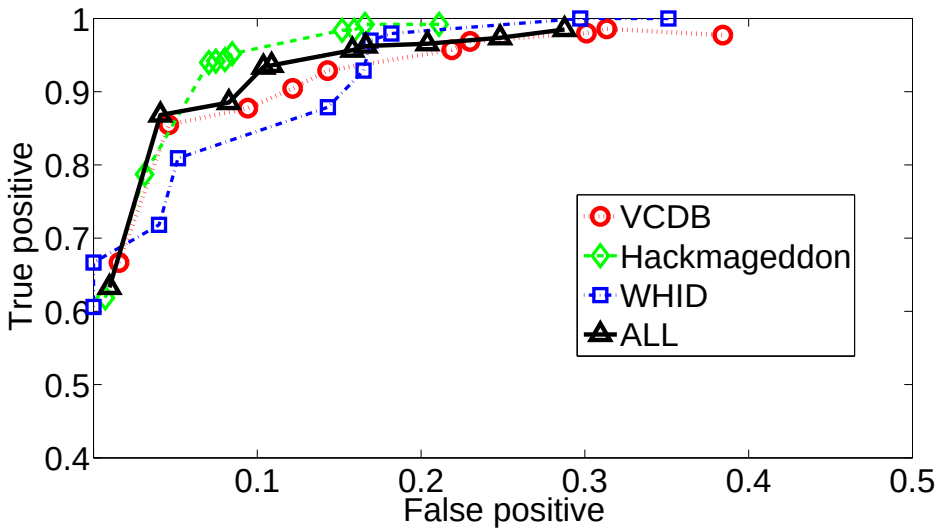

Example of desirable operating points of the classifier:

Accuracy	Hackmageddon	VCDB	WHID	All
True Positive (TP)	96%	88%	80%	88%
False Positive (FP)	10%	10%	5%	4%

Split ratio

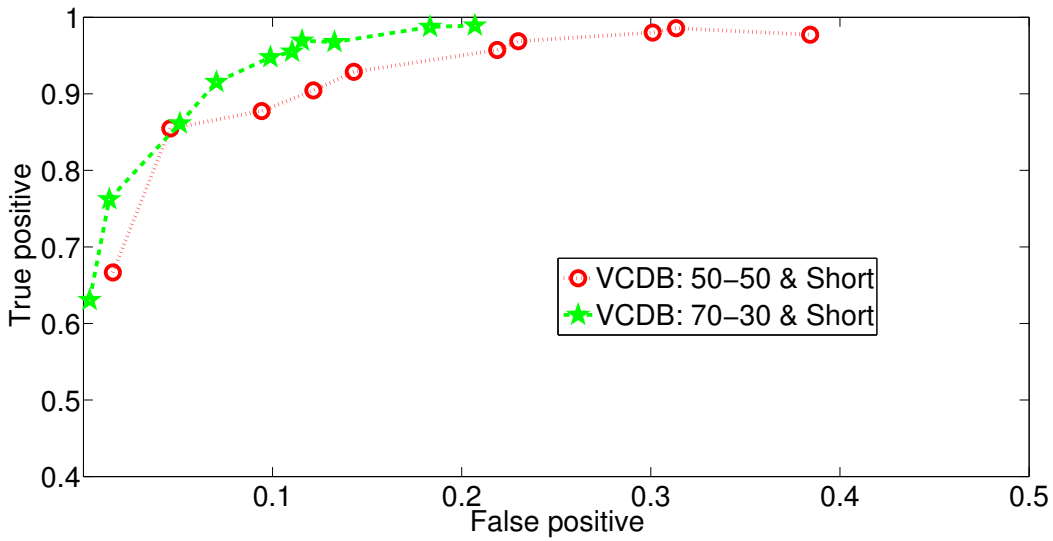

More training data gives better performance.

Short term v.s. long term prediction

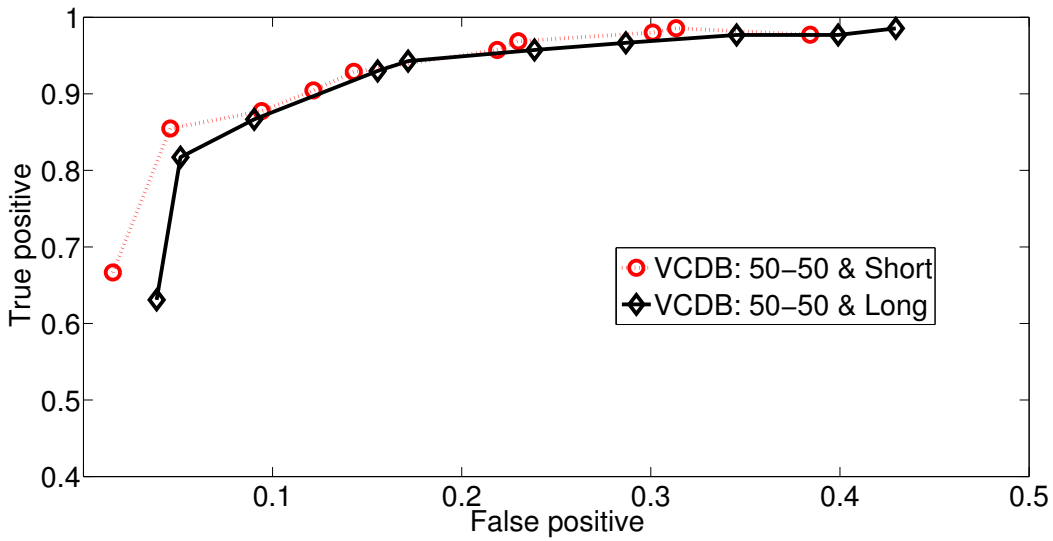

Temporal features become slighted outdated.

Importance of the Features

Top feature descriptor	Value
Untrusted HTTPS Certificates	0.1531
Frequency	0.1089
Organization size	0.0976
Open recursive resolver	0.0928

- Two mismgmt features rank in top 4.

Importance of the Features

Top feature descriptor	Value
Untrusted HTTPS Certificates	0.1531
Frequency	0.1089
Organization size	0.0976
Open recursive resolver	0.0928

- Two mismgmt features rank in top 4.

Feature category	Normalized importance
Mismanagement	0.3229
Time series data	0.2994
Recent-60 secondary features	0.2602

- Secondary features almost as important as time series data.

Importance of the Features

Top feature descriptor	Value
Untrusted HTTPS Certificates	0.1531
Frequency	0.1089
Organization size	0.0976
Open recursive resolver	0.0928

- Two mismgmt features rank in top 4.

Feature category	Normalized importance
Mismanagement	0.3229
Time series data	0.2994
Recent-60 secondary features	0.2602

- Secondary features almost as important as time series data.
- Dynamic features more important than static features.

The power of data diversity

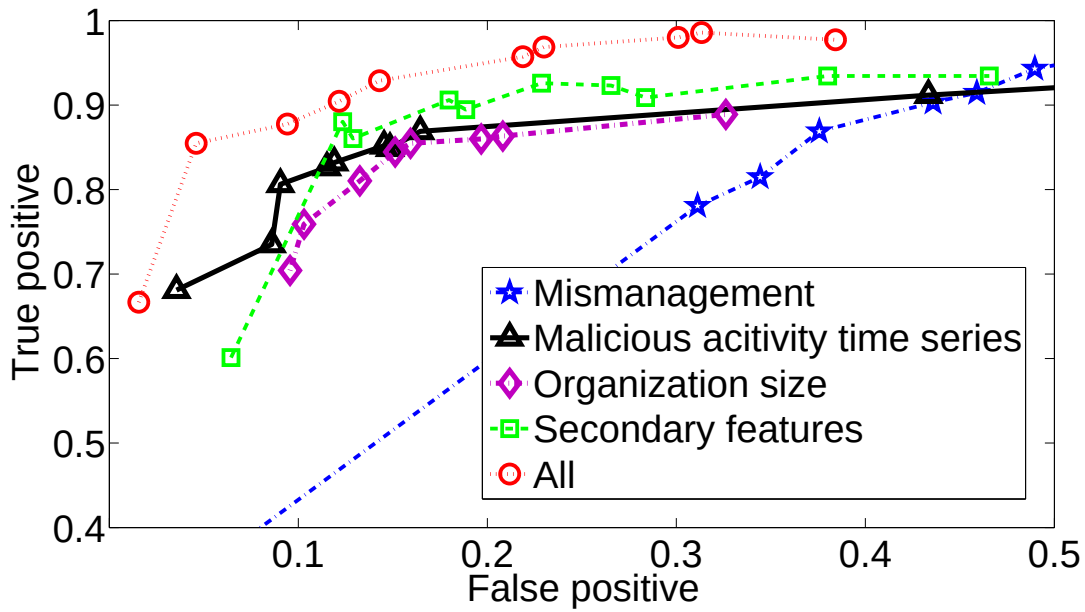

Any single data source does not hold sufficient predictive power

Case study: top data breaches of 2015

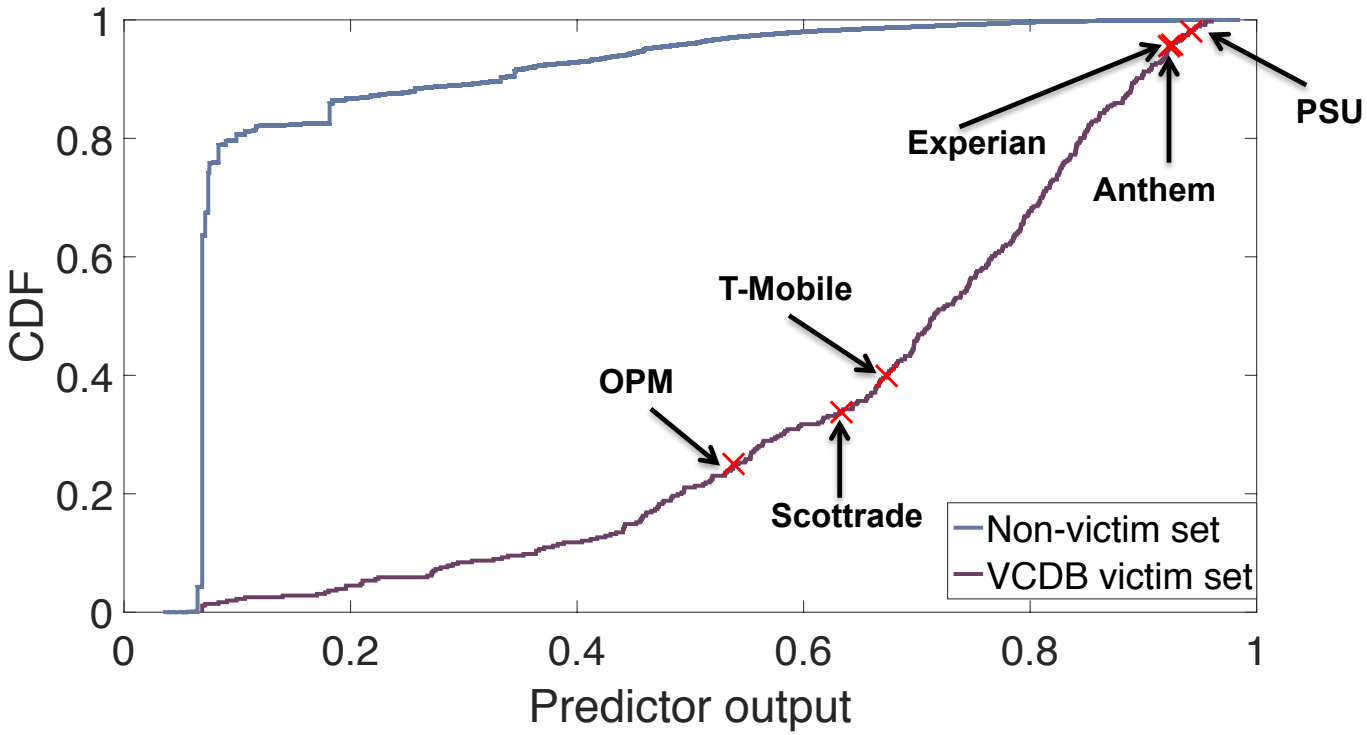

- Breaches from 2014: Sony, Ebay, Homedepot, Target, OnlineTech/JP Morgan Chase

Can we do even better?

Prediction by incident type

- Insufficient data for most of the incident types; one exception.

Can we do even better?

Prediction by incident type

- Insufficient data for most of the incident types; one exception.

Incident type	Crimeware	Cyber Esp.	Web app.	Else
VCDB	59	16	368	213

Can we do even better?

Prediction by incident type

- Insufficient data for most of the incident types; one exception.

Incident type	Crimeware	Cyber Esp.	Web app.	Else
VCDB	59	16	368	213

- Train a binary classifier: likelihood of falling victim to "web app incident".

Can we do even better?

Prediction by incident type

- Insufficient data for most of the incident types; one exception.

Incident type	Crimeware	Cyber Esp.	Web app.	Else
VCDB	59	16	368	213

- Train a binary classifier: likelihood of falling victim to "web app incident".

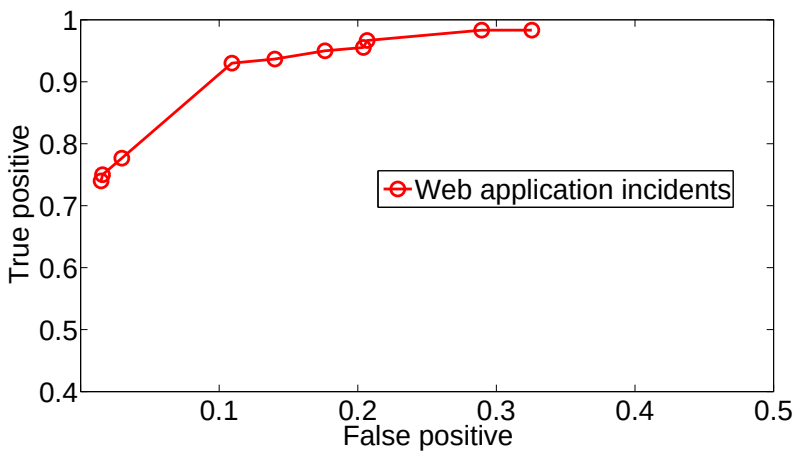

- Fine-grained predictions are possible esp. with more incident data.

Outline of the talk

- Data and Preliminaries
  - Data sources
  - Pre-processing
- Forecasting method and results
  - Feature extraction
  - Construction of the classifier
  - Prediction performance
- **Fine-grained prediction**
- Risk assessment as a form of "public monitoring"

But we don't have more incident data

Idea: conditional density estimation

- Use the preceding framework to perform "overall risk" prediction.
- Next, perform *conditional prediction*: if an incident should occur, the likelihood of its being of a particular type $\Rightarrow$ *Risk profiles*.

But we don't have more incident data

Idea: conditional density estimation

- Use the preceding framework to perform "overall risk" prediction.
- Next, perform *conditional prediction*: if an incident should occur, the likelihood of its being of a particular type $\Rightarrow$ *Risk profiles*.

Shall use VCDB (including non-cyber incidents)

- Details on the incident, actor, action, assets involved, and the victim.
- Plus information from AWIS: rank (global, regional), rank history (average, standard deviation), speed, age, locale, category, publicly traded, etc.

Challenges

Incomplete labels

- The level of details that are available vary for each report.

Challenges

Incomplete labels

- The level of details that are available vary for each report.

Selection bias and rare events

- Data incidents are largely under-reported.
- There is discrepancy in reporting.

Challenges

Incomplete labels

- The level of details that are available vary for each report.

Selection bias and rare events

- Data incidents are largely under-reported.
- There is discrepancy in reporting.

	Error	Hacking		Malware	Misuse	Physical		Social
		Comp. Cred.	Other			Theft	Other	
Overall	0.22	0.12	0.21	0.06	0.15	0.14	0.04	0.04
Manufacturing	0.08	0.09	0.33	0.13	0.22	0.13	0.00	0.02
Retail Trade	0.15	0.26	0.11	0.19	0.09	0.09	0.11	0.02
Information	0.09	0.28	0.41	0.07	0.04	0.03	0.01	0.07
Finance & Insurance	0.25	0.09	0.11	0.05	0.12	0.10	0.19	0.07
Pro., Sci. & Tech. Svcs	0.16	0.09	0.56	0.04	0.13	0.09	0.00	0.02
Educational Svcs	0.30	0.13	0.21	0.06	0.11	0.14	0.00	0.05
Health Care & Social Asst	0.25	0.08	0.03	0.02	0.23	0.38	0.02	0.01
Public Administration	0.27	0.09	0.29	0.03	0.17	0.10	0.01	0.03

Distribution of incidents by business sector.

A layered approach

To address incomplete labels:

- Train multiple binary classifiers, each estimating a portion of the risk
- Chain rule:
  $P(\text{Physical Theft}) = P(\text{Physical}) \times P(\text{Theft} \mid \text{Physical})$

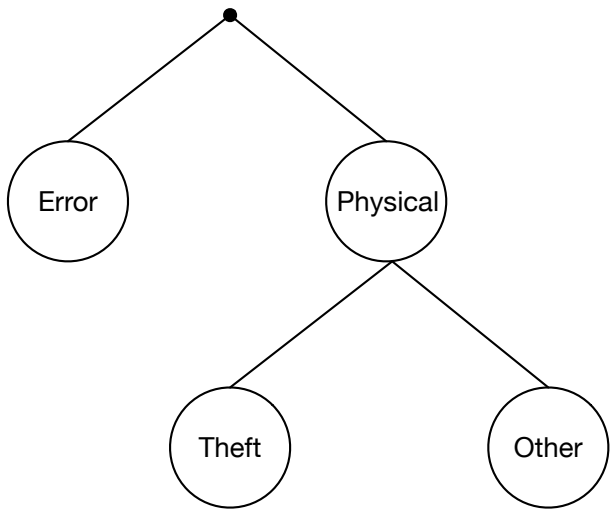

Our classifiers

- Training is done using Random Forest classifiers on 2013 incidents, and testing is performed on 2014 incidents.
- Two sets of classifiers using only business sector, and the full feature-set.

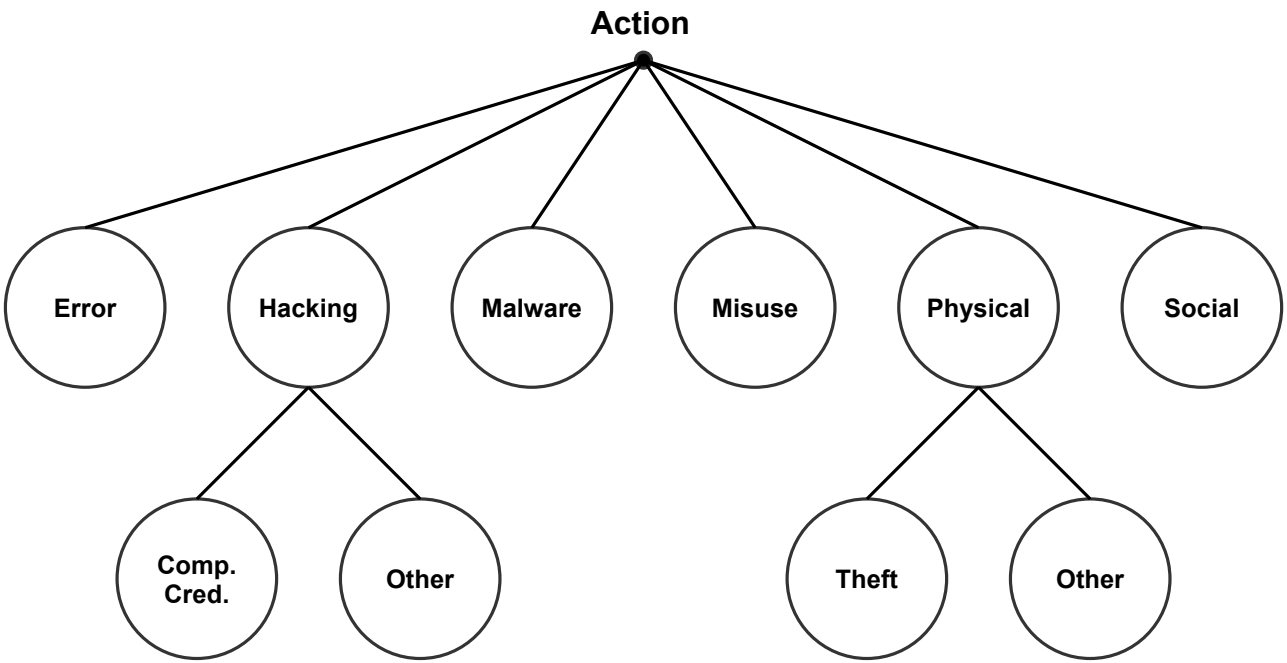

Our classifiers

- Training is done using Random Forest classifiers on 2013 incidents, and testing is performed on 2014 incidents.
- Two sets of classifiers using only business sector, and the full feature-set.

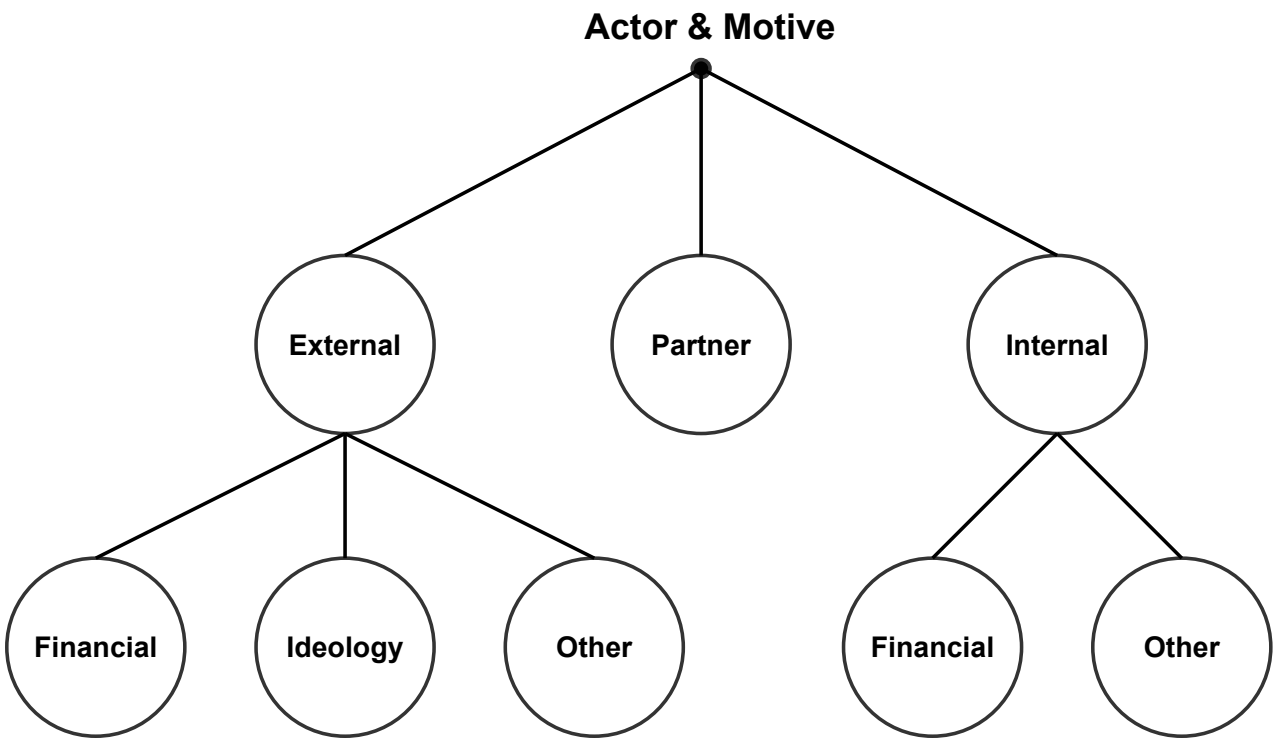

Our classifiers

- Training is done using Random Forest classifiers on 2013 incidents, and testing is performed on 2014 incidents.
- Two sets of classifiers using only business sector, and the full feature-set.

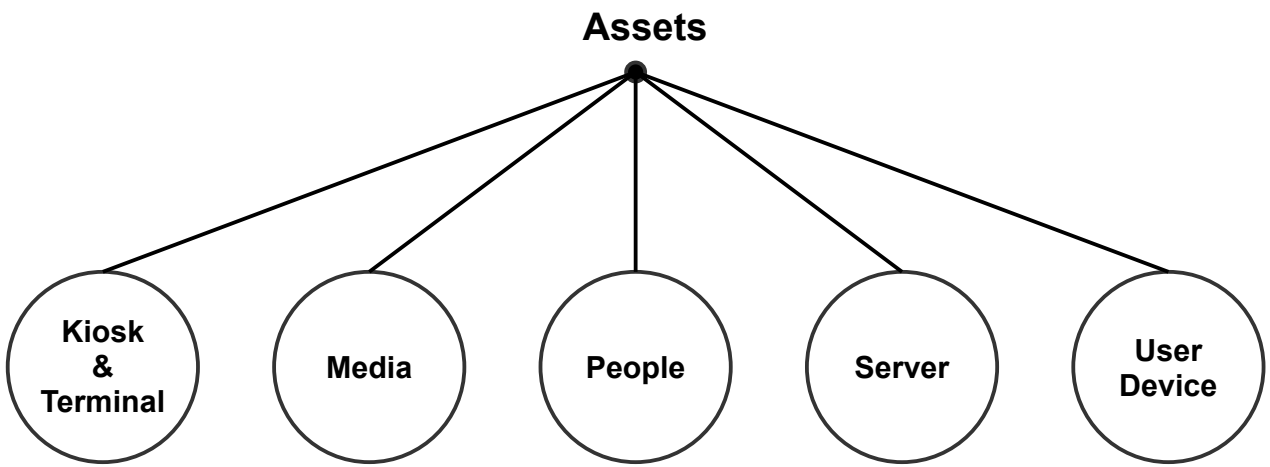

Interpreting the outputs

Converting conditional distribution to binary labels

- Identify incident types as *Risky*, or *Non-Risky*.
- Define thresholds for each classifier and apply them to the continuous output.
- Find a point on the RoC curve of the classifier that corresponds to the desired level of protection (true positive rate).

Example risk profiles

Risk profiles for sample organizations and their corresponding industries.

Organization	Error	Hacking		Malware	Misuse	Physical		Social
		Comp. Cred.	Other			Theft	Other	
Information								
Russian Radio			×					
Verizon			×					
Public Administration								
Macon Bibb County	×							
Internal Revenue Service					×			

- Gray cells signify incident types with high risk;
- Crosses indicate the actual incident.

Outline of the talk

- Data and Preliminaries
  - Data sources
  - Pre-processing
- Forecasting method and results
  - Feature extraction
  - Construction of the classifier
  - Prediction performance
- Fine-grained prediction
- **Risk assessment as a form of "public monitoring"**

Information sharing agreements among firms

Information sharing agreements among firms

Executive Order 13691 "Promoting Private Sector Cybersecurity Information Sharing"

Information Sharing and Analysis Organizations (ISAOs), Cyber Information Sharing and Collaboration Program (CISCP), Computer Emergency Readiness Team (US-CERT), etc

Information Sharing and Analysis Centers (ISACs)

FINANCIAL SERVICES Information Sharing and Analysis Center

IT ISAC

The disincentive: disclosure costs

Disclosure costs

- Drop in market values following security breach disclosure [Campbell et al. 03][Cavusoglu, Mishra, Raghunathan 04]
- Loss of consumer/partner confidence
- Bureaucratic burden

The disincentive: disclosure costs

Disclosure costs

- Drop in market values following security breach disclosure [Campbell et al. 03][Cavusoglu, Mishra, Raghunathan 04]
- Loss of consumer/partner confidence
- Bureaucratic burden

How to sustain cooperation?

- Audits and sanctions (e.g. by an authority or the government) [Laube and Bohme 15]
- Introducing additional economic incentives (e.g. taxes and rewards for members of ISACs) [Gordon, Loeb, Lucyshyn 03]

The disincentive: disclosure costs

Disclosure costs

- Drop in market values following security breach disclosure [Campbell et al. 03][Cavusoglu, Mishra, Raghunathan 04]
- Loss of consumer/partner confidence
- Bureaucratic burden

How to sustain cooperation?

- Audits and sanctions (e.g. by an authority or the government) [Laube and Bohme 15]
- Introducing additional economic incentives (e.g. taxes and rewards for members of ISACs) [Gordon, Loeb, Lucyshyn 03]
- **Inter-temporal incentives**: conditioning future cooperation on history of past interactions.

Private vs. public monitoring

- Inter-temporal incentives are based on the beliefs of participants about each others' disclosure decisions.
  - Firms, or any external observer, can only imperfectly assess the honesty and comprehensiveness of these reports.
  - Who should perform the monitoring?
- We will consider a repeated game framework.
  - Imperfect private monitoring.
  - Imperfect public monitoring.
- We illustrate the key role of a rating/assessment system in facilitating cooperation on information disclosure.

Information sharing games: stage game model

- Two firms
- $r_i \in \{0, 1\}$: (partially) concealing and (fully) disclosing
- Gain from other firm's disclosed information G
- Disclosure costs C

	1	0
1	$G - C$, $G - C$	$-C$, G
0	G, $-C$	0, 0

Information sharing games: stage game model

- Two firms
- $r_i \in \{0, 1\}$: (partially) concealing and (fully) disclosing
- Gain from other firm's disclosed information G
- Disclosure costs C

	1	0
1	$G - C$, $G - C$	$-C$, G
0	G, $-C$	0, 0

$\Rightarrow$ Prisoner's dilemma: only equilibrium of one shot game is $(0, 0)$.

Repeated games and monitoring possibilities

- Can we sustain (nearly) **efficient payoffs** in repeated games?
- Depends on whether/how deviations are detected and punished.
- Let b_i denote the **belief** of i about r_j.

Repeated games and monitoring possibilities

- Can we sustain (nearly) **efficient payoffs** in repeated games?
- Depends on whether/how deviations are detected and punished.
- Let b_i denote the **belief** of i about r_j.

Imperfect **Private** Monitoring

$$\pi(b_i|r_j) = \begin{cases} \epsilon, & \text{for } b_i = 0, r_j = 1 \\ 1-\epsilon, & \text{for } b_i = 1, r_j = 1 \\ \alpha, & \text{for } b_i = 0, r_j = 0 \\ 1-\alpha, & \text{for } b_i = 1, r_j = 0 \end{cases}$$

with $\epsilon \in (0, 1/2)$ and $\alpha \in (1/2, 1)$.

Imperfect **Public** Monitoring

$$\hat{\pi}((b_i, b_j)|(r_i, r_j)) := \pi(b_i|r_j)\pi(b_j|r_i)$$

monitoring by a rating/assessment system.

Limitations of private signals: a two-stage game

	1	0
1	$G-C,\ G-C$	$-C,\ G$
0	$G,\ -C$	$0,\ 0$

Table : Information sharing game

	H	L
H	h,h	0,0
L	0,0	ℓ,ℓ

Table : Partnership coordination

Based on the outcome of the 1^{st} stage, decide whether to form a high or low profit partnership in the 2^{nd} stage.

Limitations of private signals: a two-stage game

	1	0
1	$G-C,\ G-C$	$-C,\ G$
0	$G,\ -C$	$0,\ 0$

Table : Information sharing game

	H	L
H	h,h	0,0
L	0,0	ℓ,ℓ

Table : Partnership coordination

Pure strategies: play $r_i = 1$, then H iff $b_i = 1$ (trigger strategies).

- It is optimal for i to play H iff she believes w.p. $\geq \frac{\ell}{h+\ell}$ firm j also playing H.
- If i plays $r_i = 1$, she believes w.p. $1-\epsilon$ that j will play H. I.e., it is not sequentially rational for her to act based on her signal.

$\Rightarrow$ Following a similar argument for other pure strategies, cooperation cannot be guaranteed. Mixed strategies fare better: cooperation can happen with positive probability.

Infinitely repeated games with private monitoring

- Wanted: **a folk theorem** - a full characterization of payoffs that can be achieved in a repeated game if players are sufficiently patient.

Infinitely repeated games with private monitoring

- Wanted: **a folk theorem** - a full characterization of payoffs that can be achieved in a repeated game if players are sufficiently patient.
- **No folk theorem** for infinitely repeated games with imperfect private monitoring in general.

Infinitely repeated games with private monitoring

- Wanted: **a folk theorem** - a full characterization of payoffs that can be achieved in a repeated game if players are sufficiently patient.
- **No folk theorem** for infinitely repeated games with imperfect private monitoring in general.
  - They exist for some **modifications/subclasses**:
    - Communication (cheap talk) [Compte 98, Kandori and Matsushima 98].
    - Pubic actions, e.g., announcing sanctions [Park 11].
    - Sufficiently correlated private signals [Mailath and Morris 02].

Imperfect public monitoring: A folk theorem

[Fudenberg, Levine, and Maskin 1994]

If the imperfect public monitoring is *sufficiently informative*, s.t.:

- individual full rank: deviations by an individual player are statistically distinguishable.
- pairwise full rank: deviations by players i and j are distinct, i.e., induce different distributions over public outcomes.

Imperfect public monitoring: A folk theorem

[Fudenberg, Levine, and Maskin 1994]

If the imperfect public monitoring is *sufficiently informative*, s.t.:

- individual full rank: deviations by an individual player are statistically distinguishable.
- pairwise full rank: deviations by players i and j are distinct, i.e., induce different distributions over public outcomes.

then there exists a discount factor $\underline{\delta} < 1$, such that for all $\delta \in (\underline{\delta}, 1)$, any feasible and strictly individually rational payoff profile can be sustained by public strategies.

Our monitoring mechanism is informative

- It can be verified that our public monitoring model satisfies these two conditions.
- E.g., minmax profile has individual full rank for either firm:

$$\begin{array}{c} \mathbf{b} = \\ r_i = 0 \\ r_i = 1 \end{array} \begin{array}{c} \begin{array}{cccc} (0,0) & (1,0) & (0,1) & (1,1) \end{array} \\ \begin{pmatrix} \alpha^2 & (1-\alpha)\alpha & \alpha(1-\alpha) & (1-\alpha)^2 \\ \epsilon\alpha & (1-\epsilon)\alpha & \epsilon(1-\alpha) & (1-\epsilon)(1-\alpha) \end{pmatrix} \end{array}$$

The role of monitoring in information sharing

- The folk theorem holds with the **same monitoring technology** of that of individual firms $\Rightarrow$ the rating/assessment system facilitates coordination.
- Conclusions hold with countably finite disclosure decisions and discrete ratings by the monitoring system.
- Work remains:
  - The structure of efficient public strategies.
  - Assessment (e.g., risk predictions) that affect payoffs outside of the information sharing agreement.

Conclusion

A prediction framework for forecasting cybersecurity incidents

- Data sources, pre-processing, features, and training.
- Fine-grained prediction of incident types.

Its role in encouraging better information sharing

- As a form of public monitoring to induce inter-temporal incentives to sustain cooperation.

Conclusion

A prediction framework for forecasting cybersecurity incidents

- Data sources, pre-processing, features, and training.
- Fine-grained prediction of incident types.

Its role in encouraging better information sharing

- As a form of public monitoring to induce inter-temporal incentives to sustain cooperation.

An interesting coupling

- One's performance is only as good as one's data
  - Incidents: under-reporting, non-uniform reporting and bias.
  - Other errors/noises in the data pale in comparison.
- But even imperfect monitoring can be used to induce security information sharing.
  - Which leads to better quality data (esp. labels), which in turn improves the quality of monitoring.

Acknowledgement

Work supported by the NSF and the DHS

References:

- Y. Liu, A. Sarabi, J. Zhang, P. Naghizadeh, M. Karir, M. Bailey and M. Liu, "Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents", *USENIX Security*, August 2015, Washington, D. C.
- A. Sarabi, P. Naghizadeh, Y. Liu and M. Liu, "Prioritizing Security Spending: A Quantitative Analysis of Risk Distributions for Different Business Profiles", *WEIS*, June 2015, Delft University, The Netherlands.
- P. Naghizadeh and M. Liu, "Inter-Temporal Incentives in Security Information Sharing Agreements", *ITA*, February 2016, San Diego, CA.