Navigating Internet Neighborhoods: Reputation, How to Crowd-source It, and Its Impact on Security

Mingyan Liu

Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI

April 25, 2013

Acknowledgment

Collaborators:

- Parinaz Naghizadeh Ardabili
- Jing Zhang, Michael Bailey, Manish Karir

Threats to Internet security and availability

From unintentional to intentional, random maliciousness to economic driven:

- misconfiguration
- mismanagement
- botnets, worms, SPAM, DoS attacks, . . .

Typical operators' countermeasures: *filtering/blocking*

- within specific network services (e.g., e-mail)
- with the domain name system (DNS)
- based on source and destination (e.g., firewalls)
- within the control plane (e.g., through routing policies)

Host Reputation Block Lists (RBLs)

Commonly used RBLs:

- daily average volume (unique entries) ranging from 146M (BRBL) to 2K (PhishTank)

RBL Type	RBL Name
Spam	BRBL, CBL, SpamCop, WPBL, UCEPROTECT
Phishing/Malware	SURBL, PhishTank, hpHosts
Active attack	Darknet scanners list, Dshield

Potential impact of RBLs

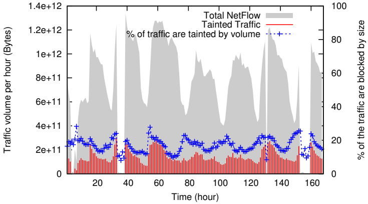

(a) By traffic volume (bytes).

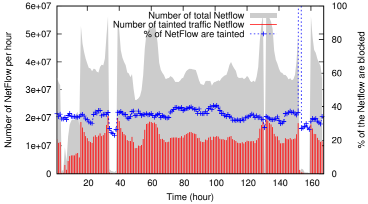

(b) By number of flows.

NetFlow records of all traffic flows at Merit Network

- at all peering edges of the network from 6/20/2012-6/26/2012
- sampling ratio 1:1
- 118.4TB traffic: 5.7B flows, 175B packets.

As much as 17% (30%) of overall traffic (flows) “tainted”

How reputation lists should be/are used

Strengthen defense:

- filter configuration, blocking mechanisms, etc.

Strengthen security posture:

- get hosts off the list
- install security patches, update software, etc.

Retaliation for being listed:

- lost revenue for spammers
- example: recent DDoS attacks against Spamhaus by Cyberbunker

Aggressive outbound filtering:

- fixing the symptom rather than the cause
- example: the country of Mexico

Limitations of host reputation lists

Host identities can be highly transient:

- dynamic IP address assignment
- policies inevitably reactive, leading to significant false positives and misses
- potential scalability issues

RBLs are application specific:

- a host listed for spamming can initiate a different attack

Lack of standard and transparency in how they are generated

- not publicly available: subscription based, query enabled

An alternative: network reputation

Define the notion of "reputation" for a network (suitably defined) rather than for hosts

A network is typically governed by consistent policies

- changes in system administration on a much larger time scale
- changes in resource and expertise on a larger time scale

Policies based on network reputation is *proactive*

- reputation reflects the security posture of the entire network, across all applications, slow changing over time

Enables risk-analytical approaches to security; tradeoff between benefits in and risks from communication

An illustration

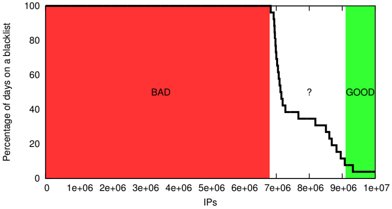

Figure: Temporal and Spatial aggregation of reputation

- Taking the union of 9 RBLs
- Left: % of days an Addr is listed (est. total of 100M)
- Right: % Addrs blacklisted within an autonomous system (est. total of 35-40K)

Many challenges to address

- What is the appropriate level of aggregation
- How to obtain such aggregated reputation measure, over time, space, and applications
- How to use these to design reputation-aware policies

- How to make the reputation measure accurate representation of the quality of a network
- What effect does it have on the network's behavior toward others and itself

Outline of the talk

Incentivizing input – crowd-sourcing reputation

- Assume a certain level of aggregation
- Each network possesses information about itself and others
- Can we incentivize networks to participate in a collective effort to achieve accurate estimates/reputation assessment, while observing privacy and self interest

Impact of reputation on network behavior

- Benefits from and cost in investing in security
- Positive externality and free rider
- Can the desire for good reputation (or the worry over bad reputation) positively alter a network's decision in investment

Crowd-sourcing reputation

- Basic setting
  - A distributed multi-agent system.
  - Each agent has perceptions or beliefs about other agents.
  - The truth about each agent known only to itself.
  - Each agent wishes to obtain the truth about others.
- Goal: construct mechanisms that *incentivize* agents to participate in a collective effort to arrive at correct perceptions.
- Key design challenges:
  - Participation is voluntary.
  - Individuals may not report truthfully even if they participate.
  - Individuals may collude.

Other applicable contexts and related work

Online review/recommendation systems:

- Example: Amazon, EBay
- Users (e.g., sellers and buyers) rate each other

Reputation in P2P systems

- Sustaining cooperative behavior among self-interested individuals.
- User participation is a given; usually perfect observation.

Elicitation and prediction mechanisms

- Used to quantify the performance of forecasters; rely on observable objective ground truth.
- Users do not attach value to realization of event or the outcome built by elicitor.

The Model

- K inter-connected networks, $N_1, N_2, \cdots, N_K$.
- Network N_i's overall quality or health condition described by a $r_{ii} \in [0, 1]$: *true* or *real quality* of N_i.
- A central *reputation system* collects input from each N_i and computes a *reputation index* $\hat{r}_i$, the estimated quality.

Main Assumptions

- N_i knows r_{ii} precisely, but this is its *private information*.
- N_i can sufficiently monitor inbound traffic from N_j to form an estimate R_{ij} of r_{jj}.
- N_i's observation is in general *incomplete* and may contain noise/errors: $R_{ij} \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2)$.
- This distribution is known to network N_j, while N_i itself may or may not be aware of it.
- The reputation system may have independent observations R_{0i} for $\forall i$.
- The *reputation mechanism* is common knowledge.

Designing the mechanism

- Goal: solution to the *centralized* problem in an *informationally decentralized* system.
- Choice parameters of the mechanism are:
  - Message space $\mathcal{M}$: inputs requested from agents.
  - Outcome function $h(\cdot)$: a rule according to which the input messages are mapped to outcomes.
- Other desirable features: budget balance, and individual rationality.

The centralized problem

Systems' Objective

Minimize estimation error for all networks.

Two possible ways of defining a reputation index:

- *Absolute index* $\hat{r}_i^A$: an estimate of r_{ii}.
- *Relative index* $\hat{r}_i^R$: given true qualities r_{ii}, $\hat{r}_i^R = \frac{r_{ii}}{\sum_k r_{kk}}$.

$$\min \sum_i |\hat{r}_i^A - r_{ii}| \quad \text{or} \quad \min \sum_i |\hat{r}_i^R - \frac{r_{ii}}{\sum_k r_{kk}}|$$

If the system had full information about all parameters:

$$\hat{r}_i^A = r_{ii} \quad \text{and} \quad \hat{r}_i^R = \frac{r_{ii}}{\sum_k r_{kk}}$$

In a decentralized system

N_i's Objective

The truth element: security

Accurate estimate $\hat{r}_j$ on networks N_j *other than itself*.

$$I_i = -\sum_{j \neq i} f_i(|\hat{r}_j^A - r_{jj}|) \quad \text{or} \quad I_i = -\sum_{j \neq i} f_i(|\hat{r}_j^R - \frac{r_{jj}}{\sum_k r_{kk}}|) \ .$$

$f_i()$'s are increasing and convex.

The image element: reachability

High reputation $\hat{r}_i$ for *itself*.

$$II_i = g_i(\hat{r}_i^A) \quad \text{or} \quad II_i = g_i(\hat{r}_i^R).$$

$g_i()$'s are increasing and concave.

Different types of networks

- *Truth type:* dominated by security concerns, e.g., DoD networks, a buyer on Amazon.
- *Image type:* dominated by reachability/traffic attraction concerns: a blog hosting site, a phishing site, a seller on Amazon.
- *Mixed type:* legitimate, non-malicious network; preference in general increasing in the accuracy of others' and its own quality estimates.

$$u_i = -\lambda \sum_{j \neq i} f_i(|\hat{r}_j^A - r_{jj}|) + (1-\lambda) g_i(\hat{r}_i^A)$$

- A *homogeneous* vs. a *heterogeneous* environment

Reputation mechanisms

Design a simple mechanism for each type of environment and investigate its incentive feature.

- Possible forms of input:
  - *cross-reports* $X_{ij}, j \neq i$: N_i's assessment of N_j's quality
  - *self-reports* X_{ii}: networks' *self-advertised* quality measure
- The qualitative features (increasing in truth and increasing in image) of the preference are public knowledge; the functions $f_i()$, $g_i()$ are private information.
- N_i is an expected utility maximizer due to incomplete information.
- Assume external observations are unbiased.
- If taxation is needed, aggregate utility of N_i defined as $v_i := u_i - t_i$.

Setting I: Truth types, absolute reputation

$$\text{(Model I)} \qquad u_i = -\sum_{j \neq i} f_i(|\hat{r}_j^A - r_{jj}|)$$

The absolute scoring (AS) mechanism:

- Message space $\mathcal{M}$: each user reports $x_{ii} \in [0,1]$.
- Outcome function $h(\cdot)$:
  - The reputation system chooses $\hat{r}_i^A = x_{ii}$.
  - N_i is charged a tax term t_i given by:

$$t_i = |x_{ii} - R_{0i}|^2 - \frac{1}{K-1} \sum_{j \neq i} |x_{jj} - R_{0j}|^2 \ .$$

Properties of the AS mechanism

Rationale: assign reputation indices assuming truthful reports, ensure truthful reports by choosing the appropriate t_i.

- Truth-telling is a *dominant strategy* in the induced game
  $\Rightarrow$ Achieves centralized solution.
- $\sum_i t_i = 0$
  $\Rightarrow$ Budget balanced.
- The mechanism is individually rational
  $\Rightarrow$ Voluntary participation.

Truth revelation under AS

Truth-telling is a dominant strategy in the game induced by the AS mechanism

$$
\begin{aligned}
E[v_i(x_{ii}, \{X_{jj}\}_{j \neq i})] &= -\sum_{j \neq i} E[f_i(|\hat{r}_j^A - r_{jj}|)] \\
&\quad - E[|x_{ii} - R_{01}|^2] + \frac{1}{K-1} \sum_{j \neq i} E[|X_{jj} - R_{0j}|^2]
\end{aligned}
$$

- x_{ii} can only adjust the 2nd term, thus chosen to minimize the 2nd term.
- By assumption, N_i knows $R_{0i} \sim \mathcal{N}(r_{ii}, \sigma_{0i}^2)$, thus optimal choice $x_{ii} = r_{ii}$.

Individual rationality under AS

The AS mechanism is individually rational.

- Staying out: reserved utility given by $-\sum_{j \neq i} E(f_i(|R_{ij} - r_{jj}|))$.
- Participating: expected utility $-\sum_{j \neq i} f_i(0)$ at equilibrium.
- $f_i(\cdot)$ is increasing and convex, thus $E[f_i(|R_{ij} - r_{jj}|)] \geq f_i(E(|R_{ij} - r_{jj}|)) = f_i(\sqrt{\frac{2}{\pi}}\sigma_{ij}) > f_i(0), \ \forall j \neq i$.
- The AS mechanism is individually rational.

Extended-AS Mechanism

- What if the system does not possess independent observations?
- Use a random ring to gather cross-observations and assess taxes.
- N_i is asked to report X_{ii}, as well as $X_{i(i+1)}$ and $X_{i(i+2)}$.

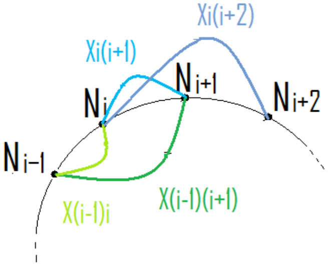

Extended-AS Mechanism

- N_i is charged two taxes:
  - on the inaccuracy of its self-report wrt what N_{i-1} says about N_i
  - on the inaccuracy of its cross-report on N_{i+1} wrt what N_{i-1} says

$$
\begin{aligned}
t_i \quad = \quad & |x_{ii} - X_{(i-1)i}|^2 - \frac{1}{K-2} \sum_{j \neq i, i+1} |X_{jj} - X_{(j-1)j}|^2 \\
& + |x_{i(i+1)} - X_{(i-1)(i+1)}|^2 - \frac{1}{K-2} \sum_{j \neq i, i+1} |X_{j(j+1)} - X_{(j-1)(j+1)}|^2
\end{aligned}
$$

- Truthful self-reports achieved by the 1st taxation term.
- Truthful cross-reports achieved by the 2nd taxation term.
- Other associations also possible: e.g., random sets.

Extended-AS results in the centralized solution

Setting II: Truth types, relative reputation

$$\text{(Model II)} \qquad u_i = -\sum_{j \neq i} f_i(|\hat{r}_j^R - \frac{r_{jj}}{\sum_k r_{kk}}|)$$

The fair ranking (FR) mechanism:

- Message space $\mathcal{M}$: each user reports $x_{ii} \in [0, 1]$.
- Outcome function $h(\cdot)$:
  - the system assigns $\hat{r}_i^R = \frac{x_{ii}}{\sum_k x_{kk}}$.
  - No taxation is used.

Properties of the FR mechanism

- Truth-telling is a Bayesian Nash equilibrium in the induced game

$$u_i(x_{ii}, \{r_{kk}\}_{k\neq i}) = -\sum_{j\neq i} f_i(|\frac{r_{jj}(x_{ii}-r_{ii})}{(x_{ii}+\sum_{k\neq i} r_{kk})(\sum_k r_{kk})}|)$$

  $\Rightarrow$ Achieves centralized solution $x_{ii} = r_{ii}$.
- The mechanism is individually rational
  $\Rightarrow$ Voluntary participation.
- Achievable without cross-observations from other networks, direct observations by the system, or taxation.

Setting III: Mixed types, relative reputation

$$(\text{Model III}) \quad u_i = \quad -\sum_{j \neq i} f_i(|\hat{r}_j^R - \frac{r_{jj}}{\sum_k r_{kk}}|) + g_i(\hat{r}_i^R)$$

- The individual's objective is no longer aligned with the system objective
- Direct mechanism possible depending on the specific forms of $f_i()$ and $g_i()$.

Setting IV: Mixed types, absolute reputation

$$\text{(Model IV)} \quad u_i = \quad -\sum_{j \neq i} f_i(|\hat{r}_j^A - r_{jj}|) + g_i(\hat{r}_i^A)$$

An Impossibility result:

- centralized solution cannot be implemented in BNE.

Consider suboptimal solution:

- use both self- and cross-reports
- forgo the use of taxation

A simple averaging mechanism

$$\text{(Model IV)} \quad u_i = \quad -\sum_{j \neq i} f_i(|\hat{r}_j^A - r_{jj}|) + g_i(\hat{r}_i^A)$$

- Solicit only cross-reports.
- Take $\hat{r}_i^A$ to be the average of all x_{ji}, $j \neq i$, and R_{0i}.
- Used in many existing online system: Amazon and Epinions.
- Truthful revelation of R_{ji} is a BNE.
  - N_j has no influence on its own estimate $\hat{r}_j^A$.
  - N_j's effective objective is to minimize the first term.
  - The simple averaging mechanism results in $\hat{r}_i^A \sim \mathcal{N}(r_{ii}, \sigma^2/K)$.
- $\hat{r}_i^A$ can be made arbitrarily close to r_{ii} as K increases.
- (Under this mechanism, if asked, N_i will always report $x_{ii} = 1$)

Can we do better?

Instead of ignoring N_i's self-report, incentivize N_i to provide *useful* information.

- Convince N_i that it can contribute to a higher estimated $\hat{r}_i^A$ by supplying input X_{ii},
- Use cross-reports to assess N_i's self-report, and threaten with punishment if it is judged to be overly misleading.

Truthful cross-reports

A mechanism in which N_i's cross-reports are not used in calculating its own reputation estimate. Then:

- N_i can only increase its utility by altering $\hat{r}_j^A$ when submitting X_{ij},
- N_i doesn't know r_{jj}, can't use a specific utility function to strategically choose X_{ij},
- N_i's best estimate of r_{jj} is R_{ij},

$\Rightarrow$ Truthful cross-reports!

Questions:

- Can N_i make itself look better by degrading N_j?
- Is it in N_i's interest to degrade N_j?

A punish-reward (PR) mechanism

Denote the output of the simple averaging mechanism by $\bar{X}_{0i}$.

$$\hat{r}_i^A(X_{ii}, \bar{X}_{0i}) = \begin{cases} \frac{\bar{X}_{0i} + X_{ii}}{2} & \text{if } X_{ii} \in [\bar{X}_{0i} - \epsilon, \bar{X}_{0i} + \epsilon] \\ \bar{X}_{0i} - |X_{ii} - \bar{X}_{0i}| & \text{if } X_{ii} \notin [\bar{X}_{0i} - \epsilon, \bar{X}_{0i} + \epsilon] \end{cases}$$

- ϵ is a fixed and known constant.
- Take the average of X_{ii} and $\bar{X}_{0i}$ if the two are sufficiently close; else punish N_i for reporting significantly differently.

$\Rightarrow$ Each network only gets to optimize its self-report, knowing all cross-reports are truthful.

Choice of self-report

Self-report x_{ii} determined by $\max_{x_{ii}} \; E[\hat{r}_i^A(x_{ii}, \bar{X}_{0i})]$, where $\bar{X}_{0i} \sim \mathcal{N}(r_{ii}, \frac{\sigma^2}{K})$ assuming common and known σ. Optimal x_{ii}, when $\epsilon = a\sigma' = a\frac{\sigma^2}{K}$, is given by:

$$x_{ii}^* = r_{ii} + a\sigma' y$$

$0 < y < 1 \quad \Rightarrow$
self-report is positively biased and within expected acceptable range.

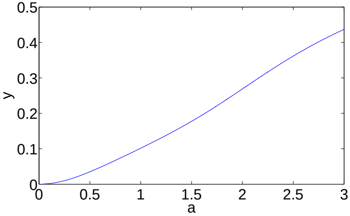

Performance of the mechanism

How close is $\hat{r}_i^A$ to the real quality r_{ii}:
$e_m := E(|\hat{r}_i^A - r_{ii}|)$

- For a large range of a values, N_i's self-report benefits the system as well as all networks other than N_i.
- Optimal choice of a does not depend on r_{ii} and σ'.

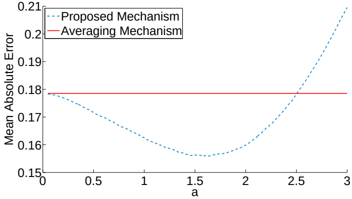

Figure: MAE for $r_{ii} = 0.75$, $\sigma^2 = 0.1$

Comparing the final reputation

Under the averaging mechanism $E[\bar{X}_{0i}] = r_{ii}$.

- For sufficiently large a, N_i benefits from providing the self-report.

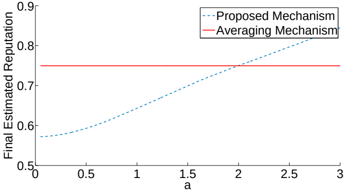

Figure: Expected reputation for $r_{ii} = 0.75$, $\sigma^2 = 0.1$

There is a mutually beneficial region $a \in [2, 2.5]$: the self-report helps N_i obtain a higher estimated reputation, while helping the system reduce its estimation error on N_i.

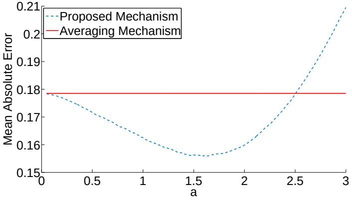

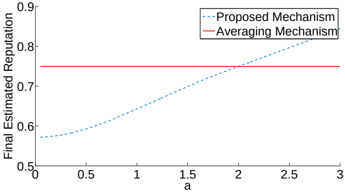

A heterogenous environment

A mix of T truth types and $K - T$ image types, using the AS mechanism

- Additional conditions needed to ensure individual rationality
  - The higher the percentage of image types, the less likely is a truth type to participate
  - The higher a truth type's own accuracy, the less interested it is to participate
  - An image type may participate if r_{ii} is small.
- The benefit of the mechanism decreases in the fraction of image types.

Handling collusion/cliques

- Absolute Scoring and Fair Ranking are naturally collusion-proof.
- PR remains functional using only the cross-observations from a subset of trusted entities, or even a single observation by the reputation system.

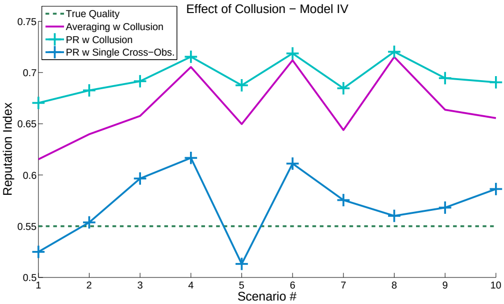

Handling collusion/cliques

- If the system lacks independent observations, introducing randomness can reduce the impact of cliques.
- E.g. extended-AS mechanism: tax determined by random matching with peers.
- Increased likelihood of being matched with non-colluding users reduces benefit of cliques.

Other aspects

- Other mechanisms, e.g., weighted mean of the cross-report, etc.
- Other heterogeneous environments
- Presence of malicious networks.

Outline of the talk

Incentivizing input – crowd sourcing reputation

- Assume a certain level of aggregation
- Each network possesses information about itself and others
- Can we incentivize networks to participate in a collective effort to achieve accurate estimates/reputation assessment, while observing privacy and self interest

Impact of reputation on network behavior

- Benefits from and cost in investing in security
- Positive externality and free rider
- Can the desire for good reputation (or the worry over bad reputation) positively alter a network's decision in investment

Interdependent Security Risks

- Security investments of a network have *positive externalities* on other networks.
- Networks' preferences are in general heterogeneous:
  - Heterogeneous costs.
  - Different valuations of security risks.
- Heterogeneity leads to under-investment and free-riding.

Network Security Investment Game

Originally proposed by [Jiang, Anantharam & Walrand, 2011]

- A set of N networks.
- N_i's action: invest $x_i \geq 0$ in security, with increasing effectiveness.
- Cost $c_i > 0$ per unit of investment (heterogeneous).
- $f_i(\mathbf{x})$ *security risk/cost* of N_i where:
    - $\mathbf{x}$ vector of investments of all users.
    - $f_i(\cdot)$ decreasing in each x_i and convex.
- N_i chooses x_i to minimize the cost function

$$h_i(x) := f_i(\mathbf{x}) + c_i x_i \ .$$

- Analyzed the suboptimality of this game.

Example: a total effort model

A 2-player total effort model: $f_1(\mathbf{x}) = f_2(\mathbf{x}) = f(x_1 + x_2)$, with $c_1 = c_2 = 1$.

$h_1(\mathbf{x}) = -f_1(x_1 + x_2) - x_1$, $h_2(\mathbf{x}) = -f_2(x_1 + x_2) - x_2$:

- Let $\mathbf{x}^o$ be the Nash Equilibrium, and $\mathbf{x}^*$ be the Social Optimum.
- At NE: $\partial h_i / \partial x_i = f'(x_1^o + x_2^o) + 1 = 0$.
- At SO: $\partial (h_1 + h_2) / \partial x_i = 2f'(x_1^* + x_2^*) + 1 = 0$.
    - By convexity of $f(\cdot)$, $x_1^o + x_2^o \leq x_1^* + x_2^* \Rightarrow$ under-investment.

An illustration

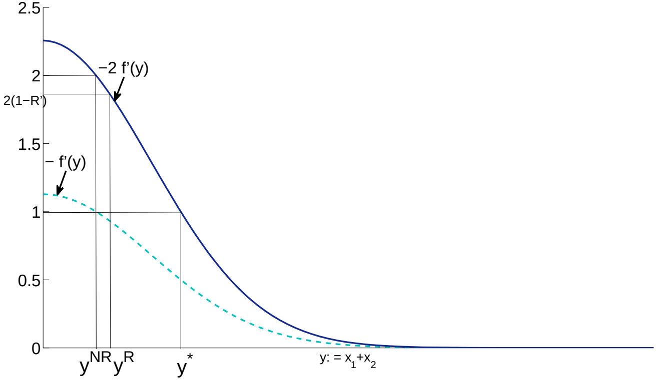

Figure: Suboptimality gap

The same game with reputation

The same model, with the addition:

- N_i will be assigned a reputation based on its investment.
- Valuation of reputation given by $R_i(\mathbf{x})$: increasing and concave.
- N_i chooses x_i to minimize the cost function

$$h_i(x) := f_i(\mathbf{x}) + c_i x_i - R_i(\mathbf{x}) \ .$$

The effect of reputation: example 1

One's reputation only depends on one's own investment:
$R_i(\mathbf{x}) = R_i(x_i)$

- $R_1(x) = kR_2(x)$, $k > 1$: N_1 values reputation more than N_2.
- $h_1(\mathbf{x}) = -f_1(x_1 + x_2) - x_1 - R_1(x_1)$,
  $h_2(\mathbf{x}) = -f_1(x_1 + x_2) - x_2 - R_2(x_2)$.
- At NE: $\partial h_i / \partial x_i = f'(x_1^R + x_2^R) + 1 - R_i'(x_i^R) = 0$.
  - $R_1'(x_1^R) = R_2'(x_2^R)$ and thus $x_1^R > x_2^R \Rightarrow$ The one who values reputation more, invests more.
  - By convexity of $f(\cdot)$, $x_1^o + x_2^o \leq x_1^R + x_2^R \Rightarrow$ Collectively invest more in security and decrease suboptimality gap.

An illustration

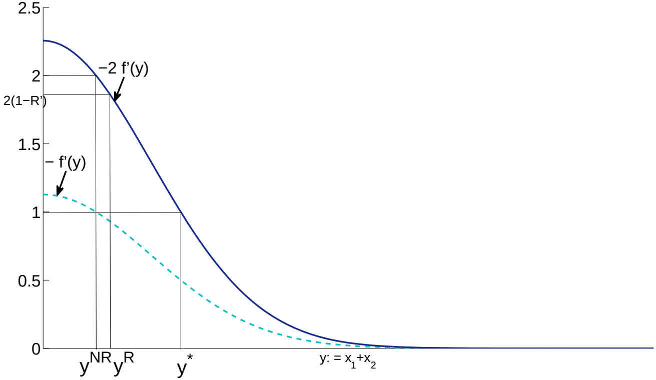

Figure: Driving equilibrium investments towards the social optimum

The effect of reputation: example 2

One's reputation depends on the total investment: $R_i(x_1 + x_2)$:

- Assume $R_1(x_1 + x_2) = kR_2(x_1 + x_2)$, $k > 1$: N_1 values reputation more than N_2.
- At NE: $\partial h_i / \partial x_i = f'(x_1^R + x_2^R) + 1 - R_i'(x_1^R + x_2^R) = 0$.
  - $x_2^R = 0 \Rightarrow$ The one who values reputation more, makes all the investment.
  - N_1 invests at the same level as they collectively did when using $R_i(x_i)$, decreasing suboptimality gap.
- Similar to the standard total effort model: agents free-ride.

Currently under investigation

- More interesting cases lie in between: $R_i(\hat{r}_i(\mathbf{x}))$ strongly depends on x_i but weakly depends on x_j.
- Mechanism design approach.
- In practice it may be hard for a network to know the security risk it is subject to; reputation estimate can be used as a proxy: $f_i(\hat{r}_1, \cdots, \hat{r}_N)$.
- Augmenting the action space: e.g., different levels of efforts in *inbound* and *outbound* traffic filtering, which may be a cheaper and more effective means of improving one's reputation but not necessarily one's true security quality.

Conclusion

Network reputation as a way to capture, encourage, and inform the security quality of policies

Incentivizing input – crowd sourcing reputation

- A number of preference models and environments
- Incentive mechanisms in each case

Impact of reputation on network behavior

- A reputation-augmented security investment game.
- Reputation can increase the level of investment and drive the system closer to social optimum.
- Many interesting open questions.