Network Economics
--
Lecture 3: Incentives in online systems II: robust reputation systems and information elicitation

Patrick Loiseau
EURECOM
Fall 2016

References

- Main:
  - N. Nisam, T. Roughgarden, E. Tardos and V. Vazirani (Eds). "Algorithmic Game Theory", CUP 2007. Chapters 27.
    - Available online: http://www.cambridge.org/journals/nisan/downloads/Nisan_Non-printable.pdf
- Additional:
  - Yiling Chen and Arpita Gosh, "Social Computing and User Generated Content," EC'13 tutorial
    - Slides at http://www.arpitaghosh.com/papers/ec13_tutorialSCUGC.pdf and http://yiling.seas.harvard.edu/wp-content/uploads/SCUGC_tutorial_2013_Chen.pdf
  - M. Chiang. "Networked Life, 20 Questions and Answers", CUP 2012. Chapters 3-5.
    - See the videos on www.coursera.org

Outline

1. Introduction
2. Eliciting effort and honest feedback
3. Reputation based on transitive trust

Outline

1. Introduction
2. Eliciting effort and honest feedback
3. Reputation based on transitive trust

Importance of reputation systems

- Internet enables interactions between entities
- Benefit depends on the entities ability and reliability
- Revealing history of previous interaction:
  - Informs on abilities
  - Deter moral hazard
- Reputation: numerical summary of previous interactions records
  - Across users – can be weighted by reputation (transitivity of trust)
  - Across time

Reputation systems operation

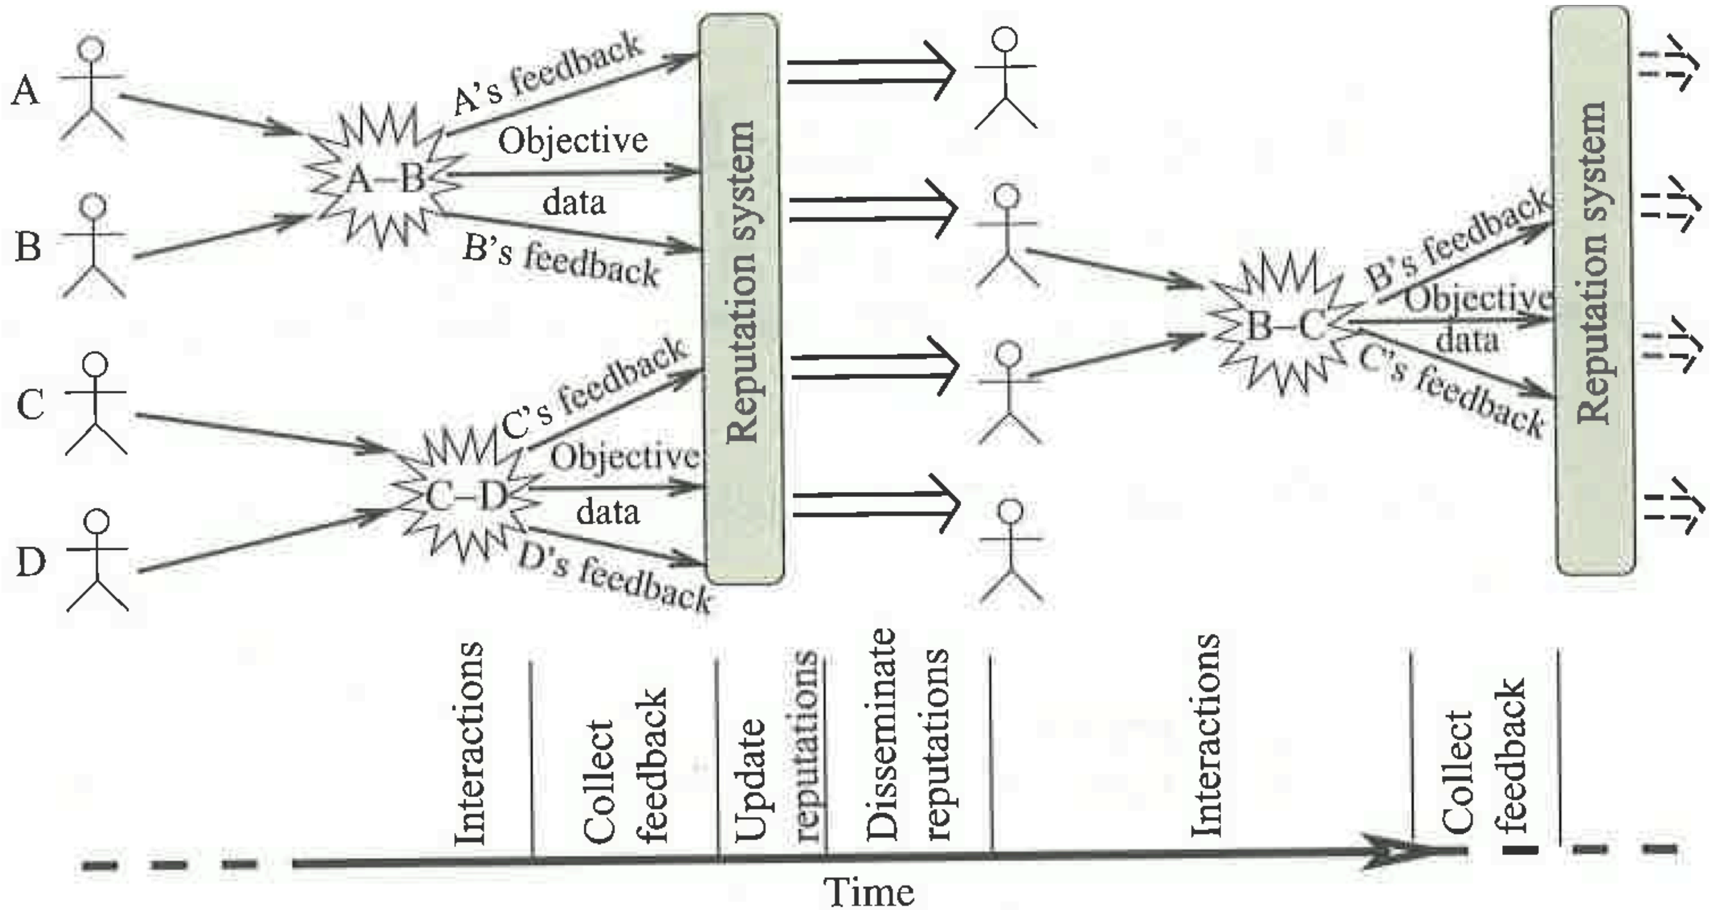

Attacks on reputation systems

- Whitewashing
- Incorrect feedback
- Sybil attack

A simplistic model

- Prisoner's dilemma again!
- One shot
  - (D, D) dominant
- Infinitely repeated
  - Discount factor δ

	C	D
C	1, 1	-1, 2
D	2, -1	0, 0

Equilibrium with 2 players

- Grim = Cooperate unless the other player defected in the previous round
- (Grim, Grim) is a subgame perfect Nash equilibrium if $\delta \geq 1/2$
  - We only need to consider single deviations

- → If users do not value future enough, they don't cooperate

Game with N+1 Players (N odd)

- Each round: players paired randomly
- With reputation (reputation-grim): agents begin with good reputation and keep it as long as they play C against players with good reputation and D against those with bad ones
  - SPNE if $\delta \geq 1/2$
- Without reputation (personalized-grim): keep track of previous interaction with same agent
  - SPNE if $\delta \geq 1-1/(2N)$

Whitewashing

- Play D and come back as new user!
- Possible to avoid this with entry fee f

Outline

1. Introduction
2. Eliciting effort and honest feedback
3. Reputation based on transitive trust

Different settings

- How to enforce honest reporting of interaction experience?

1. Objective information publicly revealed: can just compare report to real outcome
   – E.g., weather prediction

2. No objective outcome is available
   – E.g., product quality – not objective
   – E.g., product breakdown frequency – objective but no revealed

The Brier scoring rule

- Expert has belief q:
  - Sunny with proba q, rainy with proba 1-q
- Announces prediction p (proba of sunny)
- How to incentivize honest prediction?
  - Give him "score"
    - $S(p, \text{sunny}) = 1 - (1-p)^2$
    - $S(p, \text{rainy}) = 1 - p^2$
- Expected score $S(p, q) = 1-q+q^2-(p-q)^2$
  - Maximized at $p=q$

Proper scoring rules

- Definition: a scoring rule is proper if
  $$S(q, q) \geq S(p, q) \text{ for all } p$$

- It is strictly proper if the inequality is strict for all $p \neq q$

- Brier rule is strictly proper
- Other strictly proper scoring rule:
  - $S(p, \text{state}) = \log p_{\text{state}}$

Different settings

- How to enforce honest reporting of interaction experience?

1. Objective information publicly revealed: can just compare report to real outcome
   – E.g., weather prediction

2. No objective outcome is available
   – E.g., product quality – not objective
   – E.g., product breakdown frequency – objective but no revealed

Peering agreement rewarding

- Rewarding agreement is not good
- If a good outcome is likely (e.g., because of well noted seller), a customer will not report a bad experience

→peer-prediction method

- Use report to update a reference distribution of ratings (prior distribution)
- Reward based on comparison of probabilities of the reference rating and the actual reference report

Model

- Product of given quality (called type) observed with errors
- Each rater sends feedback to central processing center
- Center computes rewards based exclusively on raters indications (no independent information)

Model (2)

- Finite number of types t=1, …, T
- Commonly known prior Pr_0
- Set of raters I
  - Each gets a 'signal'
  - $S=\{s_1, \ldots, s_M\}$: set of signals
  - S^i: signal received by i, distributed as $f(.|t)$

Example

- Two types: H (high) and L (low)
  - $Pr_0(H)=.5$, $Pr_0(L)=.5$
- Two possible signals: h or l
- $f(h|H)=.85$, $f(l|H)=.15$, $f(h|L)=.45$, $f(l|L)=.55$
  - $Pr(h)=.65$, $Pr(l)=.35$

Game

- Rewards/others ratings revealed only after receiving all reports from all raters
- → simultaneous game

- x^i: i's report, $x = (x^1, \ldots, x^I)$: vector of announcements
- x^i_m: i's report if signal s_m
- i's strategy:
- $\tau_i(x)$: payment to i if vector of announcement x

Best Response

- Best response

- Truthful revelation is a Nash equilibrium if this holds for all i when $x^i_m = s_m$

Example

Scoring rules

- How to assign points to rater i based on his report and that of j?
- Def: a scoring rule is a function that, for each possible announcement assigns a score to each possible value s in S
- We cannot access s_j, but in a truthful equilibrium, we can use j's report
- Def: A scoring rule is strictly proper if the rater maximizes his expected score by announcing his true belief

Logarithmic scoring rule

- Ask belief on the probability of an event
- A proper scoring rule is the Logarithmic scoring rule: Penalize a user the log of the probability that he assigns to the event that actually occurred

Peer-prediction method

- Choose a reference rater r(i)
- The outcome to be predicted is $x^{r(i)}$
- Player i does not report a distribution, but only his signal
  - The distribution is inferred from the prior
- Result: For any mapping r, truthful reporting is a Nash equilibrium under the logarithmic scoring rule

Proof

Example

Remarks

- Two other equilibria: always report h, always report l
  - Less likely
- See other applications of Bayesian estimation by Amazon reviews in M. Chiang. “Networked Life, 20 Questions and Answers”, CUP 2012. Chapters 5.

Outline

1. Introduction
2. Eliciting effort and honest feedback
3. Reputation based on transitive trust

Transitive trust approach

- Assign trust values to agents that aggregate local trust given by others
- t(i, j): trust that i reports on j
- Graph
- Reputation values
- Determine a ranking of vertices

Example: PageRank

Example 2: max-flow algorithm

Slide in case you are ignorant about max-flow min-cut theorem

Example 3: the PathRank algorithm

Definitions

- Monotonic: if adding an incoming edge to v never reduces the ranking of v
  - PageRank, max-flow, PathRank

- Symmetric if the reputation F commutes with the permutation of the nodes
  - PageRank
  - Not max-flow, not PathRank

Incentives for honest reporting

- Incentive issue: an agent may improve their ranking by incorrectly reporting their trust of other agents
- Definition: A reputation function F is rank-strategyproof if for every graph G, no agent v can improve his ranking by strategic rating of others
- Result: No monotonic reputation system that is symmetric can be rank-strategyproof
  - PageRank is not
  - But PathRank is

Robustness to sybil attacks

- Suppose a node can create several nodes and divide the incoming trust in any way that preserves the total incoming trust
- Definition:
  - sybil strategy
  - Value-sybilproof
  - Rank-sybilproof

Robustness to sybil attacks: results

- Theorem: There is no symmetric rank-sybilproof function
- Theorem (stronger): There is no symmetric rank-sybilproof function even if we limit sybil strategies to adding only one extra node
- → PageRank is not rank-sybilproof

Robustness to sybil attacks: results (2)

- Theorem: The max-flow based ranking algorithm is value-sybilproof
  - But it is not rank-sybilproof

- Theorem: The PathRank based ranking algorithm is value-sybilproof and rank-sybilproof