The Social Cost of Cheap Pseudonyms

Game Theory Seminar

Ton van Deursen

Cooperation

How to achieve cooperation:

- Repetition: avoid negative consequences in the future with the same people.
- Reputation: avoid negative consequences in the future with different people.

Changing Identities

Changing Identities

"On the Internet, nobody knows you're a dog."

"On the internet, nobody knows that yesterday you were a dog, and therefore should be in the doghouse today."

Choices

On the internet, a person has a choice of:

- changing identity, or
- maintaining a persistent identity.

Choices

On the internet, a person has a choice of:

- changing identity, or
- maintaining a persistent identity.

The option of anonymity turns transfer of reputation information into a strategic variable, controlled by each player.

Negative Feedback

Problem: people can discard negative feedback

Solutions:

- Distrust newcomers $\Rightarrow$ Newcomers have to accept bad treatment for a while (dues-paying) $\Rightarrow$ Inefficient.
- Trust newcomers until they proved untrustworthy $\Rightarrow$ There is an incentive to misbehave and then change identifiers.
- Disallow anonymity $\Rightarrow$ Privacy issues $\Rightarrow$ Some communities would not function anymore.

Main challenge

We seek a mechanism that

- encourages players to maintain a persistent identifier, and
- does not rely on verification and revelation of identities.

The basic model

We consider a repeated game:

- with periods $t = \{0, 1, 2, \ldots\}$,
- with M active players,
- where every period αM players exit and the same number enters,
- where players play a prisoner's dilemma with payoffs:

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

- where active players may change their identifier,
- where the history of the games is known,
- where the own history of name changes is known.

The basic model (cnt.)

Player i's strategy in period t is a mapping:

$$s_i^t : H_s^t \times H_i^t \times H_E^t \to \Delta(\{C, D\})$$

The basic model (cnt.)

Player i's strategy in period t is a mapping:

$$s_i^t : H_s^t \times H_i^t \times H_E^t \rightarrow \Delta(\{C, D\})$$

The total payoff for player i, when playing strategy s is:

$$u_i(s) = \sum_{t=b(i)}^{b(i)+l(i)} u_i^t$$

where u_i^t is the payoff in period t.

Benchmark

Normalized per-period payoff:

$$\alpha E[u_i(s)]$$

Benchmark

Normalized per-period payoff:

$$\alpha E[u_i(s)]$$

Benchmark for the amount of cooperation will be the average among all the players of the expected per-period payoff:

$$V(s) = \liminf_{N\to\infty} \frac{\sum_{i=0}^{N} \alpha E[u_i(s)]}{N}$$

Benchmark

Introduction

The model
-The basic model
-LPS
-PGTS

Malicious player and trembles

Payments for identifiers

Identifier commitments

Conclusion

Normalized per-period payoff:

$$\alpha E[u_i(s)]$$

Benchmark for the amount of cooperation will be the average among all the players of the expected per-period payoff:

$$V(s) = \liminf_{N \to \infty} \frac{\sum_{i=0}^{N} \alpha E[u_i(s)]}{N}$$

- $V = 1$ if every player cooperates in every period
- $V = 0$ if every player defects in every period.

Localized punishment strategy (LPS)

Localized punishment strategy:

- Play C against a newcomer
- Play C against a veteran who complied with LPS in the previous period
- Play D against a veteran who deviated from LPS in the previous period.

Localized punishment strategy (LPS)

Localized punishment strategy:

- Play C against a newcomer
- Play C against a veteran who complied with LPS in the previous period
- Play D against a veteran who deviated from LPS in the previous period.

Average per-period payoff: $V(LPS) = 1$.

Localized punishment strategy (LPS)

Localized punishment strategy:

- Play C against a newcomer
- Play C against a veteran who complied with LPS in the previous period
- Play D against a veteran who deviated from LPS in the previous period.

Average per-period payoff: $V(LPS) = 1$.

No equilibrium if players can change identifier freely.

Public grim trigger strategy (PGTS)

Public grim trigger strategy:

- Play D if there has ever been a defection in earlier period.
- Play C otherwise.

Average per-period payoff: $V(LPS) = 1$.

The model

- A fraction ϵ of the players is malicious.
- The probability of trembles is ϵ.
- The probability of losing one's identifier is ϵ.

Let $V^*(\epsilon, M)$ be the supremum of $V(s)$, with population M.
The "stable value" is defined as

$$SV = \lim_{\epsilon \to 0} \lim_{M \to \infty} V^*(\epsilon, M)$$

For fixed identifiers we can prove the following proposition:

Proposition 1. *For all* $\alpha < 0.3$, $M > 1$, *and* $\epsilon < 0.1$, *LPS is an equilibrium with* $V(s) = 1 - O(\epsilon)$. *More precisely,* $V(S) \geq 1 - 2\epsilon$.

Localized punishment strategy (LPS)

For fixed identifiers we can prove the following proposition:

Proposition 3. *For all* $\alpha < 0.3$, $M > 1$, *and* $\epsilon < 0.1$, *LPS is an equilibrium with* $V(s) = 1 - O(\epsilon)$. *More precisely,* $V(S) \geq 1 - 2\epsilon$.

Corollary 4. *For the game with persistent identifiers* $SV = 1$.

Localized punishment strategy (LPS)

- Consider a single deviation from the equilibrium,
- When D is played when C is asked for the gain = 1,
- The next period, there will be a penalty of 2 with probability $(1-\alpha)(1-2\epsilon)$,
- Deviation is only profitable if $1 > 2(1-\alpha)(1-2\epsilon)$,
- Thus LPS is an equilibrium

Localized punishment strategy (LPS)

Per-period payoff for each player is:

- If 2 non-deviators meet:
  $(1-\epsilon)^2 \cdot 1 + \epsilon(1-\epsilon) \cdot 2 + \epsilon(1-\epsilon) \cdot -1 + \epsilon^2 \cdot 0 = 1-\epsilon$
- If 2 deviators meet:
  $(1-\epsilon)^2 \cdot 0 + \epsilon(1-\epsilon) \cdot -1 + \epsilon(1-\epsilon) \cdot 2 + \epsilon^2 \cdot 1 = \epsilon$
- If a non-deviator meets a deviator: $1+3\epsilon$
- If a deviator meets a non-deviator: $2-3\epsilon$

Localized punishment strategy (LPS)

Per-period payoff for each player is:

- If 2 non-deviators meet:
  $(1-\epsilon)^2 \cdot 1 + \epsilon(1-\epsilon) \cdot 2 + \epsilon(1-\epsilon) \cdot -1 + \epsilon^2 \cdot 0 = 1-\epsilon$
- If 2 deviators meet:
  $(1-\epsilon)^2 \cdot 0 + \epsilon(1-\epsilon) \cdot -1 + \epsilon(1-\epsilon) \cdot 2 + \epsilon^2 \cdot 1 = \epsilon$
- If a non-deviator meets a deviator: $1+3\epsilon$
- If a deviator meets a non-deviator: $2-3\epsilon$

If there were k deviations in the previous period, the average payoff will be:

$$(\frac{k}{M})^2 \cdot \epsilon + (1-\frac{k}{M})(\frac{k}{M})(-1+3\epsilon) + (1-\frac{k}{M})(\frac{k}{M})(2-3\epsilon) + (1-\frac{k}{M})^2(1-\epsilon)$$

Using the fact that $E[\frac{k}{M}] = \epsilon$, we get $1-2\epsilon+2\epsilon^2$ which is larger than $1-2\epsilon$.

Public grim trigger strategy (PGTS)

For any $\epsilon > 0$, $V(PGTS) = 0$.

Paying your dues (PYD)

Paying your dues:

- Rewards positive reputation rather than punishing negative reputation.
- When a newcomer meets a veteran, the newcomer chooses C and the veteran chooses D.

Paying your dues (PYD)

Paying your dues:

- Rewards positive reputation rather than punishing negative reputation.
- When a newcomer meets a veteran, the newcomer chooses C and the veteran chooses D.

Different forms of dues-paying:

- In the prisoner's dilemma: defection.
- In eBay: accepting poor treatment/lower prices.
- In Magic card trades: being obliged to initiate the trade.

PYD more formally

There are two types of players:

- Entrants (newcomers): somebody who has not played before.
- Veterans: somebody who has played before.

PYD more formally

There are two types of players:

- Entrants (newcomers): somebody who has not played before.
- Veterans: somebody who has played before.

Players are “in compliance” if their strategy conforms to the PYD-strategy:

- If both players are compliant: cooperate (C)
- If both players are not in compliance: defect (D)
- If the player is not in compliance: choose a new identifier.
- If a compliant veteran meets an entrant, the veteran defects (D) and the entrant cooperates (C). Only if $q < \hat{q}(\alpha, \epsilon, M)$:

$$\hat{q}(\alpha, \epsilon, M) = \frac{1 - \frac{1}{M}}{(1-\alpha)(2 - \alpha - \frac{2}{M} - \epsilon + \frac{\epsilon}{M} + \epsilon\alpha)(1 - 2\epsilon)}$$

Expected payoff

Proposition 5. *For $\alpha < 0.3$, $\epsilon < 0.1$, $M > 11$, and $\hat{q}(\alpha, \epsilon, M) \leq 1$, PYD is an equilibrium of the game with impersistent identities, where $V(s) = 1 - \frac{\alpha}{2-\alpha} - O(\epsilon) - O(1/M)$.*

Expected payoff

Proposition 7. *For* $\alpha < 0.3$, $\epsilon < 0.1$, $M > 11$, *and* $\hat{q}(\alpha, \epsilon, M) \leq 1$, *PYD is an equilibrium of the game with impersistent identities, where* $V(s) = 1 - \frac{\alpha}{2-\alpha} - O(\epsilon) - O(1/M)$.

Corollary 8. *For the game with impersistent identities* $SV \geq 1 - \frac{\alpha}{2-\alpha}$

Improvements

The PYD scheme can be slightly improved by adjusting the strategy:

- Omit dues for newcomers in any period following one where there are no deviations.
- Make dues that are to be paid dependent on the collective behavior of the veterans in the previous period.

Improvements

The PYD scheme can be slightly improved by adjusting the strategy:

- Omit dues for newcomers in any period following one where there are no deviations.
- Make dues that are to be paid dependent on the collective behavior of the veterans in the previous period.

While there can be improvements over the PYD equilibrium, the improvements are slight and the bound is tight:

- Veterans must receive expected payoffs that are sufficiently larger than entrants' payoffs to prevent someone from defecting and then returning in the following period.
- The "most efficient" way to create a differential between begin a veteran and begin an entrant is having the veteran defect against the entrant.

Payment for identifiers

Payment for identifiers:

- Makes dues paying explicit.
- Attains full efficiency.
- Prevents players from defecting and starting over with a new identifier.

Payment for identifiers:

- Makes dues paying explicit.
- Attains full efficiency.
- Prevents players from defecting and starting over with a new identifier.

Redistribution:

- Entry fees in $t+1$ distributed evenly among players, playing in period t.
- Players stay too long in the game.
- Not fair if players' expected lifetimes are heterogeneous.

Redistribution to players

If we do not redistribute entry fees to the players:

- Assume players' varying wealth causes them to value money differently.
- $\lambda \in (0, 1]$ where $\lambda = 1$ indicates poor players and $\lambda = 0.01$ indicates wealthy players.
- Expected payoff is $\frac{V(s)}{\alpha} - \lambda F$ where F is the entry fee.

The entry fee must be sufficiently large to prevent wealthy players from deviating, but will deter others from entering.

Intermediaries

We assume an intermediary:

- Assigns regular ID's to players
- Does not reveal which players received which ID's
- Issues once-in-a-lifetime ID's

Intermediaries

We assume an intermediary:

- Assigns regular ID's to players
- Does not reveal which players received which ID's
- Issues once-in-a-lifetime ID's

Strategy: Play D against regular ID's in the original strategy

- Using a once-in-a-lifetime ID signals commitment to keep using that ID
- Using a regular ID signals that you are untrustworthy

Intermediaries

The need for trust in the intermediary can be reduced using blind signatures.

Different arenas (societies) can have different intermediaries.

Trade-off

There is a trade-off between anonymity and accountability:

- Broader arenas increase accountability
- Broader arenas decrease anonymity

Conclusion

- In the Internet, positive reputations are valuable, but negative reputations do not stick
- Natural convention is to distrust or even mistreat newcomers
- It is better to create an environment in which newcomers are trusted until proven otherwise
- There is an inherent social cost to free name changes