Peer Monitoring, Ostracism and the Internalization of Social Norms

with Rohan Dutta and Salvatore Modica

Introduction

- build on work showing the importance of self-enforcing social norms in enabling groups to overcome public goods problems (Olson, Ostrom)
- social norms are endogenous: Boyd-et-al cross-cultural experiments
- Indonesian whale hunters: need to share the catch, status determined by gift-giving, measured in experiments where ultimatum bargainers reject good offers

Our Model

elaborate on the model of peer incentives from Kandori, Levine/Modica and Levine/Mattozzil

an environment where monitoring is difficult (few monitors)

- individual behavior: Nash equilibrium with respect to selfish preferences
- collective decisions: groups can coordinate on a mutually advantageous equilibrium
- monitoring and penalties for anti-social behavioral
- internalization of social norms
- stickiness of social norms

Issues

- cost of punishing the monitor depends social closeness of monitor and producer: trade-off between information and incentives; rotation, supervisor versus peer review, police versus doctors
- optimality of social norms outside the laboratory may lead to the failure of procedures such as double-blind designed to reduce or eliminate possibility of outside influence
- tradeoff between social benefit and the social cost of monitoring: external incentives for public good contribution – substitute or complement? Perverse effects with fixed costs
- more general Lucas critique of experiments (lab, field, natural) – interventions may (or may not) change social norms depending upon circumstances
- does internalization complement or substitute incentives?
- cultural norms and strategic subsidies of internalization

The Base Model

- large group where monitoring is difficult in the sense that each production decision is observed by at most one other person.
- continuum of pairs with a unit mass
- pair consists of a producer and monitor

Technology

producer effort $e \in \{0, 1\}$ with cost ec where $c > 0$

value of public good: fraction of pairs producing ϕ per capita benefit ϕV

monitor costlessly observes noisy signal $z \in \{0, 1\}$: with probability π the signal is wrong; makes report $x \in \{0, 1\}$

social interaction: population is rematched into social subgroups of size $N \geq 4$; producer and monitor in same subgroup h

exactly one of the N members of each subgroup randomly chosen to be presenter and may volunteer to share an interesting story

$N - 1$ members of anonymous audience observe the report by or about the presenter and vote whether to ostracize; $1 < K < N - 1$ votes in favor lead to ostracism

presentation has value of N to the presenter and βN to each audience member

Truthful Strategies

truthful strategy:

- choice of whether or not to produce as a producer
- whether to send the message equal to the signal if a monitor
- always volunteer a story conditional on having one
- rule for ostracizing the presenter

social norm: a truthful strategy that if followed by everyone is a Nash equilibrium

collective decision: group chooses *optimal social norm* that maximizes the ex ante per capita utility of the identical group members (*social utility*)

Two Types of Social Norms

default norm

no effort

all stories to be volunteered

nobody ostracized

utility from only the social interaction $U = 1 + \beta(N-1)$.

implementation of production

$e = 1$

monitor tells the truth

all stories are volunteered

incentive compatible ostracism rule

note that all ostracism rules are incentive compatible for the audience because nobody is decisive

Implementing Production

potential social norms denoted by s correspond to ostracism probabilities $p(x), q(x)$ as function of the report $x \in \{0, 1\}$.

ostracizing one member of a pair imposes in expectation a cost of 1 on that person and a cost of $h\beta$ on the partner.

per capita probability of ostracism [on the equilibrium path]

$$\Pi(s) = (1/2)[\pi p(0) + (1-\pi)p(1)] + (1/2)[\pi q(0) + (1-\pi)q(1)].$$

social utility $W(s)$ is per capita payoff from production V minus the per capita cost of production (half the producer cost) plus utility from the social interaction minus the expected cost of ostracism:
$W(s) = V - c/2 + (1 - \Pi(s))U$.

Cost of Implementing Production

$\mathrm{W(s)=U+V-[\Pi}(s)U + c/2]$

cost of implementation $C(s) = \Pi(s)U + c/2$

monitoring cost plus production cost

optimal social norm must minimize implementation

implementation will be optimal if and only if $V \geq \min_s C(s)$.

Mechanism Design

principal (stand in for group) with two agents, monitor and producer

producer who chooses $e \in \{0, 1\}$ at a personal cost of ce and utility of $(V - c/2)e$ to the principal,

monitor who observes z and reports x to the principal.

principal can choose to punish either or both of the two agents.

punishment of either one has a cost to the punished of 1, a cost to the partner of $h\beta$ and a cost to the principal of $U/2$

principal - can precommit to punishment probabilities $p(x), q(x)$ as a function of the report of the monitor.

Cost Minimizing Social Norms

Theorem: *If and only if the* implementation condition

$$\frac{c}{(1-2\pi)(1-h^2\beta^2)} \leq 1$$

is satisfied can production be implemented. In the cost minimizing social norm producers who are reported to have taken the bad action ($x=0$*) are ostracized with probability* $p(0)=P$ *and monitors who report the good action (*$x=1$*) are ostracized with probability* $q(1)=Q$ *and there is no other ostracism. The ostracism probabilities are*

$$P=\frac{c}{(1-2\pi)(1-h^2\beta^2)}, \quad Q=h\beta P,$$

and the cost of implementation is

$$C=\left[\frac{U}{2}\frac{\pi+(1-\pi)h\beta}{(1-2\pi)(1-h^2\beta^2)}+\frac{1}{2}\right]c.$$

Further Discussion

- note the discontinuity: implementation fails abruptly
- feedback effect: a bigger punishment for the producer implies a bigger punishment for the monitor. The feedback effect is that the latter reduces the incentive for the producer to produce: by not producing she can reduce the probability the monitor is punished for sending a good report.
- must punish the monitor for good reports even though that is the only kind submitted and they are known to be true
- only way to get the monitor to tell the truth is to make her indifferent between the two reports. There is no mechanism or social norm in which the monitor strictly prefers to tell the truth
- malicious gossip is valued in the sense that a monitor is less likely to be ostracized for filing a bad report.
- cost of implementation is proportional to c the incentive to cheat on the social norm; standard result in peer monitoring

Rotation and Expertise

assume a trade-off of the form $\pi = f(h)$

f twice continuously differentiable with $f'(h) < 0$ and $f"(h) > 0$

(more social interaction between producer and monitor = better signal)

Theorem: *Let $C(h, \pi)$ denote the least cost of implementation if the implementation condition is satisfied and ∞ otherwise. If there exists a $h, f(h)$ such that the implementation condition is satisfied then there is a unique minimum of $C(h, \pi)$ subject to $\pi = f(h)$ and the optimum satisfies*

1. h is decreasing, π increasing in β

2. if h_f, π_f are the solutions of the cost minimization problem and $\overline{f}$ satisfies $\pi_f = \overline{f}(h_f)$ and greater signal sensitivity than f in the sense that $|\overline{f}'(h)| > |f'(h)|$ then $h_{\overline{f}} > h_f$ and $\pi_{\overline{f}} < \pi_f$.

Police versus Surgeons

surgeons require a high level of specialized knowledge: sensitivity of f to h is much greater for surgeons than for police officers

outsiders unlikely to have the specialized knowledge needed to evaluate “surgical output”; not so difficult for outsider to evaluate “police output.”

social network of surgeons sparser in the sense there are more about fifteen times as many police than surgeons

so good friends of police officers are more likely to be among other police officers than good friends of surgeons among other surgeons: β lower for surgeons than police officers

theorem says higher h for surgeons than for police officers.

indeed: police use supervisor evaluation and rotation to achieve low h while surgeons are self-policing

Alternative Monitoring Technologies

a fraction of monitors randomly assigned to a fraction of producers

producer may have no monitors, one monitor, or many monitors, randomly determined

who knows what about whom?

two extremes:

1. very few monitors so that the number of monitors per producer can as a good approximation be taken to be either zero or one, with the producer unaware of whether a monitor is present,

2. very many monitors all of whom observe exactly the same signal

our benchmark case lies between these two extremes

Few Monitors

η probability monitor is present to witness a production decision

only effect is to change the incentive constraint for the producer

$$\eta P = \frac{1}{1 - h^2\beta^2} \frac{c}{1 - 2\pi},$$

implementability accordingly harder to satisfy, but implementation cost does not change since larger punishments are used with smaller frequency

Many Monitors

many monitors who observe exactly the same signal

ostracize all monitors with probability one for disagreement

if all tell the truth all strictly prefer to tell the truth

in equilibrium no punishment of monitors

same as $h = 0$.

Urban Slum versus Poor Rural Village

urban slum: π large so no public goods (trash in the streets)

rural village π small

two types of transactions:

likely to be seen by many people

unlikely to be seen except perhaps by one

former case h is effectively zero: all the matters is that π is small

likely to see public good production in this case (Ostrom: water projects and so forth)

where monitoring is difficult h and β are large so implementability fails

expect rural villages to be like urban slums for public goods where production is hard to observe

What Economic Theory is For

hard to monitor: cheating of outsiders or tourists – one-on-one transactions with outsiders in a shop, hotel, or restaurant

public good element: cheating strangers gives the village a bad reputation so few tourists

modern technology has made it easier to monitor one-on-one transactions for hotels and restaurants – on line review services such as Trip Advisor not only allow tourists to avoid places they are likely to be cheated, but allow villagers to observe that a particular individual is engaging in cheating

better online information about hotels and restaurants should lead to social norms that discourage the cheating of outsiders in hotels and restaurants but not in shops (jewelry, souvenirs, clothing, art).

Double-Blind in the Laboratory

it is believed that participants behave altruistically in laboratory dictator experiments to make a good impression on experimenter

double-blind treatment used to eliminate this ("what happens in Vegas stays in Vegas")

we believe that what participants are "worried" about getting discovered to have violated a social norm from outside the laboratory

unlike the literature we do not think that representations of double-blind are blindly believed

1. Mistakes happen. If hackers can obtain confidential and damaging emails from Yahoo, what are the chances the experimental records are so secure that they will never leak to the outside world?

2. Even if identities are protected – for example through double-blind – there is a long history of deception in experiments by psychologists who have systematically lied to their subjects. What, for example, is to keep a deceptive experimenter from using a secret camera to record supposedly confidential placement of money into an envelope?

Double Blind Model

only a chance η of being monitored (the probability of a leak) and $h = 0$ since monitor incentives are not relevant when there is a public release of information

through instructions, design, and reputation, the perceived value of η may be made small but not zero

subjects have some concern that if they behave selfishly in the laboratoryword of this will get back to their friends outside the laboratory and they will then have an unfortunate reputation for behaving badly when they think nobody is looking

theory says that a reduction in η that is not sufficiently great will simply raise the probability of ostracism but have no effect on behavior

in other words: no effect until enough effort is made, then selfishness

data from dictator meta-studies suggests this is in fact the case

Cost Versus Benefit with Subsidies

choice of production level or quality $0 \leq \theta \leq \Theta$ where Θ is large

cost of producing is θ^2 ; producer produces $e\theta$

$e = 1$ still means the norm is followed; same signalling technology

producer who chooses not to follow the social optimally deviates to 0

cost coefficient of public good production

$$\mu = \frac{U}{2} \frac{\pi + (1-\pi)h\beta}{(1-2\pi)(1-h^2\beta^2)} + \frac{1}{2}$$

group is maximizing $V\theta - \mu\theta^2$

Always Produce

optimal social norm is given by $\theta = V/(2\mu)$

social utility $V^2/(4\mu)$

strictly positive so it is always better to implement production rather than use the default social norm

obvious result that θ is strictly decreasing in μ

Subsidies

cost of production is defrayed by subsidy $r\theta$ taken from the value of the public good $V\theta$

outsiders have better information than the group as they directly observe θ (example the IRS)

optimal deviation for an individual is no longer to produce 0, rather it is to maximize utility net of punishment $r\theta - \theta^2$, that is, to produce $r/2$ and receive a utility net of punishment of $r^2/4$.

Direct cost of production $d = \theta^2$.

Utility gain from deviating $c = r^2/4 - (r\theta - \theta^2)$.

producer incentive constraint

$$P = \frac{c}{(1-2\pi)(1-h^2\beta^2)}.$$

Effect of Subsidy

Theorem: *It is always optimal to implement production, at the level*

$$\theta = \left(V + (\mu - 1/2)\, r\right) / (2\mu)$$

which as expected is increasing in r*. The social utility advantage of implementing least cost production over the default equilibrium where* $r/2$ *is produced is*

$$G(r) \equiv (V + (\mu - 1/2)r)^2/(4\mu) + (1-\mu)r^2/4 - Vr/2$$

which satisfies $G(0) = V^2/(4\mu)$*,* $G(2V) = 0$ *and* $G'(r) < 0$ *for all* $r < 2V$*.*

Fixed Cost

natural to think that organizing a non-default mechanism with active monitoring an punishments has a fixed cost $F > 0$ associated with it

so implement production only if the utility gain over the default exceeds the fixed cost

Corollary: Let $\overline{r}$ be the positive solution of $G(r) = F$ if one exists, 0 otherwise. Note that for F sufficiently small a positive solution always exists and that $\overline{r} < 2V$. Then for $r < \overline{r}$ it is optimal to implement production and output is $\theta = (V + (\mu - 1/2)\, r)\, /(2\mu)$ while for $r > \overline{r}$ the default social norm is optimal and output is $r/2 < (V + (\mu - 1/2)\, r)\, /(2\mu)$. Hence for $\overline{r} = 0$ output increases in r while for $\overline{r} > 0$ output increases in r up to $\overline{r}$, drops discontinuously, and then increases again. In either case social utility is always increasing in r.

Incentives and Experiments

this actually happens

- experiments beginning with Gneezy Rustichini show that introducing modest incentives can discourage the activity it is designed to promote
- their experiment: parents picking up children late at day care
- but notice: while in our theory fining parents for showing up late increases lateness – but also welfare

Generalized Lucas Critique

small interventions are unlikely to change social norms hence conclusions drawn from small interventions may mislead as the effect of large interventions

for example: subsidizing mosquito netting in a few villages is unlikely to change religion practices, but doing over an entire region may

the point is: in doing interventions it is generally assumed social norms are fixed and have no particular reason for being what they are

in fact: religious practices may be a well-chosen social norm to respond to circumstances

Investment in Social Norms

as before the group/principal announces a pure strategy σ called the social norm.

after this announcement and before matching, production and monitoring individuals may choose to invest (or specialize) in a pure strategy s of their choice

cost investment:

$a \geq 0$ if the strategy chosen $s = \sigma$ is the social norm,

$a + \Gamma$ if the strategy chosen $s \neq \sigma$ is not the social norm, where $\Gamma \geq 0$ is the benefit of conformity

it is less costly to learn the language used by everyone else than to invent your own language

choice of investment is known only to the investor: no punishment is possible based on the investment decision

Consequences of Investment

an investor gets utility from the strategy invested in

if s is chosen and the terminal node is consistent with s the investor receives a bonus of $B \geq 0$ the *value of commitment*

we assume $a + \Gamma \leq B$ so that investing in a strategy and following it is profitable

internalization means that individuals choose to invest in the social norm

observe that the group/principal should never choose a social norm that will not be internalized: it would always be better to announce as the social norm the equilibrium strategy chosen by members

Essential versus Inessential Indifference

solution of the basic model involved several forms of indifference

the producer is indifferent between producing and not producing

- inessential: can be made strict by punishing a little more for a bad signal

the monitor is indifferent between reporting 0 and 1

- essential: cannot be made strict; model not robust to introducing a small cost of observing the signal

the audience members are indifferent to ostracizing or not ostracizing

- essential: cannot be made strict; weakly dominant not to ostracize; model not robust to small probability unanimity is required for ostracism.

$B, \Gamma > 0$ makes all indifference inessential and the model robust

Conformity and Ostracism: Complements or Substitutes?

Theorem: *Suppose* $B, c > 0$. *Define*

$$\overline{P} = \frac{c}{(1-2\pi)(1-h^2\beta^2)}$$

$$\underline{P} = \overline{P} \cdot \max\left\{0,\, 1-(1+(1-2\pi)h\beta)\,\frac{B}{c},\, (1-\frac{B}{c})(1-h^2\beta^2)\right\}.$$

If $\underline{P} > 1$ *implementation of production is not possible. If* $\underline{P} \leq 1$ *then there exists* $(1+\pi)B/2 \geq \overline{\Gamma} \geq \underline{\Gamma} \geq 0$ *such that production can be implemented if and only if* $\Gamma \geq \underline{\Gamma}$ *where* $\underline{\Gamma} = 0$ *if and only if* $\overline{P} \leq 1$. *If* $B \geq c$ then $\overline{\Gamma} = c/2, \underline{P} = 0$ *and for* $\Gamma \geq \overline{\Gamma}$ *there is* complete internalization*: production is implemented without ostracism*

Illustration of the Remainder of the Theorem

If $\Gamma \geq \underline{\Gamma}$ the cost minimizing internalized social norm implements production and, for generic parameter values, ostracism probabilities are given by unique continuous piecewise linear functions $\hat{P}(\Gamma)$ and $\hat{Q}(P)$; if $\pi \leq .333$ then for $\hat{P}(\underline{\Gamma}) < P < \hat{P}(\overline{\Gamma})$ the producer strictly prefers to produce, and

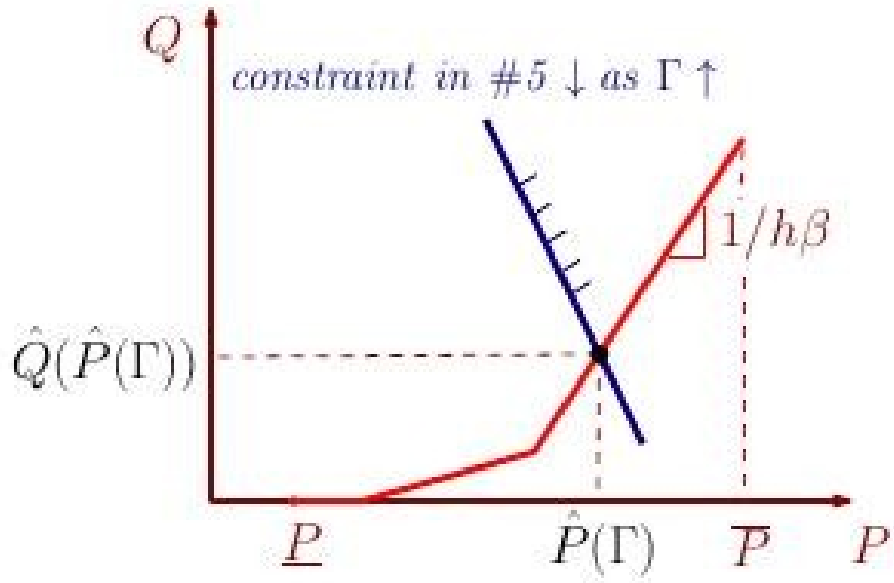

Observations

both B, Γ are needed

if both are large enough there is complete internalization

suppose production cannot be implemented without internalization and $B < c$ so that monitoring is needed to implement production: then both internalization and monitoring are needed: they are complements

as the benefit of conformity is increases internalization reduces the need for monitoring and they are substitutes

benefits of being able to implement may be disproportionate: even if B, Γ are quite small if they enable implementation the gain is on the order of the value of production V which can be very large

the value of commitment loosens the incentive constraints

the benefit of conformity is like credit that can be spent on either the monitor or producer: in the not too noisy signal case it should be spent on the monitor: the producer strictly prefers to produce

Social versus Cultural Norms

- individuals choose social norms
- cultural norms are generally derived at a young age from others, especially parents and peers
- cultural norms require a much larger investment and have a much greater value of commitment
- should be part of the same theory as that of social norms
- investment in strategies can be subsidized by interested parties
- public schools teach national myths; fight over curriculum is over history, language, religion – not arithmetic or reading
- combine with Bisin-Verdier horizontal/vertical models?