Peer Monitoring, Ostracism and the Internalization of Social Norms

with Rohan Dutta and Salvatore Modica

Introduction

- builds on work showing the importance of self-enforcing social norms in enabling groups to overcome public goods problems creating incentives through ostracism (Olson, Ostrom)
- builds on political economy models of peer discipline
- social norms are endogenous: (Boyd-et-al cross-cultural experiments)
- social norms often change slowly: distrustful/dishonest norms often survive for centuries (Bigoni-et-al Italian north/south divide)
- social norms may change rapidly: 9/11 story

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
- stickiness of social norms
- internalization of social norms

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 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 $U = N$. (note normalization)

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 on the partner (note normalization)

per capita probability of ostracism [on the equilibrium path] $\Pi(s)$

group objective

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

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 $V \geq \min_s C(s)$.

Cost Minimizing Social Norms

Theorem: *If and only if the* implementation condition

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

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*

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

and the cost of implementation is

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

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.
- malicious gossip is valued in the sense that a monitor is less likely to be ostracized for filing a bad report.

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 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."

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

message: the coziness of surgeons really is a problem – they will get away with more bad stuff (implementation only for large V)

less skilled professions will be subject to greater discipline because "it can be done"

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

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

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 $h = 0$.

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

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 – 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 laboratory word 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 $e\theta$

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)

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

Incentives and Experiments

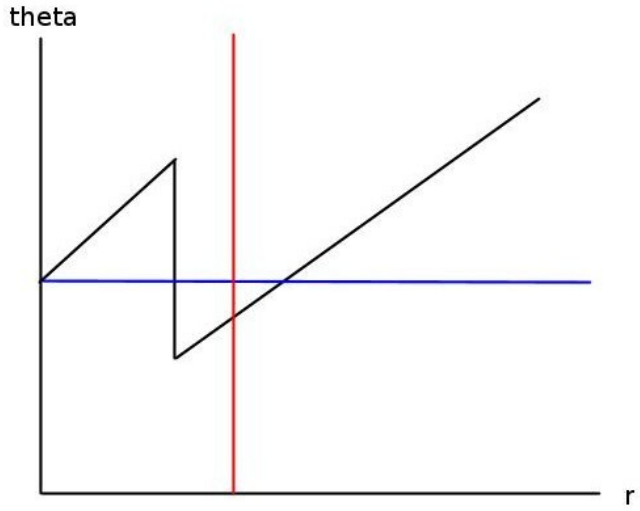

this actually happens: Gneezy Rustichini picking up kids on time at daycare

but: welfare is increased!!!

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

cost of investment: $a(s) \geq 0$

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

benefit of conformity $\Gamma = (b(\sigma) - a(\sigma)) - (b(s) - a(s))$ constant and non-negative

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

Internalization

internalization means that individuals choose to invest in the social norm

observe that the group 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 – can lead to huge welfare gain if stakes are high

Basics

- you need both B and Γ but for intermediate levels of both they can substitute; if both are large enough complete internalization is achieved
- you may need both internalization and incentives, but once feasibility is established they substitute
- Γ is like a budget that can be spent either reducing incentives for producer or monitor
- if $\pi \leq .333$ then use incentives for producer, internalization for monitor
- if $\pi \geq .382$ then use incentives for monitor, internalization for producer

Social versus Cultural Norms

really a continuum: think of cultural norms as requiring higher investment and involve larger stakes

now we assume a variable investment technology $b(s) = (1/\beta)a$ where $\beta > 1$ so you never invest

the group can subsidize the cultural norm with social cost $(1/\beta)a$ (perfect observability of contribution to this public good)

group imposes a limit A

so: $B = \Gamma = (1/\beta)A$

how should the group choose B

revisit the variable quality model: $V\theta - (1/2)\theta^2 - C(\theta, B) - \beta B$

Expected Conclusion

as V increases get more θ, B, P, Q

is there more internalization?

Ache: “Successful hunters often leave their prey outside the camp to be discovered by others, carefully avoiding any hint of boastfulness.”
what about the lab where there is a much smaller v than the one the cultural norm is based on?

In the Laboratory

as V increases for fixed v we have lab θ goes up, B is already fixed so cannot change

ultimatum bargaining (Henrich et al)

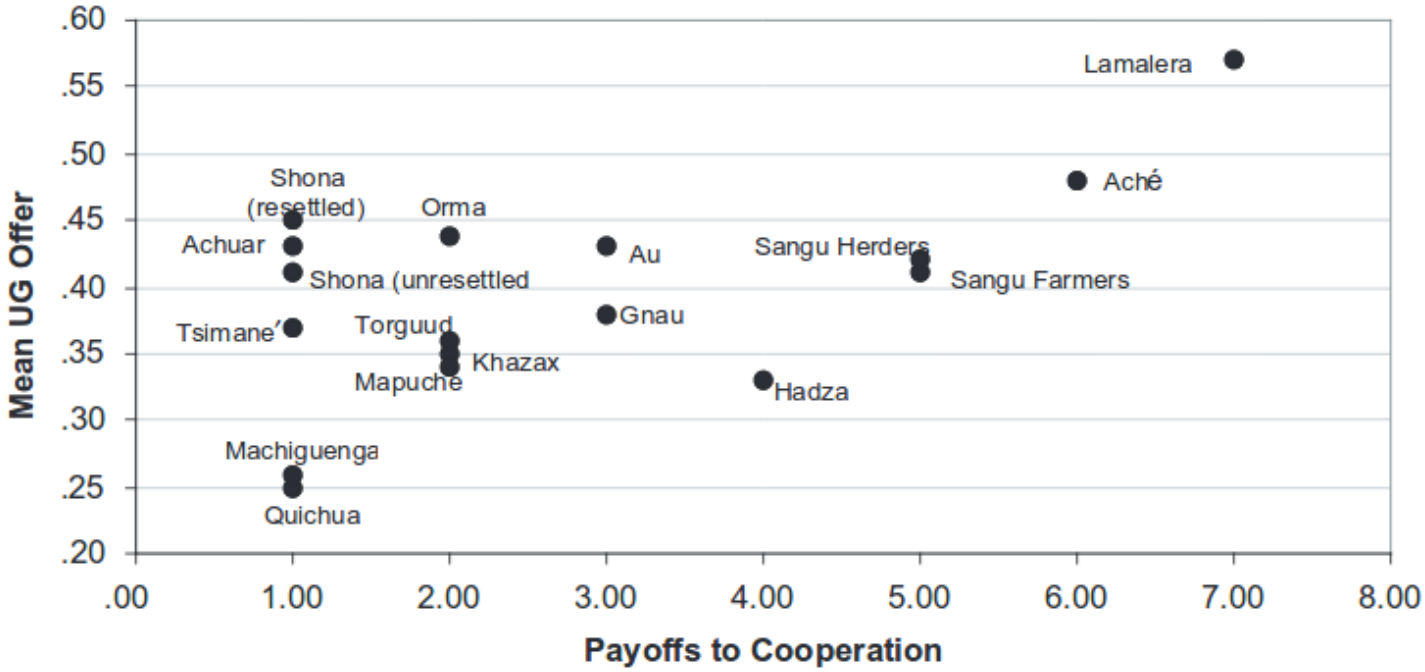

How About Punishment?

in the lab for very low V the social norm is the default, so no punishment for any v; for very large V internalization B is very large so no punishment in the lab (U-shaped punishment curve in the lab)

ultimatum bargaining

- Machigeuenga/Quichua/Ache: 1 rejection
- (5%) in Machigeuenga (out of 21 pairs)
- average for all others 12% except Lamalera
- deception was used in the Lamalera case: sham offers were reported to responder – way out of range of any offers in the data
- why you should never use deception...

careful use of theory enables us to make effective use of laboratory data – and also to design better experiments

Strategy and Conflict

- 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?