Renunciation Games

Jan van Eijck
CWI & ILLC, Amsterdam

IRIT workshop talk, Toulouse, September 3, 2013

Abstract

Renunciation games are games where an individual (or small group) is pitted against a collective. The setup of the games is such that the social optimum of a game can only be reached at the expense of a single individual (or small group). When will an individual sacrifice his/her own interest to save society? It turns out that the nature of the renunciation game changes crucially depending on the temptation offered to the renouncer. Eijck [to appear, 2013]

Overview

- Game Theory and the Structure of Society
- Terminology for Strategic n Player Games
- Punishment and Reward as Game Transformations
- Proportionality of Punishment, Social Harm
- Repression Cost of a Game
- From Punishment to Reward: Welfare Redistribution
- Civilization Cost of a Game
- Renunciation Games: Self-Sacrifice to Save Society
- Further Questions

Game Theory and the Structure of Society

- GT: an analysis tool using extremely abstract/crude models of 'social reality.'
- Hopefully we can still learn something from GT analysis of what goes on in the social world that we live in.
- Cf. Parikh [2002].
- Use (strategic) GT to analyze mechanisms of Punishment and Reward

Some Terminology for Strategic n Player Games

A strategic game G is a tuple

$$(\{1, \ldots, n\}, \{S_i\}_{i \in \{1,\ldots,n\}}, \{u_i\}_{i \in \{1,\ldots,n\}}),$$

where

- $\{1, \ldots, n\}$ with $n > 1$ is the set of players,
- each S_i is a set of strategies (the strategies for player i),
- each u_i is a function from $S_1 \times \cdots \times S_n$ to $\mathbb{R}$ (the utility function for player i).

Use N for $\{1, \ldots, n\}$, S for $S_1 \times \cdots \times S_n$ and u for $\{u_i\}_{i \in \{1,\ldots,n\}}$.
Then (N, S, u) denotes a game.

Strategy Profiles

A member of $S_1 \times \cdots \times S_n$ is a **strategy profile**: each player i picks a strategy $s_i \in S_i$.

- Use s to range over strategy profiles.
- Use s_{-i} for the strategy profile that results by deleting strategy choice s_i of player i from s.
- Let (s'_i, s_{-i}) be the strategy profile that is like s for all players except i, but has s_i replaced by s'_i.
- Let S_{-i} be the set of all strategy profiles minus the strategy for player i (the product of all strategy sets minus S_i). Note that $s_{-i} \in S_{-i}$.

Best Response, Nash Equilibrium

A strategy s_i is a best response in s if

$$\forall s_i' \in S_i \; u_i(s) \geq u_i(s_i', s_{-i}).$$

A strategy profile s is a (pure) Nash equilibrium if each s_i is a best response in s:

$$\forall i \in N \; \forall s_i' \in S_i \; u_i(s) \geq u_i(s_i', s_{-i}).$$

Let $\text{nash}(G) = \{s \in S \mid s \text{ is a Nash equilibrium of } G\}$.

A game G is Nash if G has a (pure) Nash equilibrium.

Social Welfare, Social Optimum

Define a social welfare function $W : S_1 \times \cdots \times S_n \to \mathbb{R}$ by setting

$$W(s) = \sum_{i=1}^{n} u_i(s).$$

A strategy profile s of a game $G = (N, S, u)$ is a social optimum if

$$W(s) = \sup\{W(t) \mid t \in S\}.$$

For a finite game, s is a social optimum if $W(s)$ is the maximum of the welfare function for that game.

Punishment and Reward as Game Transformations

Let us start with the Prisoner's Dilemma (PD) Game:

	II cooperates	II defects
I cooperates	$3, 3$	$0, 4$
I defects	$4, 0$	$1, 1$

Suppose a 'social software engineer' has to design a policy to make defection less profitable in a PD situation.

One thing she could do is put a penalty P on defection.

This does not have an immediate effect. A penalty can only be imposed if the cheater gets caught. Suppose the probability of getting caught is γ. In case the cheater gets caught, she gets the penalty, otherwise she gets what she would have got in the original game.

Transforming the PD Game

Adopting the punishment policy amounts to a change of utilities.

The policy change can be viewed as a game transformation that maps strategic game G to strategic game $G^{\gamma P}$, where $G^{\gamma P}$ is like G except for the fact that the utility function is replaced by:

$$\begin{aligned} u_I^{\gamma P}(c, c) &= u_I(c, c), \\ u_I^{\gamma P}(d, c) &= \gamma P + (1 - \gamma) u_I(d, c), \\ u_I^{\gamma P}(c, d) &= u_I(c, d), \\ u_I^{\gamma P}(d, d) &= \gamma P + (1 - \gamma) u_I(d, d), \end{aligned}$$

and similarly for u_{II}^{pP}.

The new utility of cheating if the other is honest amounts to P in case you get caught, and to the old utility of cheating in case you can get away with it.

Compare also the social ties utilities function in Chapter 6 of Moison [2013].

Something is missing still ...

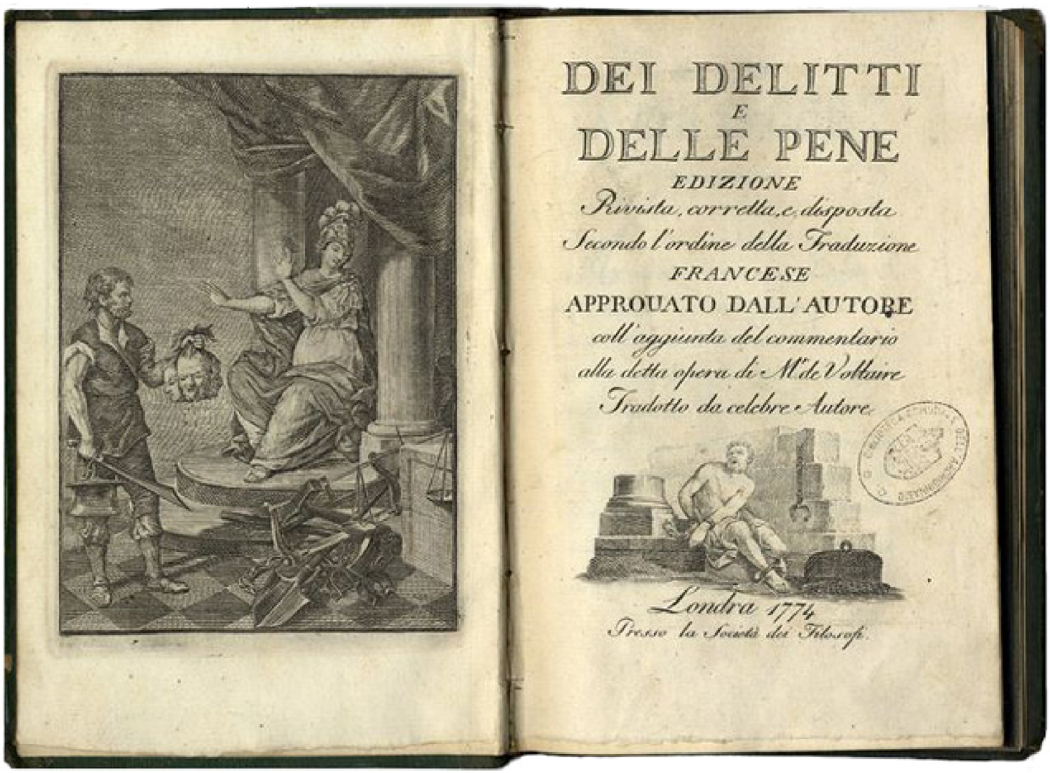

Proportionality of Punishment

> If an equal punishment be ordained for two crimes that injure society in different degrees, there is nothing to deter men from committing the greater as often as it is attended with greater advantage. [Beccaria, 1764, Ch 6]

A Measure for Social Harm Caused by the Strategy of an Individual Player

Let a game $G = (N, S, u)$ be given.

For any $i \in N$, define the individual harm function $H_i : S \to \mathbb{R}$, as follows:

$$H_i(s) = W(s) - \sup_{s'_i \in S_i} W(s'_i, s_{-i}).$$

This gives the difference between the social welfare for the profile s and the best outcome for society as i unilaterally deviates from her current strategy (against her own interest).

That is, $H_i(s)$ gives a measure for how much player i harms society by playing s_i rather than the alternative s'_i that ensures the maximum social welfare.

Clearly, in case s is a social optimum, $H_i(s) = 0$ for any i.

Appropriate Punishment Transformations of a Game

If $G = (N, S, u)$ is a strategic game, $\gamma \in [0, 1]$, then G^γ is the game (N, S, u^γ), where u^γ is given by:

$$u_i^\gamma(s) := (1 - \gamma)u_i(s) + \gamma H_i(s).$$

Think of γ as the probability of getting caught.

For the PD example:

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

$\Rightarrow$

	c	d
c	$3, 3$	$0, 4 - 6\gamma$
d	$4 - 6\gamma, 0$	$1 - 3\gamma, 1 - 3\gamma$

Explanation: social harm caused by playing d in case the other plays c is -2, social harm caused by playing d in case the other plays d is also -2.

The Repression Cost of a Game

The repression cost of a game G is the least γ for which the move from G to G^γ turns a social optimum into a Nash equilibrium. In case G has no social optimum, the repression cost is undefined.

Intuition behind this: a strong police force is needed to make detection and punishment very probable, so γ can be taken as a measure for repression.

Games that have a Nash social optimum have repression cost equal to 0: no game transformation is needed to create a social optimum that is Nash.

A game G with a social optimum, but without Nash social optimum can always be transformed into a game G^γ with a social optimum that is Nash.

Illustration: The Repression Cost of the PD Game

Take another look at the transformation:

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

$\Rightarrow$

	c	d
c	$3, 3$	$0, 4 - 6\gamma$
d	$4 - 6\gamma, 0$	$1 - 3\gamma, 1 - 3\gamma$

To turn (c, c) into a Nash equilibrium, we have to demand

$$3 \geq 4 - 6\gamma.$$

This gives: $\gamma \geq \frac{1}{6}$, so PD$^{\gamma}$ has a social optimum that is Nash iff $\gamma \in [\frac{1}{6}, 1]$, and the repression cost for the PD game is $\frac{1}{6}$.

From Punishment to Reward: Welfare Redistribution

Instead of punishing offenders directly, one can also punish them indirectly …

… by raising taxes for everybody, and distributing the tax revenue uniformly over the players.

Welfare redistribution can be represented by a map from payoff functions to new payoff functions, i.e., by a game transformation.

The Welfare Redistribution Map

The map for welfare redistribution is $G \mapsto G[\gamma]$, where $\gamma \in [0, 1]$, and the payoff u_i^γ in the new game $G[\gamma]$ is computed from the payoff u_i in G (assuming there are n players) by means of:

$$u_i^\gamma(s) = (1 - \gamma)u_i(s) + \gamma \frac{W(s)}{n}.$$

Here $W(s)$ gives the result of the welfare function on s in G.

Thus, player i is allowed to keep $1 - \gamma$ of her old revenue $u_i(s)$, and gets an equal share $\frac{1}{n}$ of $\gamma W(s)$, which is the part of the welfare that gets redistributed.

This definition is mentioned (but not used) in Chen and Kempe [2008].

Note that for any game G, any $\gamma \in [0..1]$, any s:

$$W(s) = W^\gamma(s).$$

The Civilization Cost of a Game

The civilization cost of a game G is the least γ for which the move from G to $G[\gamma]$ turns a social optimum into a Nash equilibrium. In case G has no social optimum, the civilization cost is undefined.

> The players are not "selfish"; rather the preferences of each player i are represented by the payoff function $m_i(a)+\alpha m_j(a)$, where $m_i(a)$ is the amount of money received by player i when the action profile is a, j is the other player, and α is a given non-negative number. [Osborne, 2004, exercise 27.1]

Compare Apt and Schaefer [2012], who define the selfishness level of a game. Let $G(\alpha)$ be the result of adding $\alpha W(s)$ to each payoff in s. The selfishness level of G is the least α for which the move from G to $G(\alpha)$ turns some social optimum of G into a Nash equilibrium.

Illustration: The Civilization Cost of the PD Game

Redistribution of part γ of social welfare transforms the PD game.

Note that if welfare is already equally distributed, the payoff transformation changes nothing.

In the case $W(c,d) = W(d,c) = 4$, we get that the payoffs of the players get changed, into $(1-\gamma)4 + \frac{4}{2}\gamma = 4 - 2\gamma$ for the player who defects, and into $\frac{4}{2}\gamma = 2\gamma$ for the player who cooperates.

	c	d
c	$3,3$	$0,4$
d	$4,0$	$1,1$

$\Rightarrow$

	c	d
c	$3,3$	$2\gamma, 4-2\gamma$
d	$4-2\gamma, 2\gamma$	$1,1$

Thus, the social optimum of $G[\gamma]$ is Nash iff $3 \geq 4 - 2\gamma$ iff $\gamma \geq \frac{1}{2}$. So the civilization cost of the PD game is $\frac{1}{2}$.

Illustration: The ToC Game

The Tragedy of the Commons game scenario was first analyzed in Gordon [1954] and was made famous in an essay by Garrett Hardin:

> The tragedy of the commons develops in this way. Picture a pasture open to all. It is to be expected that each herdsman will try to keep as many cattle as possible on the commons. Such an arrangement may work reasonably satisfactorily for centuries because tribal wars, poaching, and disease keep the numbers of both man and beast well below the carrying capacity of the land. Finally, however, comes the day of reckoning, that is, the day when the long-desired goal of social stability becomes a reality. At this point, the inherent logic of the commons remorselessly generates tragedy. [Hardin, 1968]

Formal Version

ToC can be viewed as a multi-agent version of the PD game. Assume there are n players. Chapter 1 of Vazirani et al. [2007] proposes the following model. The players each want to have part of a shared resource. Setting the value of the resource to 1, each player i has to decide on the part of the resource x_i to claim, so we can assume that $x_i \in [0, 1]$.

Stipulate the following payoff function. Let N be the set of agents. If $\sum_{j \in N} x_j < 1$ then the value for player i is

$$u_i = x_i(1 - \sum_{j \in N} x_j).$$

The benefit for i decreases as the resource gets exhausted. If $\sum_{j \in N} x_j \geq 1$ (the demands on the resource exceed the supply), the payoff for the players becomes 0.

Nash Equilibrium of ToC Game

Take the perspective of player i. Let D be the total demand of the other players, i.e., $D = \sum_{j \in N, j \neq i} x_j < 1$. Then strategy x_i gives payoff $u_i = x_i(1 - (D + x_i))$, so the optimal solution for i is $x_i = (1 - D)/2$.

Since the optimal solution for each player is the same, this gives $x = \frac{1-(n-1)x}{2}$, and thus $x = \frac{1}{n+1}$ as the optimal strategy for each player.

This gives $D+x = \frac{n}{n+1}$, and payoff for x of $u = \frac{1}{n+1}(1-\frac{n}{n+1}) = \frac{1}{(n+1)^2}$.

The total payoff is $\frac{n}{(n+1)^2}$, which is roughly $\frac{1}{n}$. This means that the social welfare in the Nash equilibrium for this game depends inversely on the number of players.

Social Optimum of the ToC Game

If the players had agreed to leave the resource to a single player, the total payoff would have been $u = x(1 - x)$, which is optimal for $x = \frac{1}{2}$, yielding payoff $u = \frac{1}{4}$.

If the players had agreed to demand only equal shares of $\frac{1}{2}$ of the resource, they would have had a payoff of $\frac{1}{4n}$ each, which is much more than $\frac{1}{(n+1)^2}$ for large n.

Illustration: The Civilization Cost of the ToC Game

A social optimum s in the ToC game satisfies $W(s) = \frac{1}{4}$.

We can now calculate just how much welfare we have to distribute for a given alternative to social optimum s to lose its appeal for i. A tempting alternative s' for i in s loses its appeal for i in s when the following holds:

$$u_i^\gamma(s') \leq u_i^\gamma(s).$$

Write out the definition of u_i^γ:

$$(1-\gamma)u_i(s') + \gamma\frac{W(s')}{n} \leq (1-\gamma)u_i(s) + \gamma\frac{W(s)}{n}.$$

Solving for γ yields $\ldots \gamma = 1$.

Since the social optimum s was arbitrary, it follows that the cost of civilization for the tragedy of the commons game is 1.

Renunciation Games: Self Sacrifice to Save Society

Some new games (not from the textbooks) where an individual is pitted against a collective. The setup of the games is such that the social optimum of the game can only be reached at the expense of one single individual.

When will an individual sacrifice his or her own interest to save society?

Pure Renunciation Game

The pure renunciation game has n players, who each choose a strategy in $[0, 1]$, which represents their demand. If at least one player renounces (demands 0), then all other players get as payoff what they demand. Otherwise, nobody gets anything. The payoff function for i is given by:

$$u_i(s) = \begin{cases} s_i & \text{if } \exists j \neq i : s_j = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Analysis

This game has n social optima

$$(0, 1, \ldots, 1), (1, 0, 1, \ldots, 1), \ldots, (1, \ldots, 1, 0),$$

where the social welfare W equals $n - 1$. The social optima are also Nash equilibria. No need for welfare redistribution, no need for punishment.

The repression cost and the civilization cost of this game are both 0.

The situation changes if there is a temptation for the renouncer in the game.

Renunciation Game With Mild Temptation

This renunciation game has n players, who each choose a strategy in $[0, 1]$, which represents their demand. If at least one player renounces (demands 0), then all other players get as payoff what they demand. Otherwise, if there is one player i who demands less than any other player, i gets what she demands, and the others get nothing. In all other cases nobody gets anything. The payoff function for i is given by:

$$u_i(s) = \begin{cases} s_i & \text{if } \exists j \neq i : s_j = 0 \\ & \text{or } \forall j \neq i : 0 < s_i < s_j \\ 0 & \text{otherwise.} \end{cases}$$

Analysis: Repression Cost

This game has n social optima

$$(0, 1, \ldots, 1), (1, 0, 1, \ldots, 1), \ldots, (1, \ldots, 1, 0),$$

where the social welfare W equals $n - 1$.

The social optima are not Nash equilibria. For in a social optimum, the player who renounces (and receives nothing) can get any q with $0 < q < 1$ by playing q. That's the temptation.

Suppose s is the profile where player i plays $q > 0$ and all other players play 1. Then $W(s) = q$, so the social harm i does by not renouncing is $q - (n - 1)$.

How can the social optimum where i renounces and all other players play 1 be turned into a Nash equilibrium by punishment?

By picking a value for γ that yields:

$$u_i^{\gamma}(s) = (1-\gamma)q + \gamma(x+1-n) \geq 0.$$

Solving for γ gives:

$$\gamma \geq \frac{q}{n-1}.$$

Taking the supremum for $q \to 1$ gives:

$$\gamma \geq \frac{1}{n-1}.$$

The repression cost for this game is $\frac{1}{n-1}$.

Analysis: Civilization Cost

The cost of civilization for the Renunciation Game is $\gamma = \frac{1}{2n-2}$.

Focus on player 1 and compute the least γ for which the social optimum $(0, 1, \ldots, 1)$ turns into a Nash equilibrium in $G[\gamma]$. The payoff function for player 1 in $G[\gamma]$ satisfies:

$$u_1^\gamma(0, 1, \ldots, 1) = \gamma \frac{n-1}{n}.$$

For the social optimum to be Nash, this value has to majorize

$$u_1^\gamma(q, 1, \ldots, 1) = (1-\gamma)q + \frac{\gamma}{n}q.$$

Since q can be arbitrarily close to 1, we get $u_1^\gamma(q, 1, \ldots, 1) < (1 - \gamma) + \frac{\gamma}{n}$, Therefore $(0, 1, \ldots, 1)$ is a social optimum in $G[\gamma]$ iff $\gamma \frac{n-1}{n} \geq (1-\gamma) + \frac{\gamma}{n}$. Solving this for γ gives $\gamma \geq \frac{1}{2n-2}$.

Renunciation Game With Heavy Temptation

The situation changes drastically if there is heavy temptation.

This renunciation game has n players, who each choose a strategy q in $[0, 1]$, which represents their demand. If at least one player renounces (demands 0), then all other players get as payoff what they demand. Otherwise, if there is one player i who demands less than any other player, i gets $n - 1$ times what she demands, and the others get nothing. In all other cases nobody gets anything. The payoff function for i is given by:

$$u_i(s) = \begin{cases} s_i & \text{if } \exists j \neq i : s_j = 0 \\ (n-1)s_i & \text{if } \forall j \neq i : 0 < s_i < s_j \\ 0 & \text{otherwise.} \end{cases}$$

Analysis: Repression Cost

Social optima are the same as before.

To compute the repression cost of the game, take the social optimum s where player i renounces. If i yields to temptation, the social harm is modest (but the harm to the other players is considerable). Let $s' = (q, s_{-i})$, with $0 < q < 1$. Then

$$H_i(s') = q(n-1) - (n-1) = (q-1)(n-1).$$

Putting in the equation for u_i^γ and solving for γ yields $\gamma = 1$. The repression cost for the game is 1.

Analysis: Civilization Cost

The civilization cost for Renunciation With Heavy Temptation is also 1.

We have to compute the least γ that turns social optimum $(0, 1, \ldots, 1)$ into a Nash equilibrium in $G[\gamma]$. The constraint on the payoff function for player 1 is:

$$u_1^\gamma(q, 1, \ldots, 1) = (1 - \gamma)(n - 1)q + \frac{\gamma}{n}(n - 1)q.$$

Since q can be arbitrarily close to 1, this gives

$$u_1^\gamma(q, 1, \ldots, 1) < (1 - \gamma)(n - 1) + \frac{\gamma}{n}(n - 1).$$

This puts the following constraint on γ:

$$\gamma\frac{n - 1}{n} \geq (1 - \gamma)(n - 1) + \frac{\gamma}{n}(n - 1).$$

Solving for γ gives $n\gamma \geq n$, and it follows that $\gamma = 1$.

Some Questions

- What is the formal relation between repression cost and civilization cost?
- Can we make sense of the notion of a social optimum in abstract strategic games, where utilities are replaced by preference orderings?
- Is it possible to define analogues of repression cost and civilization cost for abstract strategic games? Maybe we have to blend in some social choice theory to compare strategy profiles?
- What does this kind of analysis teach us about what goes on in social reality?
- Where does Logic come in?

References

K. R. Apt and G. Schaefer. Selfishness level of strategic games. In M. Serna, editor, Algorithmic Game Theory, volume 2012 of Lecture Notes in Computer Science, pages 13–24. Springer, 2012.

C. Beccaria. On Crimes and Punishment. Marco Coltellini, 1764. In Italian: Dei delitti e delle pene.

P.-A. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In Proceedings 10th ACM Conference on Electronic Commerce, pages 140–149, 2008.

J. v. Eijck. PDL as a multi-agent strategy logic. In B. C. Schipper, editor, TARK 2013 – Theoretical Aspects of Reasoning About Knowledge, Proceedings of the 14th Conference – Chennai, India, pages 206–215, 2013.

J. v. Eijck. Strategies in social software. In J. van Benthem, S. Ghosh, and R. Verbrugge, editors, Modeling Strategic Reasoning, Texts in Logic and Games. Springer, to appear.

H. S. Gordon. The economic theory of a common-property resource: The fishery. Journal of Political Economy, 62:124–142, 1954.

G. Hardin. The tragedy of the commons. Science, 162:1243–48, 1968.

F. Moison. The Bonds of Society: An Interdisciplinary Study of Social Rationality. PhD thesis, Université de Toulouse, 2013.

M. J. Osborne. An Introduction to Game Theory. Oxford University Press, New York, Oxford, 2004.

R. Parikh. Social software. Synthese, 132:187–211, 2002.

V. V. Vazirani, N. Nisan, T. Roughgarden, and E. Tardos. Algorithmic Game Theory. Cambridge University Press, 2007.