The Normative Power of Sanction. A Simulation Model

Giulia Andrighetto (giulia.andrighetto@istc.cnr.it)
European University Institute, Florence, Italy and Institute of Cognitive Science and Technology, CNR, Rome, Italy

Daniel Villatoro (dvillatoro@iiia.csic.es)
IIIA - Artificial Intelligence Research Institute, CSIC, Barcelona, Spain

Abstract

As specified by Axelrod in his seminal work An *Evolutionary Approach to Norms* (Axelrod, 1986), punishment is a key mechanism to achieve the necessary social control and to impose social norms in a self-regulated society. In this paper, we distinguish between two enforcing mechanisms punishment and sanction, focusing on the specific ways in which they favor the emergence and maintenance of cooperation. To achieve this task, we have developed a normative agent able to punish and sanction defectors, and we have run a proof-of-concept simulation to test our hypotheses.

Keywords: Punishment, cooperation, social norms, cognitive modelling, agent-based simulation

Introduction

Theoretical, empirical and ethnographic studies have demonstrated that punishment in human societies promotes and sustains cooperation in large groups of unrelated individuals and more generally plays a crucial role in the maintenance of social order (Fehr & Gachter, 2002; Ostrom, Walker, & Gardner, 1992; Boyd & Richerson, 1992; Boyd, Gintis, & Bowles, 2010; Sigmund, 2007; Herrmann, Thoni, & Gachter, 2008). Although these studies have provided key insights to the understanding of punishment, they have largely looked at this mechanism from the classical economic perspective as a way of changing people's conduct by increasing the costs of undesired behaviour (Becker, 1968). The model of decision making advocated by this perspective is that of the rational actor influenced only by economic incentives.

We claim that this way of considering punishment is incomplete and not likely to mantain large-scale cooperation at least at reasonable costs for the social system. Instead, we argue that punishment is effective in regulating people's behaviour not only through economic incentives, but also for the normative request it asks people (Giardini, Andrighetto, & Conte, 2010; Hirschman, 1984; Xiao & Houser, 2005). In some situations, the punisher informs violators (and the public) that the targeted behaviour is not approved and that it violates a social norm, thus focusing individuals' attention on that norm. We claim that when this happens, cooperation is more stable and the costs for achieving and maintaing it lower then when only economic incentives are used.

Works in psychology suggests that *focusing* people's attention on the norm is a crucial factor in producing norm-compliant behavior. Under this perspective, the normative content elicited by punishment can induce norm compliance, even more - or at least in a more durable way - than the economic incentives imposed by it (R. B. Cialdini, Reno, & Kallgren, 1990; Bicchieri, 2006; R. Cialdini & Goldstein, 2004).

In previous work (Giardini et al., 2010; Andrighetto, Villatoro, & Conte, 2010), we have contributed to the understanding of enforcing strategies by clarifying the cognitive mechanisms undelying them. In particular, we distinguished between revenge, punishment and sanction pointing out the specific mental representations - beliefs, goals, and emotions - characterizing them and the relative ways in which these strategies aim to influence people's conduct. We have used the term punishment to refer to the enforcement mechanism aimed at obtaining deterrence only by changing the costs and benefits of a particular situation; while we used sanction to indicate the mechanism aimed at changing people' conduct by informing violators (and the public) that the targeted behaviour is not approved and that it violates a social norm.

Clearly, in real life situations there is often an overlap - even if very slight - between these two mechanisms; analysing punishment and sanction in isolation however allows us to explore the specific contribution of each of them to the achievement and maintenance of cooperation and possibly to design actions aimed to highlight and exploit such contributions.

Recently, researchers have conducted several experiments designed to explore the norm-signalling effect of sanction, analysing what factors might impact the expressive power of this mechanism (Xiao & Houser, 2005; Masclet, 2003; Noussair & Tucker, 2005), but to our knowledge, the work presented here is the first simulation study that focuses specifically on this topic. Simulation experiments allow us to isolate in vitro punishment and sanction, verify their relative effects on cooperation, and perform what-if analyses that allow to address policy design issues.

In particular, in this paper we explore the hypothesis that cooperation is more stable and less costly for society if individuals are enforced by sanctions: this enforcing strategy has the effect of activating people's normative motivation to cooperate, leading to a more durable cooperation than if people are driven only by the instrumental motivation to avoid punishment. More specifically, the norm-signaling component of sanction allows social norms to be activated and to spread more quickly in the population than if it were enforced only by mere punishment. This normative elicitation has the effect of increasing pro-social behaviours and consequently cooperation within the population.

The article is organized as follows: in Section *Punishment vs Sanction*, punishment and sanction will be analyzed and distinguished on the basis of the specific ways in which they work in order to obtain deterrence. In Section *Agent Archi-*

tecture, we present a rich normative agent architecture, which allows agents to be influenced by punishment and sanction and to process the normative information communicated by the latter. Finally, some simulation results aimed to compare the effectiveness of punishment and sanction in the achievement and maintenance of cooperation and their relative costs for the system are presented and discussed. Future work and conclusions follow.

Punishment vs Sanction

As said in the *Introduction*, we distinguish between two different enforcing strategies, punishment and sanction. On the one hand, we refer to punishment as a practice consisting in imposing a cost on the wrongdoer, with the aim of deterring him from future offenses. Deterrence is achieved by modifying the relative costs and benefits of the situation, so that wrongdoing becomes a less attractive option. The effect of punishment is achieved by influencing the *instrumental* mind of the individual, by shaping his material payoffs (Kreps, Milgrom, Roberts, & Wilson, 1982). This approach to punishment is in line with the economic model of crime, also known as the rational choice theory of crime (Becker, 1968), claiming that the deterrent effect of punishment is caused by increasing individuals' expectations about the price of non-compliance. A rational comparison of the expected benefits and costs guides criminal behaviors and this produces a disincentive to engage in criminal activities.

This view of punishment has been criticized by several scholars stating that it considers *citizens just as consumers with unchanging or arbitrarily changing tastes in matters civic as well as commodity-related behavior* ((Hirschman, 1984)). These researchers criticize the idea that human behaviour is influenced only by economic incentives. Moreover, this idea is questioned also by a large set of empirical evidences showing that punishment can increase cooperation also if it is purely *symbolic* and merely expresses social disapproval, without any material consequences for the punished individual (Noussair & Tucker, 2005).

On the other hand, with sanction we indicate the enforcing strategy intentionally aimed at *informing* the target and the public both of the existence and the violation of a social norm (Giardini et al., 2010; Hirschman, 1984; Xiao & Houser, 2005; Andrighetto et al., 2010) [1] and at *asking* them to comply with it in the future.

The sanctioner ideally wants that the sanctioned changes his conduct not just to avoid the penalty but because he recognizes that there is a norm and wants to respect it. Sanction mixes together material and symbolic aspects and it is aimed at changing the future behaviour of an individual by influencing both its *instrumental* and *normative* mind. In order to decide how to behave, the individual will take into consideration not only a mere costs and benefits measure but also the norm.

Often the sanctioner uses scolding to reign in free-riders, or expresses indignation or blame, or simply he mentions that the targeted behaviour violated a norm. Through these actions, he aims to focus people' attention on different normative aspects, such as: (a) the existence and violation of a norm; (b) the high rate of norm surveillance in the social group; (c) the causal link between violation and sanction: "you are being sanctioned because you violated that norm"; (d) the fact that the sanctioner is a norm defender [2].All these normative messages have a key effect in producing norm compliance and favouring social control as well.

Works in psychology suggest that the influence of a norm is crucially related to the degree to which individuals' attention is focused on the norm. Even a strong personal commitment to a norm does not predict behaviour if that norm is not activated or focus of attention (Bicchieri, 2006; Xiao & Houser, 2005; R. B. Cialdini et al., 1990). Furthermore, the more these norms are made *salient*, the more they will elicit a normative conduct. Norm salience indicates to an individual how operative and relevant a norm is within a group and a given context (Andrighetto et al., 2010). It is a complex function, depending on several social and individual factors. On the one hand, the actions of others provide information about how important a norm is within that social group, in particular it depends on: (1) the amount of compliance and the cost people are willing to spend to comply (R. B. Cialdini et al., 1990); (2) the surveillance rate, the frequency and intensity of punishment (Haley, 2003) and the enforcement typology (private or public, 2nd and 3rd party, punishment or sanction, etc.) (Masclet, 2003); (3) the efforts and costs spared to educate the population to a certain norm; (4) the visibility and explicitness of the norm (R. B. Cialdini et al., 1990); (5) the credibility and legitimacy of the normative source (Sacks, Levi, & Tyler, In Press). On the other hand, norm salience is also affected by the individual sphere, it depends on the degree of entrenchment with beliefs, goals, values and previously internalized norms of the agent (Deci & R.M., 2000).

We claim that both punishment and sanction favor the increment of cooperation in social systems, but sanction achieves cooperation in a more stable way and at a lower cost for the system. Cooperation is expected to be more robust if agents' decisions are driven not only by instrumental considerations but are also based on normative ones. Moreover, an individual that complies with the norm for internal reasons is also more willing to exercise a special form of social control as well, reproaching transgressors and reminding would-be violators that they are doing something wrong. In the following Sections, an agent based simulation aimed to test these hypotheses are presented and some results are discussed.

[1] Clearly, also punishment can have a norm-signallig effect as an unintended by-product, but only the sanctioner intentionally has this norm-defense goal.

[2] Focusing agents' attention on the fact that the sanction is a consequence of a norm violation, and not of a personal damage, has possibly the effect of encouraging the sanctionee and the observers to accept it as an *entitled* act, thus avoiding reiterated aggression.

Simulation model

In order to capture the specific dynamics of punishment and sanction and to test their relative effects in the achievement and maintenance of cooperation a simulation model has been developed. In this model, agents play a variation of the classic Prisoner's Dilemma (PD), where we included an extra stage to the game: after deciding whether to cooperate or not, agents can also choose whether they want (or not) to punish or sanction the opponents who defected.

We assume that agents are located in a social network, which determines a fixed interaction topology [3]. Each timestep of the simulation is structured in 4 phases, that are repeated until convergence is reached (or for a fixed number of timesteps). More specifically, these phases consist in:

1. Partner Selection: Agents are paired with other agents randomly chosen from their neighbors.[4].
2. First Stage: Agents play a PD game, with the following payoffs: $P(C,C) = 3,3; P(C,D) = 0,5; P(D,C) = 5,0; P(D,D) = 1,1$.
3. Second Stage: Agents decide whether to punish/sanction or not the opponents who defected. Only agents who have recognized that there is a norm of cooperation governing their group (see Section *Agent Architecture*) use sanction to enforce others' behaviours; otherwise punishment is used. Punishment works by imposing a cost to the defector, this way affecting its payoffs. On the other hand, sanction also informs the target (and possibly the audience) that the performed action violated a social norm, thus having an impact both on agents' payoffs and on the process of norm recognition and norm salience. [5]
4. Strategy Update: As agents have mixed strategies[6], these strategies are updated on the basis of agents' decisions and of the social information acquired.

In the Section *Decision Making and Strategy Update*, a description of how agents update their decision making is provided.

Agent Architecture

Unlike the vast majority of simulation models in which heterogeneous agents interact according to simple local rules, in our model all the agents are endowed with normative architectures, allowing them: (a) to recognize norms; (b) to generate new normative representations and to act on the basis of them; (c) to influence other agents by direct communication and the use of punishment or sanction. We based our architecture on a simplified version of EMIL-I-A (Andrighetto et al., 2010).

Information	Weight
Self Norm Compliance/Violation	(+/-)0.99
Observed Norm Compliance	(+) 0.33 $\times n$[9]
Non Punished Defectors	(-) 0.66 $\times n$
Punishment Observed/Given/Received	(+) 0.33 $\times n$
Sanction Observed/Given/Received	(+) 0.99 $\times n$
Norm Invocation Listened/Received	(+) 0.99 $\times n$

Table 1: Norm Salience Meter: Cues and Weights

Our normative architecture has two important components: the *norm recognition module* and the *salience meter*. The *norm recognition module* allows agents to interpret a social input as a norm. In order for agents to recognize the existence of a norm, they have to listen by consistent agents at least *two* normative messages, such as "you should not take advantage of your group members by shirking" [7] and observe *ten* normative actions compliant with the norm or aimed to defend it (i.e. cooperation, punishment and sanction, observed or received) [8]. Once these conditions are fulfilled, our agents generate a normative belief that will activate a normative goal (the normative drive) to comply with the norm.

The *salience meter* indicates to the agent how salient a certain norm is. This mesure is updated (interaction after interaction) according to both the personal decisions taken by the agents (individual norm-salience) and the normative information that they infer interacting with their neighbours (social norm-salience).

Each of these cues (see Table 2) are aggregated with different weights, and a higher weight is given to those that are highly related to normative requests [10]. For example, all the behaviors that explicitly mention the norm, such as norm invocations or sanctions, have a stronger impact on norm salience, than actions in which the normative request is not as much explicit, such as punishment. Everytime an agent complies with the norm, the norm salience increases as well. On the other hand, observing non punished/sanctioned defectors makes decrease norm salience.

The resulting salience measure ($salience \in [0-1]$, 0 representing minimum salience and 1 maximum salience) is subjective for each agent thus providing flexibility and adaptability to the system.

This norm salience meter enables the agents to *dynamically* monitor the normative scene and to adapt according to it[11].

[3] Agents can only observe and interact with their direct neighbors.

[4] In certain configurations of topologies, this policy might lead to unpaired agents, but the randomness of the partner selection ensures that all agents interact: those with higher degree will be more likely to interact than those with lower degrees.

[5] If an agent decides not to punish/sanction and it is a norm-holder (i.e. an agent with an highly salient norm of cooperation stored in its mind), it can send an educational message to its opponent.

[6] Differently to a pure strategy, mixed strategies have a probability with which an certain action will be chosen.

[7] An agent is consistent if when choosing to punish, he has before cooperated in the PD.

[8] These values are provisional and need to be fine tuned by experimentation with human-subjects.

[10] These values have been extracted from (R. B. Cialdini et al., 1990)

[11] It is interesting to note that this mechanism allows agents to record the social and normative information, without necessarily proactively exploring the word (e.g. with a trial and error procedure).

For example, in an unstable social environment, if the norm enforcement suddenly decreases, agents having highly salient norms are less inclined to violate them. A highly salient norm is a reason for which an agent continues to comply with it even in the absence of punishment. It guarantees a sort of inertia, making agents less prompt to change their strategy to a more favorable one. Vice versa, if a specific norm decays, our agents are able to detect this change, ceasing to comply with it and adapting to the new state of affairs.

Decision Making and Strategy Update

In this model, agents have two take two decisions at two different stages: to cooperate or defect and to punish/sanction or not, and both of them are probability driven. These decisions are influenced by and aggregation of economic, social and normative considerations. More specifically, the decision to cooperate or to defect is affected by the following drives:

(1) Self-Interested Drive: it motivates agents to maximize their individual utility independently of what the norm asks. The self-interested drive is updated according to (a) the calculation of the marginal reward obtained during the last timestep, and (b) the actual action taken. A proportional and normalized value of the marginal reward obtained indicates how the agent's *cooperation probability* will change. For example, if defecting an agent improved its payoff of three units wrt the last timestep, its probability of cooperating will decrease with an intensity relative to 3 [12].

(2) Social Drive[13]: Agents are influenced by what the majority of their neighbors do.

(3) Normative Drive: once the cooperation norm is recognized, agents decisions are influenced also by the normative drive. The normative drive is affected by the norm salience: the more salient the norm is, the more higher the motivation to cooperate.

The agents who cooperated can decide to punish/sanction defectors. As we said, only agents having recognized that there is a cooperation norm regulating their group can sanction, otherwise they will just use punishment. The punisher and the sanctioner are driven by different motivations. The former punishes in order to induce the future cooperation of others, thus expecting a future pecuniary benefit from its acts ((Kreps et al., 1982)). On the other hand, the sactioner is driven by a normative motivation: he sanctions to favor the generation and spreading of norms within the population. Given these differences, the probability governing the decision of punishing or sanctioning are modified by different factors and they change in the following way:

- *Punishment Drive*: Agents change their tendency to punish on the basis of the relative amount of defectors with respect to the last round. If the number of defectors increased, agents' motivation to punish will decrease accordingly.

- *Sanction Drive*: Agents change their tendency to sanction on the basis of the norm salience. The more salient the norm is, the more higher the probability to sanction defectors.

Experimental Design

In order to analyze the specific effects of punishment and sanction on the achievement of cooperation and their relative costs for maintaining it, we have performed an exhaustive experimental analysis. To reduce the search space (and save computation costs), we have prefixed some parameters (that do not affect the results obtained)[14]. In all the simulations the population is composed by 100 agents, located in a fully connected network[15].

In the following experimental sections, we compare the results obtained in simulations where punishment is used with situations in which sanction is used.

We want to remind the reader that in this work we are not interested in analysing the emergence of norms, therefore some agents already enowed with the cooperation norm are initially loaded into the simulation: we refer to them as norm's holders[16].

Emergence of Cooperation

In the first experiment, we pay attention to the relative effects of punishment and sanction on the achievement of cooperation. In Fig. 1, the different levels of cooperation obtained by using punishment and sanction are shown. The x-axis represents the timesteps of the simulation, the y-axis the level of cooperation achieved and the z-axis the initial amount of norm holders.

In Fig. 1, it is possible to observe that different damages (i.e. the amount of punishment/sanction imposed to the target) affect the cooperation level differently. As expected, with a damage of 5 agents' motivation to defect decreases in a much stronger way that when a lower damage of 3 is imposed to the defectors [17]. A damage of 3 is sufficient to achieve cooperation only in a population enforced by sanction (and in which there is at least an initial amount of 70 norm's hold-

[12] In case the marginal reward is 0 (this and last timestep reward are the same), agents would change their strategy with an inertial value in the same direction it last changed its probability.

[13] Even though we model this drive at the theoretical level, we have decided not to include it in the actual platform yet in order to have clearer results.

[14] The initial cooperation probability for all agents is $0,8$ and a punishment probability of $0,5$.

[15] Different social networks of interaction would definetly produce different dynamics in the system that at this moment we are not interested in analyzing. We refer the reader to (Villatoro, Sen, & Sabater, 2009).

[16] The amount of norm holders varies in each simulation, and they are specified in each figure.

[17] These values chosen as both 3 and 5 punishment damages turn the cooperative action more attractive in terms of payoff. A damage of 3 produces a *slight* improvement of the cooperative action (Payoff = 3) over the defection (Payoff 5 - 3 = 2). On the other hand, a damage of 5 produces a *stronger* difference between cooperation (Payoff = 3) and defection (Payoff = 0).

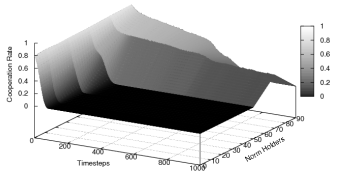

(a) Punishment. Damage 3.

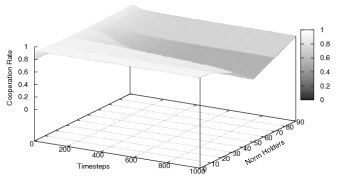

(b) Punishment. Damage 5.

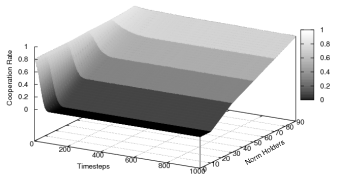

(c) Sanction. Damage 3.

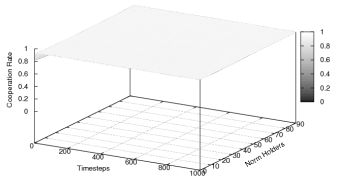

(d) Sanction. Damage 5.

Figure 1: Effects of Punishment and Sanction on the Emergence of Cooperation

ers [18]), while when punishment is used the same damage is too low (see Figures 1 *a* and*c*). As said in Section *Decision Making and Strategy Update*, the agents' probability of cooperating is affected both by the self-interested and the normative drive: sanction - thanks to its signalling component - influences the normative drive more than punishment. In order to obtain deterrence, punishment exploits the power of norms much lesser than sanction, that is why it needs to impose higher damages on its targets.

Relative Costs of Punishment and Sanction

The simulation experiments shown in Fig. 1 provide us also some data on the relative costs of punishment and sanction in the achievement and maintenance of cooperation. In order to obtain the levels of cooperation shown in Fig. 1 (b) and (d), the use of sanctions is 20,93% less costly for the system with respect to punishment. In other words, when using sanction, the amount of sanctioning acts and consequently the associated costs are reduced of 1/5 (see Table 2). This is an interesting result that confirms our idea that sanctioning combines high efficacy in discouraging defectors with lower costs for society as compared to punishment.

[18]In future versions of this work, we will study the right proportion of norm's holders and their correct location in the topology to obtain a stronger effect on cooperation.

	Occurrences	**Global Costs**
Sanction Damage 5	31.221	51.515
Punishment Damage 5	37.757	62.300

Table 2: Punishment and Sanction Occurrences and Relative Costs.

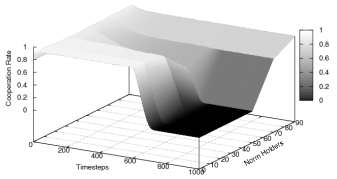

(a) Punishment. Damage 5.

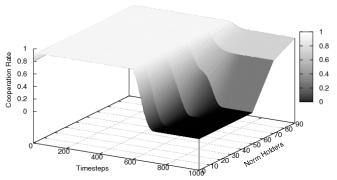

(b) Sanction. Damage 5.

Figure 2: No punishment and sanction after timestep 600

Experiment: What happens when punishing/sanctioning is not possible?

This experiment is aimed to test the hypothesis that sanction makes the population more resilient to change than if it were enforced only by mere punishment.The idea is that if defection turns into an attractive option, for example because it becomes very unlikely that defectors are discovered or even more because there is no social control, we suppose that defectors will take longer to invade again the population in which sanction has been used. In this population a larger amount of cooperation norms have spread, this having a refraining effect on the decision of abandoning the cooperative strategy. To ricreate a situation with no social control, after the timestep 600 of the simulation, we deactivated the possibility to punish/sanction defectors.

Comparing Figures 2 *a* and *b*, we can observe that, when suddenly control stops, agents enforced by sanction will continue to comply with the norm for a longer period compared to agents enforced by punishment. The explanation of this

phenomena is again in the close relationship between sanctions (executed, observed and received) and their impact on norm's salience. Agents having in mind highly salient norms of cooperation cooperate even in absence of deterrent penalties. One of the main advantages of this inertial effect of sanction is that policy makers and system designers can take advantege of this delay in order to restablish the state of the system.

Conclusions and future work

The simulations results presented in this paper clarify the relative ways in which punishment and sanction affect the emergence of cooperation. More specifically, these results verify our hypotheses that the signaling component of sanction allows this mechanism (a) to be more effective in the achievement of cooperation; (b) to make the population more resilient to environmental change than if enforced only by mere punishment; (c) to reduce significantly the costs for cooperation to emerge. Follow-ups of this work will introduce improvements, regarding both theoretical and technical aspects.

First, now that the relationship between the self-interested and the normative drive has been analysed in detail, we are also interested in observing the dynamics introduced by the social drive. Could, in the right conditions, the behavior of peers motivate agent's behavior? Psychological experiments (Asch, 1955) suggest so. Moreover, in order to confront our results with those obtained by experimental economists, we plan to run simulations in which agents play Public Good Games.

Finally, we plan to understand the differences between second and third party punishment. We hypothesize that by allowing agents to evolve their enforcing strategies from second to third party punishment/sanction, social costs will be significantly reduced.

References

Andrighetto, G., Villatoro, D., & Conte, R. (2010). Norm internalization in artificial societies. *AI Communications*.

Asch, S. (1955). Opinions and social pressure. *Scientific American*, *193*(5), 31–35.

Axelrod, R. (1986). An evolutionary approach to norms. *The American Political Science Review*, *4*(80), 1095-1111.

Becker, G. S. (1968). Crime and punishment: An economic approach. *The Journal of Political Economy*, *76*(2), 169–217.

Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. New York: Cambridge University Press.

Boyd, R., Gintis, H., & Bowles, S. (2010, April 30). Coordinated Punishment of Defectors Sustains Cooperation and Can Proliferate When Rare. *Science*, *328*(5978), 617–620.

Boyd, R., & Richerson, P. (1992, May). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, *13*(3), 171–195.

Cialdini, R., & Goldstein, N. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*, 591– 621.

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990, June). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, *58*(6), 1015-1026.

Deci, E. L., & R.M., R. (2000). The "what" and "why" of of goal pursuits: Human needs and the self-determination of behaviour. *Psychological Inquiry*.

Fehr, E., & Gachter, S. (2002). Altruistic punishment in humans. *Nature*, *415*, 137–140.

Giardini, F., Andrighetto, G., & Conte, R. (2010). A cognitive model of punishment. In *Cogsci 2010, annual meeting of the cognitive science society 11-14 august 2010,* (p. -). Portland, Oregon.

Haley, D. F. . K. J. (2003). Genetic and cultural evolution of cooperation. In ed. P. Hammerstein (Ed.), (p. 7-36). Cambridge, MA: The MIT Press.

Herrmann, B., Thoni, C., & Gachter, S. (2008, March 7). Antisocial Punishment Across Societies. *Science*, *319*(5868), 1362–1367.

Hirschman, A. O. (1984, May). Against parsimony: Three easy ways of complicating some categories of economic discourse. *American Economic Review*, *74*(2), 89-96.

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, *27*(2), 245 - 252.

Masclet, D. (2003, June). *L'analyse de l'influence de la pression des pairs dans les quipes de travail* (CIRANO Working Papers No. 2003s-35). CIRANO.

Noussair, C., & Tucker, S. (2005). *Combining monetary and social sanctions to promote cooperation* (Open Access publications from Tilburg University No. urn:nbn:nl:ui:12-377935). Tilburg University.

Ostrom, E., Walker, J., & Gardner, R. (1992). Covenants with and without a sword: Self-governance is possible. *The American Political Science Review*, *86*(2), 404–417.

Sacks, A., Levi, M., & Tyler, T. (In Press). Conceptualizing legitimacy: Measuring legitimating beliefs. *American Behavioral Scientist.*.

Sigmund, K. (2007, November). Punish or perish? Retaliation and collaboration among humans. *Trends in Ecology & Evolution*, *22*(11), 593–600.

Villatoro, D., Sen, S., & Sabater, J. (2009). Topology and memory effect on convention emergence. In *Proceedings of the international conference of intelligent agent technology (iat).* IEEE Press.

Xiao, E., & Houser, D. (2005, May 17). Emotion expression in human punishment behavior. *Proc Natl Acad Sci U S A*, *102*(20), 7398–7401.