

Making an Example: Signalling Threat in the Evolution of Cooperation

Theodor Cimpanu

School of Computing, Engineering & Digital Technologies
Teesside University
Middlesbrough, UK
T.Cimpanu@tees.ac.uk

The Anh Han

School of Computing, Engineering & Digital Technologies
Teesside University
Middlesbrough, UK
T.Han@tees.ac.uk

Abstract—Social punishment has been suggested as a key approach to ensuring high levels of cooperation and norm compliance in one-shot (i.e. non-repeated) interactions, however it does not usually emerge if it is not also cost-efficient. Signalling retribution hearkens back to medieval sovereignty, insofar as the very word for gallows in French stems from the Latin word for *power* and serves as a grim symbol of the ruthlessness of high justice. Here, we introduce the mechanism of signalling an act of punishment and a special type of defector emerges, one who can recognise this signal and avoid punishment by way of fear. We describe the analytical conditions under which threat signalling can maintain high levels of cooperation. Moreover, we perform extensive agent-based simulations so as to confirm and expand our understanding of the external factors that influence the success of social punishment. We show that our suggested mechanism catalyses cooperation, even when signalling is costly or when punishment would be impractical. We observe the preventive nature of advertising retributive acts and we contend that the resulting social prosperity is a desirable outcome in the contexts of AI and multi-agent systems. To conclude, we argue that fear acts as an effective stimulus to pro-social behaviour.

Index Terms—Emergent behaviour; Modelling for agent-based simulation; Social simulation

I. INTRODUCTION

Punishment has been suggested as one of the most relevant explanations to understanding how selfish individuals self-organise and enforce cooperation or compliance to social norms in various societies [16, 30, 6]. Numerous empirical studies show human proclivity towards punishing unjust behaviour or violations of social norms, often at great cost to their own selves [16, 12]. Although in modern societies, sanctioning systems have been widely implemented in the hopes of enforcing laws, many social norms continue to be upheld by the effects of private sanctions [16]. Moreover, third-party punishment has also been implemented in various online systems, such as virtual agent societies [36] or vendor marketplaces [25], as a method of enhancing pro-social behaviour and norms compliance, by both customers and sellers [25].

Cooperation often emerges under the influence of social punishment (i.e. punishing wrongdoers) [11, 16, 18], but this fails to explain how punishment evolves, especially if it is costly to punish others. Indeed, it has been concluded that punishment is often maladaptive within its respective games

[14, 28, 41] and that punishment can only evolve if it is cost-effective to do so (i.e. the offender suffers much more from retribution compared to the aggrieved party) [14, 28, 41].

Refusing low offers in the ultimatum game (another form of punishment) in the presence of observers, made the wilful (punishers) more likely to receive higher offers in future interactions [15]. The fear of having a low offer refused increased the tendency to present higher offers to obstinate individuals and this may help explain which mechanism(s) allow the promotion of punishment when it is costly to do so. In addition, pre-play signalling has been shown to open new avenues for cooperation to emerge, even when such signals are meaningless [35]. One area, to our knowledge, which has not yet been explored is using signalling [23] to explain the emergence of punishment as a viable strategy in evolutionary games, whereby punishers make their retributive deeds well-known in the population, as a deterrent to malefactors. Threat of punishment has also been indicated as one form of making credible commitments [26, 17], which becomes another reasonable explanation to the dilemma of social punishment.

In this paper, we propose and analyse a novel approach towards explaining the evolutionary advantage of punishers in the context of anonymous interactions [37] (without relying on reputation). We make use of evolutionary game theoretic models [37] (see Section III) to show that signalling acts of punishment can promote the emergence of cooperation in the selfish environment of the one-shot Prisoner's Dilemma (PD) [37]. This game is a popular underlying agent interaction framework for studying self-regarding agents and it is also the most difficult pairwise social dilemma for cooperation to emerge in. Threat of punishment can reduce defection from others without having to punish and we show that social welfare in this regime is much higher than what traditional social punishment models suggest (see Section IV-B). We provide a comprehensive view of the outcomes of external factors, such as cost of signalling or effectiveness of punishment, and we show that expensive signalling can still provide meaningful gains to cooperation when punishing others is costly.

The effect of threat of punishment by costly signalling may provide key insights into policy making in the context of distributed systems or artificial intelligence. Indeed, it has been concluded that increasing the probability of developing super-

This work was supported by the Future of Life Institute (grant RFP2-154).

intelligent agents is incompatible with using safety methods that incur delays or limit performance [5]. What is more, when technological supremacy can be achieved in the short to medium term, the significant advantage gained from underestimating or even ignoring ethical and safety precautions could lead to serious negative consequences [1, 8]. A proposed solution to mitigating this dangerous behaviour is to look towards intrinsic measures of encouraging AI research communities to want to pursue safe, beneficial design methodology [3]. Our results show that threat signalling may serve as one intrinsic factor to prevent catastrophic consequences in that regard.

II. RELATED WORK

Punishment has been a major explanation for the evolution of cooperation in the context of the one-shot interaction [16, 12, 30] (for other explanations, see a survey in [37]). A critical condition for cooperation to be sustainable in evolutionary models [6, 20], as well as observable in lab experiments, requires the punishment to be cost-efficient, i.e. the effect it has on the wrongdoer should be sufficiently large compared the cost issued towards the punisher.

Signalling within and between organisms has been investigated using game theoretic models in areas of biology, economics and philosophy and it has been suggested that certain qualitative aspects are common to many real-world interactions [23]. Furthermore, it has been shown that signalling is a robust mechanism for promoting cooperative action in certain collective quorums [29]. In the presence of meaningless (no pre-defined meaning or behaviour) pre-play signals, cooperation has been shown to emerge as a result of individuals learning to discriminate between different signals and reacting accordingly [35]. Pre-play signalling has also been studied in the context of the evolution of honest signalling, showing that honest signalling only emerges when signalling is costly [7]. To the best of our knowledge, no work so far has studied how signalling theory could explain the prevalence of social punishment by advertising acts of punishment after the fact.

Reputation has been suggested as an approach towards addressing this puzzle [13, 31], where agents' actions consolidate in the eyes of observers some assumption of future behaviour. In this manner, social punishers can benefit indirectly through maintaining a reputation of punishing unjust behaviour [21]. However, the assumption that agents' actions are not anonymous proves unrealistic in many social contexts or application domains [37], i.e. in very large societies or when observation is difficult.

The simple presence of an audience has been shown experimentally to increase human propensity for moralistic punishment, causing an increase in costly punishment as a response to perceived moral violations [24]. Participants did not expect to encounter audience members again and the results hold for anonymous interactions or when the only observer was the experimenter. This suggests that there is at least some type of benefit to increasing the observability of one's willingness to punish, beyond reputation. Participants

were generally not self-aware of the reasons for which they decided to punish, so in the context of self-organised societies [4], this would explain why some individuals act towards the interests of society as a whole, irrespective of their intentions to do so [38]. To this end, we are further motivated to study inherent normative mechanisms that have developed as a result of indirect evolutionary advantages.

Survey data on contribution norms in homogenous and heterogeneous groups has demonstrated that uninvolved individuals hold well defined, yet conflicting normative views on equality, equity and efficiency [34]. That being the case, it has also been shown experimentally that punishment can help groups overcome this collective action problem, through the emergence of strong and stable contribution norms [34]. With regard to self-organised systems, punishment may help self-organising agents come to collective agreements on normative standards for efficiency, equity and equality.

Finally, punishment and sanctioning have been studied extensively in the multi-agent system (MAS) literature [40, 2]. Differently from our work, these studies aim at using the cooperation enforcing power of the mechanism for the purpose of regulating individual and collective behaviour, formalizing different relevant aspects of these mechanisms (such as norms and conventions) in a MAS. Moreover, to the best of our knowledge, no work exists in the literature that analyses how sending costly threat of punishment can improve cooperation. As we show later, this mechanism can significantly enhance cooperation even when punishment is not highly cost efficient.

III. MODELS AND METHODS

We adapt the Prisoner's Dilemma (PD), first by integrating the option of costly punishment as a benchmark and we follow by describing the main model and the different configurations, which we explore using replicator dynamics and simulations. By choosing the most competitive social dilemma [22], we explore the toughest environment for the emergence of cooperation, therefore increasing the relevance of any observed effects.

A. Prisoner's Dilemma (PD)

The one-shot PD is characterised by the following payoff matrix:

$$\begin{array}{c} \text{C} \\ \text{D} \end{array} \begin{array}{cc} \text{C} & \text{D} \\ \left(\begin{array}{cc} R & S \\ T & P \end{array} \right) \end{array}$$

Players experience, in pairs, a cooperation dilemma. In an interaction, individuals can decide whether to cooperate (play C) or defect (play D). Mutual cooperation (mutual defection) yields the reward R (penalty P), whereas unilateral defection provides a defector with the temptation T and the cooperator with the sucker's payoff S ($T > R > P > S$) [37]. The game is considered one-shot, in other words there is no memory of past actions or prior knowledge about the interaction. We note that an act of cooperation, i.e. *playing* C is different from a player adopting the strategy C. For the latter, a player will

always play C and this is likewise true for acts of defection and D players.

B. Social Punishment without Threat

We extend the PD by allowing a special type of C player the option of costly punishment, thereby becoming a punisher (P). After the normal interaction has taken place, a P player chooses to punish those opponents who played D during the interaction. A punishment act consists in paying a cost p to make their opponents incur a penalty q . Contrary to previous work that focuses mostly on efficient punishment [32, 6], we include the case where $p > q$, in order to better understand whether and when highly costly or inefficient punishment can still act as a promoting mechanism of cooperation. The newly defined P strategy always cooperates with C (as well as other P) players and always punishes D players. By including this strategy, we can analyse the evolutionary dynamics of punishment strategies and their viability in the evolution of cooperation. The 3×3 payoff matrix for the strategies P, D and C (for row player), is given by:

$$\begin{array}{c} \text{P} \\ \text{D} \\ \text{C} \end{array} \begin{array}{ccc} \text{P} & \text{D} & \text{C} \\ \left(\begin{array}{ccc} R & S-p & R \\ T-q & P & T \\ R & S & R \end{array} \right) \end{array}$$

C. Social Punishment with Threat

We extend the standard social punishment model by introducing signalling an act of punishment and responding to such signals. Firstly, we consider a new type of punisher (denoted by PT) who, upon punishing a defector, can advertise this act by paying a cost θ , thereby alerting future opponents to their willingness to punish (and to the consequences of defecting against them). As such, a new type of defector arises (denoted by DT), who, upon recognising the signal that denotes the threat of being punished, will react and thus cooperate with the signalling punishers (to avoid punishment). PTs cooperate with each other, whereas DTs defect against each other, in a similar fashion to P against D players. For infinite populations, we can derive the 4×4 payoff matrix for PT, D, DT and C (for row player) as follows:

$$\begin{array}{c} \text{PT} \\ \text{D} \\ \text{DT} \\ \text{C} \end{array} \begin{array}{cccc} \text{PT} & \text{D} & \text{DT} & \text{C} \\ \left(\begin{array}{cccc} R & S-p-\theta & R & R \\ T-q & P & P & T \\ R & P & P & T \\ R & S & S & R \end{array} \right) \end{array}$$

In order to derive the payoff matrix for infinite populations, notice that we can disregard the initial encounter between a PT player and either type of defector. Given some probability dependent on the composition of the population, the PT player can enact a punishment upon a DT player. We explain this interaction in-depth in Section III-D2 and provide average payoffs in the case of finite populations (the above payoff matrix for infinite populations can then be recovered at the

limit of increasing the population size to infinity). As this population is infinitely large, the infinitesimally small initial interaction can be safely omitted for the sake of simplification.

D. Methods

All the analysis and numerical results in this paper are obtained using evolutionary game theoretic methods, using replicator dynamics for infinite populations [22] and agent-based simulation for finite populations [37, 27]. In this setting, the payoff for each agent represents their fitness or social success. Evolutionary dynamics are then shaped by social learning [37, 22], whereby the most successful individuals tend to be imitated more often by others.

1) *Replicator Dynamics*: Replicator dynamics are used to study the growth of each fraction (of strategies) in the population, as a function of their frequency and relative fitness, where the fitness in this case corresponds to their payoffs [37, 22]. If we consider a triple strategy game with PT, D and DT, we denote x_{PT} , x_D and x_{DT} the fraction of each strategy, respectively. Therefore, $x_{PT} + x_D + x_{DT} = 1$. The average payoff (Π) for each strategy reads:

$$\begin{aligned} \Pi_{PT} &= (1 - x_D) * R + x_D * (S - p - \theta), \\ \Pi_D &= x_{PT} * (T - q) + (1 - x_{PT}) * P, \\ \Pi_{DT} &= x_{PT} * R + (1 - x_{PT}) * P. \end{aligned} \quad (1)$$

In order to calculate the relative fitness, we require the average fitness ($\bar{\Pi}$) in the population:

$$\bar{\Pi} = x_P * \Pi_{PT} + x_D * \Pi_D + x_{DT} * \Pi_{DT}. \quad (2)$$

We can then calculate the gradients of selection for each strategy with the replicator equations:

$$\begin{aligned} \dot{x}_{PT} &= x_{PT} * (\Pi_{PT} - \bar{\Pi}), \\ \dot{x}_D &= x_D * (\Pi_D - \bar{\Pi}), \\ \dot{x}_{DT} &= x_{DT} * (\Pi_{DT} - \bar{\Pi}). \end{aligned} \quad (3)$$

According to replicator dynamics, whenever a gradient is positive (i.e. $\dot{x} > 0$), the frequency of that particular strategy grows in the population. We can similarly extract the replicator dynamics for P, D, C and PT, D, DT, C.

2) *Agent-Based Simulations*: For our simulations, we adopt a population size $N = 100$ agents. At the beginning of the game, each agent is randomly assigned a strategy from all the available strategies for that experiment. We assume sequential interactions between players (ordered encounters), which is randomised at runtime. At each time step (generation), each agent plays the PD with every other agent in the population. The fitness for each agent is the sum of their payoffs from each interaction.

Social learning is modelled using the pairwise comparison rule [39], a standard approach in studying evolutionary dynamics in evolutionary game theory, which states that a player A with fitness f_A will imitate another player B with fitness f_B with a probability given by the Fermi function, i.e. $P_{A,B} = (1 + e^{-\beta(f_B - f_A)})^{-1}$, where β represents the intensity of selection, i.e. how strongly the agents value the difference

in fitness between them and their opponents. For $\beta = 0$, we obtain neutral drift (random decisions), whereas large values of β lead to increasingly deterministic imitation. We assume at most one imitation can happen per generation (asymmetric update).

In the absence of exploration or mutations, evolution inevitably leads to monomorphic states. Once such a state has been reached, it cannot be escaped solely through imitation. Thus, we assume that, with a certain mutation probability μ , an agent can randomly choose another strategy to adopt without necessarily imitating an existing agent.

We simulate the evolutionary process for 10^4 generations and average our measurements over the final 10^3 steps for a clear and fair comparison (for example due to cyclic patterns). Furthermore, the results for each combination of parameters are obtained from averaging 500 independent realisations, with the exception of typical run patterns. As we include mutations in this work, it is important to note that no simulation reaches a monomorphic state.

As the scenario we are exploring involves a well-mixed population, we simplify the model using a statistical average of the payoffs for each conditional strategy, as opposed to carrying out the random ordering of interactions at the start of each game. The average payoffs for social punishment without threat remain the same as previously discussed in Section III-D1.

To derive the average payoffs, we consider two distinct sequences of events, for each agent acting out a conditional strategy (PT and DT). Firstly, we consider the case when one PT player encounters a D player at the start of the generation. In this instance, the PT will punish the D player, while also incurring the cost θ for signalling the act of punishment. Each subsequent interaction with DT players will result in a reward for both the players, as DTs will react to the signal and avoid defecting against that PT player. Conversely, we consider how the payoffs change when a PT player encounters a DT player as the first defection against that PT in that generation. As the PTs signal is unbeknownst to the DT, it will defect. In turn, the PT will carry out their act of punishment, causing both players to miss the opportunity of cooperating.

The probability of either sequence happening first is dependent on the composition of the population at each first interaction for PTs. The payoffs for all other strategies remain unaffected. Let n_1 , n_2 and n_3 denote the numbers of PT, D and DT players in the population, respectively. We have $n_1 + n_2 + n_3 = N$. We denote $\Pi_{A,B}$ the payoff received by a player following the strategy A when facing players following strategy B (some payoffs are equivalent e.g. $\Pi_{C,C} = \Pi_{PT,C} = \Pi_{C,PT} = \Pi_{PT,PT} = R$). The average payoffs for PT, D and

DT read:

$$\begin{aligned}\Pi_{PT} &= \frac{1}{N-1} \left((n_1 + n_3 - 1) * \Pi_{C,C} + n_2 * \Pi_{PT,D} + \frac{n_3 * (\Pi_{PT,D} - \Pi_{C,C})}{n_2 + n_3} \right), \\ \Pi_D &= \frac{1}{N-1} \left(n_1 * \Pi_{D,PT} + (n_2 + n_3 - 1) * \Pi_{D,D} \right), \\ \Pi_{DT} &= \frac{1}{N-1} \left(n_1 * \Pi_{C,C} + \frac{n_3 * (\Pi_{D,PT} - \Pi_{C,C})}{n_2 + n_3} + (n_2 + n_3 - 1) * \Pi_{D,D} \right).\end{aligned}\quad (4)$$

Note that at the limit of infinite population size, $N \rightarrow \infty$, $x_P = \frac{n_1}{N}$, $x_D = \frac{n_2}{N}$, $x_{DT} = \frac{n_3}{N}$, we recover the equations in (1) for infinite population sizes.

For the four-strategy game (PT, D, DT, C), we introduce n_4 as the number of C players. Therefore, we have $n_1 + n_2 + n_3 + n_4 = N$. The payoffs for PT, D, DT and C then become:

$$\begin{aligned}\Pi_{PT} &= \frac{1}{N-1} \left((n_1 + n_3 + n_4 - 1) * \Pi_{(C,C)} + n_2 * \Pi_{PT,D} + \frac{n_3 * (\Pi_{PT,D} - \Pi_{C,C})}{n_2 + n_3} \right), \\ \Pi_D &= \frac{1}{N-1} \left(n_1 * \Pi_{D,PT} + (n_2 + n_3 - 1) * \Pi_{D,D} + n_4 * \Pi_{D,C} \right), \\ \Pi_{DT} &= \frac{1}{N-1} \left(n_1 * \Pi_{C,C} + \frac{n_3 * (\Pi_{D,PT} - \Pi_{C,C})}{n_2 + n_3} + (n_2 + n_3 - 1) * \Pi_{D,D} + n_4 * \Pi_{D,C} \right), \\ \Pi_C &= \frac{1}{N-1} \left((n_1 + n_4 - 1) * \Pi_{C,C} + (n_2 + n_3) * \Pi_{C,D} \right).\end{aligned}\quad (5)$$

Note that the payoffs for unconditional strategies are never affected by the ordering of interactions. Analytically, the payoff for punishers who threaten depends on the number of defectors who respond to threatening signals in the population, specifically the ratio between the two different defecting strategies (i.e. $\frac{n_3}{n_2+n_3}$). We contend that this could remain true in practical scenarios. In other words, it is better for signalling punishers when future defectors can discriminate signals precisely, and this indirectly increases the payoff for sensitive defectors, as well. This would suggest that there is a synergistic relationship between signalling punishers and fearful defectors and, to an extent, neither would prevail without the other.

IV. RESULTS

We study the potential of punishment and signalling strategies and their effects on evolutionary dynamics using three aforementioned scenarios: no threat (P, D, C), threat without cooperators (PT, D, DT) and threat with freely available strategies (PT, D, DT, C).

A. Replicator Dynamics for Infinite Populations

Following our analysis on infinite populations, we find that introducing threat signalling introduces a type of beneficial dynamic which fosters cooperation. The relationship between signalling punishers and defectors remains very similar to the one found in standard social punishment models (see Figures




Fig. 1. Phase diagram comparison between standard social punishment (left column) and social punishment with threat (right column) using replicator dynamics. Vertices represent specific strategies, whereas solid and empty dots stand for stable and unstable rest points, respectively. The colours represent speed of motion under the dynamic (lighter is faster, darker is slower). Parameters: $T = 2$; $R = 1$; $P = 0$; $S = -1$; $p = 1$; $q = 3$; $\theta = 1$.

1a and 1c). Increasing the efficiency ratio of punishment (i.e. $\frac{q}{p}$) is the only way in which social punishment can remain relevant and this can often lead to undesirable consequences. On the other hand, the dynamic created between PTs and DTs naturally promotes cooperation. DTs lose out to PTs, as DT players do not cooperate amongst each other, but they do better than their defecting brethren, by reaping the rewards of their reluctance to defect against PTs. Even when the fraction of cooperators ($x_P + x_D$) becomes very low, the existence of DTs catalyses the conversion towards cooperation (see Figure 1b). By increasing the efficiency in the case of threat signalling, we found that the range of compositions which lead to all defectors is reduced even further. The ratio q/p also favours DTs over Ds, which can provide another avenue towards cooperation that does not exist in the absence of signalling.

In Figure 1d, we show that the results remain robust when we introduce C players, and in fact that the model is more resilient to compositions with high x_C compared to traditional social punishment models, in which C outperforms P (when p is high).

Analytically, we confirm the results from our replicator dynamics analysis by computing the rest points for each system of equations. For P, D, C, we found the following stable rest point on the PC edge which we can see mirrored in Figure 1a (at $x_P = \frac{1}{2}$):

$$\{x_P = \frac{P - S + p}{R - S - T + P + p + q}, x_D = 1 - x_P, x_C = 0\}.$$

Studying the case of PT, D, DT hints at the interesting dynamics seen in Figure 1b. We observe the following edge

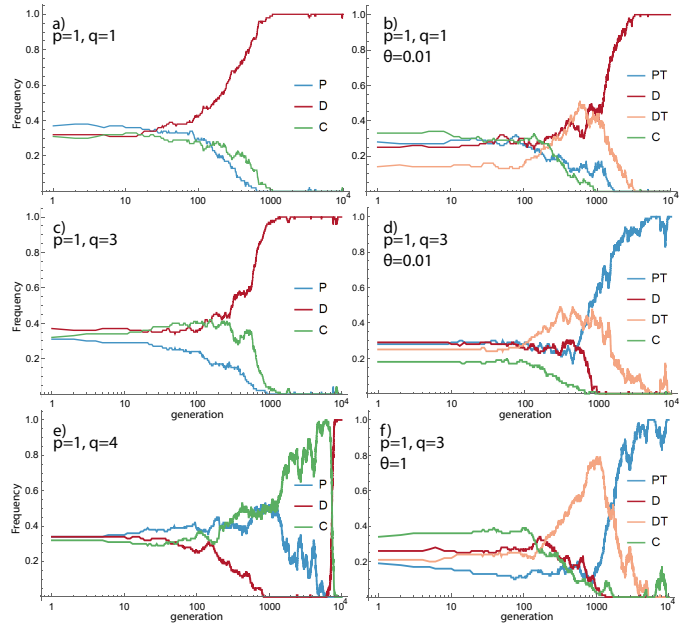


Fig. 2. Typical evolution of frequency over time before (left column) and after (right column) introducing signalling of threat. Note that we only address the typical cases for clear comparisons. Parameters: $T = 2$; $R = 1$; $P = 0$; $S = -1$; $\mu = 0.001$; $\beta = 1$.

rest points:

$$\begin{aligned} \{x_{PT} &= \frac{P - S + p + \theta}{R - S - T + P + p + q + \theta}, \\ x_D &= \frac{R - T + q}{R - S - T + P + p + q + \theta}, x_{DT} = 0\}, \\ \{x_{PT} &= 0, x_D = \frac{R - P}{R - S + p + \theta}, x_{DT} = \frac{P - S + p + \theta}{R - S + p + \theta}\}, \\ \{x_{PT} &= 0, x_{DT} = 1 - x_D\}. \end{aligned}$$

We add that for the parameter values seen in Figure 1, we can observe a slight evolutionary advantage of PT against standard defectors, compared to traditional social punishers. For the same values of p and q , while also having to pay a cost of signalling $\theta = p$, we observe the points $x_P = \frac{1}{2}$ without signalling and $x_{PT} = \frac{3}{5}$ after signalling. This shows that signalling social punishers have higher chances of survival against wrongdoers, even if the aggregated costs they expend towards punishment are larger than the ones paid by standard punishers. Note that we only discuss edge rest points for clear and concise presentation. Vertex rest points do exist in most scenarios, and they can be clearly seen in Figure 1.

B. Agent-Based Simulations for Finite Populations

The initial results from the agent-based simulations confirm the trend we found using replicator equations (see Figure 1). For typical runs, DT outperforms D which leads to the rapid extinction of D players and invites a booming growth of PTs (see Figure 2). When q/p is low enough, we observed some exceptions to the norm, in which the initial conditions led to a population of predominantly DTs, which allowed PTs to flourish which, in turn, led to further defection. This type




Fig. 3. Effect of the efficiency ratio of punishment on frequency of cooperation (punishers and cooperators alike) and average population payoff, before (left column) and after (right column) introducing inexpensive signalling. Parameters: $T = 2$; $R = 1$; $P = 0$; $S = -1$; $\mu = 0.001$; $\beta = 1$; $\theta = 0.01$.

of cyclic behaviour only happened when punishing an act of defection and/or signalling it were extremely costly. We note that the results are robust even for costly signalling.

When signalling is not costly, we found that cooperation is greatly increased across virtually all ranges of p and q . In the case of simple social punishment, punishment is only effective at increasing cooperation when it is also efficient ($\frac{q}{p} > 3$). On the other hand, even when punishment is very inefficient, we find a high propensity for pro-social behaviour in the presence of threat. Even when punishing is extremely costly ($\frac{q}{p} < 1$), we observe close to 50% cooperation. At moderate efficiency ($\frac{q}{p} > 1.5$), the frequency of cooperation is very high ($\approx 70\%$), comparable to highly efficient punishment without threat. We confirm the trends for higher values of θ (see Figures 2f and 4a), but also observe one interesting outcome of signalling that may suggest a direct benefit towards the evolutionary promotion of punishing behaviour. Even at much higher ratios of efficiency of punishment ($\frac{q}{p}$), it is usually cooperators who prevail over the population, as seen in Figure 2e. Conversely, fearful defectors only cooperate with punishers, in the case of signalling, which naturally promotes the evolutionary advantage of punishers (see Figure 2f). One valuable side effect of this phenomenon is that the system is more resilient to mutation and exploration, a large population of punishers is more able to deal with defectors when compared to a large population of cooperators, where a single defector acts akin to a wolf among lambs.

Welfare in the population (i.e. average population payoff) was overall much higher in our extended social punishment model (see Figures 3c, 3d, 4b and 4c). Of note, when punish-




Fig. 4. How costly signalling influences cooperation for (PT, D, DT, C) versus no threat (P, D, C). Figure a) shows the frequency of all cooperative strategies for various values of θ and for the standard case. Panels b) and c) show average payoffs as a function of efficiency of punishment and cost of signalling for standard no threat (left) and threat (right). Parameters: $T = 2$; $R = 1$; $P = 0$; $S = -1$; $\mu = 0.001$; $\beta = 1$; $p = 1$.

ment was very expensive ($p > 2$), average payoff decreases dramatically in the extended model. Intuitively, this happens because PTs survive by cooperating with Cs and DTs, but also incur great costs in order to punish defectors and to deter DTs from following that trend. For higher values of $\frac{q}{p}$, we see another stark difference. With the exception of a very small region of values, welfare decreases ($\bar{\Pi} \approx -25$) as cooperation goes up, whereas the opposite is true for signalling ($\bar{\Pi} \approx 75$). The *power* difference between Ps and Ds proves damaging to social welfare - Ds do not cooperate, which causes Ps to lose fitness and in return, they pay a further cost, compounding the losses, causing Ds to incur even more loss. While this behaviour fosters cooperation, it greatly decreases social welfare. Inversely, a single act of punishment is enough to convert the entire population of DTs to cooperation, which is a qualifying factor in the growth of average payoff.

Our comprehensive study of the external factors under which cooperation emerges, in regards to efficiency of punishment and the cost of signalling shows that fear of punishment enhances cooperation for almost all configurations (with the notable exception of highly efficient punishment coupled with expensive signalling). The results suggest that the transparency of social punishment, specifically the awareness agents have regarding acts of retribution, coupled with the ease of advertising said acts, behaves as a fulcrum towards cooperation. Fear of punishment, therefore, is most effective when awareness of who is or is not a punisher is high. On the other hand, the more deleterious an act of punishment is, the more likely it becomes for standard costly punishment to lead towards

satisfactory outcomes.

We also show that social welfare increases when signalling is not very costly, irrespective of the punishment efficiency. Interestingly, at higher θ values, social welfare lowers even as cooperation increases. We show that when punishment is effective, signalling can lead to lower levels of welfare. Intuitively, this suggests that advertising an ineffective act of punishment is productive even when signalling is expensive.

V. CONCLUSIONS

Punishment used as a deterring mechanism to prevent further damaging actions against the punisher or their peers appears to be a commonly found behaviour in human society and even in some animal hierarchies [11]. Much of recent literature has concluded, however, that punishment may have evolved for reasons other than the promotion of cooperation, because significant benefits to punishers could typically not be found in the context of game theory [14, 33]. Indeed, it may be the case that even if punishing defectors incurs an immediate cost, it discourages observers from repeating said action, as long as the accumulated costs of punishment are outweighed by the additional acts of cooperation evoked over long runs [13]. Our models suggest that this does not only happen in repeated interactions and that punishment can evolve through advertising the acts of punishment.

We show that signalling acts as a catalyst for the emergence of cooperation when defectors are fearful of the punishers who advertise themselves as such. Furthermore, we argue that exhibiting deeds of punishment can explain the success of punishers, when traditional social punishment mechanisms would otherwise fail due to external factors, such as lowly efficient acts of punishment. Indeed, it seems to be the case that fearing punishment can discourage future defectors even more than the evolutionary dynamics associated with inexpensive, deleterious deeds of retribution. Moreover, we show how the traditionally damaging effects of social punishment upon social welfare can be mitigated by way of threat. Because signalling punishers cooperate indiscriminately, they outperform fearful defectors who are always vying for higher status at the expense of others, including themselves.

The prosperity of the population observed under threat of punishment speaks for the preventive nature of advertising acts of justice. Undeniably, it is a beneficial outcome for wicked ventures not to occur in the first place, but contexts such as the development of AI or climate change provide us with unparalleled incentive to prevent potentially disastrous consequences. Given the importance of intrinsic factors that guide the decisions of researchers and policy makers in the field [3], we aim to explore further how the concept of threat, and the self-preservation associated with it, could help guide the current literature on this issue. Additionally, implementing this type of signal response could improve safety conditions in MA systems such as artificial societies [2, 40], especially in cases where the transparency of interactions is reduced.

These observations raise important questions around the coexistence of various types of punishers with different pro-

clivities to signalling, as well as thresholds under which they decide that advertising their deeds of punishment would be appropriate. Reciprocally, defectors could evolve to decide when avoiding punishment is worth the act of justice and which punishers they can exploit even as they signal their propensity towards justice. Perhaps having a loud voice would be more conducive to the emergence of cooperation than the ease with which one can act revenge upon their enemies. Analytically, we have suggested the synergistic mutuality between signalling punishers and fearful defectors. We reason that the high sensitivity of defectors to signals of threat may allow less expensive signalling, whereas lowly responsive defectors require more obvious (and inherently costly) displays.

Future work on this topic will aim to include the full breadth of strategies, including anti-social punishment and the dynamics between normal and signalling punishers. This could be expanded to include memory-decay of threat (such as in [9]), which would require constant reinforcement of the signal for it to remain relevant. We are also interested in how signalling of punishment enables more cost-efficient cooperation-enhancing interference mechanisms [19, 10], reducing costly incentivisation. Assuredly, it would contribute to the breadth of this research to study how network dynamics influence these observations and whether this model for social punishment is robust across a wide range of network topologies and population structures.

REFERENCES

- [1] Stuart Armstrong, Nick Bostrom, and Carl Shulman. Racing to the precipice: a model of artificial intelligence development. *AI & SOCIETY*, 08 2015.
- [2] Tina Balke and Daniel Villatoro. *Operationalization of the Sanctioning Process in Utilitarian Artificial Societies*, pages 167–185. 01 2012.
- [3] Seth Baum. On the promotion of safe and socially beneficial artificial intelligence. *AI & SOCIETY*, 32, 09 2016.
- [4] Jérémy Boes and Frédéric Migeon. Self-organizing multi-agent systems for the control of complex systems. *Journal of Systems and Software*, 134:12 – 28, 2017.
- [5] Nick Bostrom. Strategic implications of openness in ai development. *Global Policy*, 02 2017.
- [6] Robert Boyd, Herbert Gintis, Samuel Bowles, and Peter J. Richerson. The evolution of altruistic punishment. *Proc. Natl. Acad. Sci. USA*, 100(6):3531–3535, 2003.
- [7] David Catteeuw, Bernard Manderick, and T. Anh Han. Evolutionary stability of honest signaling in finite populations. In *2013 IEEE Congress on Evolutionary Computation, CEC 2013*, pages 2864–2870, 2013.
- [8] Stephen Cave and Seán S. ÓhÉigeartaigh. An ai race for strategic advantage: Rhetoric and risks. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 36–40, 2018.
- [9] Raymond Chiong and Michael Kirley. Promotion of cooperation in social dilemma games via generalised indirect reciprocity. *Connection Science*, 27, 09 2015.

- [10] Theodor Cimpanu, The Anh Han, and Francisco C Santos. Exogenous rewards for promoting cooperation in scale-free networks. In *ALIFE 2019*, pages 316–323. MIT Press, 2019.
- [11] T.H. Clutton-Brock and Geoff Parker. Punishment in animal societies. *Nature*, 373:209–16, 02 1995.
- [12] Dominique JF De Quervain, Urs Fischbacher, Valerie Treyer, Melanie Schellhammer, et al. The neural basis of altruistic punishment. *Science*, 305(5688):1254, 2004.
- [13] Miguel dos Santos, Daniel Rankin, and Claus Wedekind. The evolution of punishment through reputation. *Proceedings. Biological sciences / The Royal Society*, 278:371–7, 02 2011.
- [14] Anna Dreber, David G Rand, Drew Fudenberg, and Martin A Nowak. Winners don’t punish. *Nature*, 452(7185):348–351, 2008.
- [15] Ernst Fehr and Urs Fischbacher. The nature of human altruism. *Nature*, 425:785–91, 11 2003.
- [16] Ernst Fehr and Simon Gächter. Altruistic punishment in humans. *Nature*, 415:137–140, 2002.
- [17] T. A. Han, L. M. Pereira, and F. C. Santos. The emergence of commitments and cooperation. In *AAMAS’2012*, pages 559–566, 2012.
- [18] The Anh Han. Emergence of social punishment and cooperation through prior commitments. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI’2016)*, pages 2494–2500, Phoenix, Arizona, USA, 2016.
- [19] The Anh Han and Long Tran-Thanh. Cost-effective external interference for promoting the evolution of cooperation. *Scientific reports*, 8(15997), 2018.
- [20] C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund. Via freedom to coercion: The emergence of costly punishment. *Science*, 316:1905–1907, 2007.
- [21] Christian Hilbe and Arne Traulsen. Emergence of responsible sanctions without second order free riders, antisocial punishment or spite. *Scientific reports*, 2(458), 2012.
- [22] J Hofbauer and K Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [23] Simon Huttegger, Brian Skyrms, Pierre Tarrès, and Elliott Wagner. Some dynamics of signaling games. *Proc Natl Acad Sci U S A.*, 111(3):10873–10880, 2014.
- [24] Robert Kurzban, Peter DeScioli, and Erin M. O’Brien. Audience effects on moralistic punishment. 2007.
- [25] Tomasz Michalak, Joanna Tyrowicz, Peter McBurney, and Michael Wooldridge. Exogenous coalition formation in the e-marketplace based on geographical proximity. *Electronic Commerce Research and Applications*, 8(4):203–223, 2009.
- [26] Randolph M. Nesse. Natural selection and the capacity for subjective commitment. In Randolph M. Nesse, editor, *Evolution and the capacity for commitment*, pages 1–44. New York: Russell Sage, 2001.
- [27] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650, 2004.
- [28] Hisashi Ohtsuki, Yoh Iwasa, and Martin A. Nowak. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature*, 457(7601):79–82, 2009.
- [29] Jorge M. Pacheco, Vítor V. Vasconcelos, Francisco C. Santos, and Brian Skyrms. Co-evolutionary dynamics of collective action with signaling for a quorum. *PLOS Computational Biology*, 11(2):1–12, 02 2015.
- [30] Simon T Powers, Daniel J Taylor, and Joanna J Bryson. Punishment can promote defection in group-structured populations. *Journal of theoretical biology*, 311:107–116, 2012.
- [31] Nichola J Raihani and Redouan Bshary. The reputation of punishers. *Trends in ecology & evolution*, 30(2):98–103, 2015.
- [32] David G Rand, Joseph J Armao IV, Mayuko Nakamaru, and Hisashi Ohtsuki. Anti-social punishment can prevent the co-evolution of punishment and cooperation. *Journal of Theoretical Biology*, 265(4):624–632, 2010.
- [33] Daniel J Rankin, Miguel dos Santos, and Claus Wedekind. The evolutionary significance of costly punishment is still to be demonstrated. *Proc. Natl. Acad. Sci. USA*, 106(50):E135 LP – E135, dec 2009.
- [34] Ernesto Reuben and Arno Riedl. Enforcement of contribution norms in public good games with heterogeneous populations. *Games and Economic Behavior*, 77(1):122 – 137, 2013.
- [35] Francisco C. Santos, Jorge M. Pacheco, and Brian Skyrms. Co-evolution of pre-play signaling and cooperation. *Journal of Theoretical Biology*, 274(1):30 – 35, 2011.
- [36] Bastin Tony Roy Savarimuthu, Maryam Purvis, Martin Purvis, and Stephen Cranefield. Social norm emergence in virtual agent societies. In *Declarative Agent Languages and Technologies VI*, pages 18–28. Springer, 2009.
- [37] Karl Sigmund. *The Calculus of Selfishness*. Princeton University Press, 2010.
- [38] John E. Stewart. Evolutionary possibilities: Can a society be constrained so that “the good” self-organizes? *World Futures*, 74(1):1–35, Oct 2017.
- [39] A. Traulsen, M. A. Nowak, and J. M. Pacheco. Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:11909, 2006.
- [40] Daniel Villatoro, Giulia Andrighetto, Jordi Sabater-Mir, and Rosaria Conte. Dynamic sanctioning for robust and cost-efficient norm compliance. In *IJCAI*, volume 11, pages 414–419, 2011.
- [41] J J Wu, B Y Zhang, Z X Zhou, Q Q He, X D Zheng, R Cressman, and Y Tao. Costly punishment does not always increase cooperation. *Proc Natl Acad Sci U S A.*, 106(41):17448–17451, 2009.