

Towards Understanding the Requirements and Limitations of Reputation-Based Systems

Mohamed Ahmed and Stephen Hailes

University College London, UK

m.ahmed@cs.ucl.ac.uk, s.hailes@cs.ucl.ac.uk

Abstract. Reputation-management, as proposed for dynamic and open systems aims to provide mechanism for analysing the behaviour of agents, and to distribute this information so that the impact of the actions of those acting against the interests of a community can be limited. We study the assumptions that underpin this decision-making role for reputation-management and highlight its limitations with regard to the incentives required to realise the benefits that are claimed for it. Moreover, we show that the benefits claimed for it may not be realisable without enforcing tight constraints on the behaviour and the expectations of agents with respect to the definition of the interaction model and the incentives it presents.

1 Introduction

The motivation behind reputation-based trust management is straightforward: although cryptographic techniques underpin many of the available security solutions for networks, a complete reliance on such approaches to provide system security is inadvisable. In particular, such approaches: (i) have stringent management requirements, that are particularly difficult to meet in massively distributed and open environments (ii) are costly to manage (iii) lack reactivity to the behaviour of agents, relying largely on centralised rather than autonomic control, and (iv) as a consequence, are brittle when faced with uncertainty.

Reputation management aims to provide an endogenous mechanism to mediate the interactions between agents in what are typically assumed to be open and non-cooperative environments. The traditional view of trust is as something policed centrally by authorities that determine whether or not an individual is trustworthy (cf. Equifax etc.). For this to work, several premises must hold: (i) there is widespread trust in such authorities (ii) the penalties that the central authorities can impose are sufficiently severe to discourage bad behaviour (iii) it is not possible for individuals to avoid such penalties easily (for example by changing their identity). However, it is questionable whether these premises are met in existing networks and more doubtful in emerging networking environments. A more distributed approach, to complement that already in existence, necessitates (i) information collection and aggregation services to compensate for the limited local view agents have and (ii) a credible threat to discourage digressions from the standard expectations.

2 Related Work

There is a growing volume of work that seeks to address trust issues that result from open distributed environments. In general, such work aims to augment traditional security mechanisms with support for reactivity by using behavioural analysis techniques and ranges from intrusion detection systems [19], through to reputation-management. In the latter, the aim has been to provide two basic functions: (i) the capacity to learn (on-line) from and adapt to the behaviour of other nodes and (ii) to leverage the experiences of others, so that learning can be bootstrapped and a more accurate estimate of the behaviour of nodes under observation can be built.

Thus, for example, Ganeriwal and Srivastava [14], study the problem of generating reliable information in a sensor network and utilise a Bayesian-inspired approach to attach a subjective probability of (a measure of) integrity to the information produced by nodes. These subjective probabilities are updated by first hand experience and with the recommendations of a community of nodes. In similar fashion [6, 7, 22, 24], examine the related problems of selfish and malicious behaviour in routing of data in wireless sensor networks and mobile ad-hoc networks, and provide mechanisms for differentiating between nodes based on their observed and reported reliability. Lastly, Boukerche and Li [5] study the issues raised for reputation analysis and management when system constraints such as power and bandwidth are taken into account.

These works build on the implicit assumption that reputation-based management is a viable concept for rational agents, and that the problems lie in devising algorithms to compute the necessary trust values from the available information. There is an established body of work in the field of collaborative filtering (see [1]), but the more challenging problem is to ensure that cooperative behaviour in a collection of individual agents is a dominant strategy. For example, Levin [21] shows why simply isolating selfish nodes is insufficient to deter selfish behaviour in the case of multi-hop routing for wireless networks.

In general, with the notable exception of consideration of ephemeral identities [9, 12, 26], work that address the incentive structures or the limitation of reputation based systems [17, 23, 25] is significantly smaller in extent than the number of proposed solutions. Therefore, there is a need to study the circumstances under which the basic premises of information sharing and collective enforcement are viable and realise the dominant strategies.

This paper presents a study of some of the underlying incentive structures that influence the notion of reputation mechanisms, highlighting the need for both filtering and enforcement. Its major contribution is to provide an analysis of some the dynamics and constraints that may influence the decision to use reputation-based mechanisms to reason about the behavioural disposition of agents. We assume that the social norms required to manage and run a reputation-based mechanism are directly influenced by the prospective payoffs they entail, i.e. rational agents will only follow the rules if they are incentivised to do so (see Section 3). Given this, we show how the behavioural and structural constraints of an environment, such as the degree of observability and likelihood of repeat interaction (see Section 4), affect the incentive structure for reputation management and act either to undermine or to reinforce its viability.

3 Trust and Reputation

Following Conte and Paolucci [8], we will refer to the evaluations made on the basis of direct experiences as an *image* and define a *reputation* as “meta-belief” that is propagated between agents and is acquired indirectly. Instead of presenting parametric values that optimise for specific situations, for comparison we use the average discounted payoffs from following community defined social norms. To assess the credibility of the social norms, we look at the enforcement time they require in order to maintain a credible threat of a future loss in utility - relative to the loss incurred.

In the often cited definition of trust, Gambetta [13] asserts that “[*Trust*] (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent will perform a particular action, both before [we] can monitor such action (or independently of our capacity of ever be able to monitor it) and in a context in which it affects [our] own action”. This definition informs us that trust is subjective and dependent on: (i) the competence of the agent to fulfil a task/contract, and (ii) the subjective interpretation of the result by a principal. The competence of an agent to fulfil a task is a complicated issue and may involve numerous external parameters e.g. whether the requisite resources for the task are adequate. The “subjective probabilities” that are assigned as the capacity to fulfil a task/contract depend on the strength of the incentives that drive the process of information exchange and punishment. For example, if the payoff for reneging on a contract/task is equal to the payoff gained for fulfilling it, then, without any exogenous factors, a rational agent would be indifferent to such a contact/task.

The answer to the question of “*why should an agent attempt to build up and maintain a reputation?*” must, therefore, be driven by the fact that to do so is a strategy that leads to a higher payoff than the alternatives. Kreps et al. [20] show that when there is two-sided information asymmetry about the strategy sets available to agents (i.e. when each agent has private information about its strategies), cooperation can form a sequential equilibrium in the finitely repeated prisoners dilemma. For example, by mimicking an irrational agent, e.g. playing a Tit-for-Tat strategy, rather than unconditionally defecting (as would be deduced through a backwards induction analysis), a rational agent is able to a rational principal that it is playing against an irrational agent, and that it is therefore worthwhile for the principal to forego the backwards induction argument that leads to mutual defection until towards the end of the game - in effect reaping the rewards of cooperation at the early stages of the game and a possible windfall from end-stage defection.

The *convincing* process here is the development of an image which conveys a reputation for a given characteristic. For the agent, the building and maintenance of such an image is worthwhile since being mistaken for an irrational agent leads to a higher average discounted payoff. For the principal, monitoring for and factoring a reputation into its expectation is also a worthwhile, since it is able to delay the onset of the backwards deduction argument until towards the end of the game. The significance of this work is that it shows the value in a reputation with respect to the cost of building and maintaining it.

3.1 The Model

To reason about the merit of a reputation-based mechanism, we situate our discussion in an environment inhabited by N interacting agents. In each time step t , a pair of agents (i, j) , each with the feasible action set $A : \{a_1, \dots, a_n\}$ are matched with a uniform probability $\alpha_{ij} = \frac{1}{N-1}$, to play a symmetric stage game Γ and receive the payoffs $g : A^2 \rightarrow \mathbb{R}^2$. Being a symmetric game, the payoff matrices for the agent and principal are identical, therefore $g(a, a') = g'(a', a)$, whereby $g(a, a')$ denotes the payoffs received when the principal plays a and the agent plays a' .

To reason about punishment, we follow convention and define the mini-max action¹ m , with the payoff $g(m, m') = \min_{a \in A} (\max_{a' \in A} g(a, a'))$ (referred to in shorthand as g_m) and assume that m is either a pure action or an observable mixed action. Since m is not necessarily the best response to m' , for example playing a mini-max strategy in response to an opponents mini-max action may actually leave a principal worse off, we assume that $g_m \leq 0$ [4, 18]. In this way, the mini-max action attains a negative utility and is therefore always costly. We also define $\bar{g} = \max_{a \in A} g(a, a')$ to be the maximum possible payoff from interacting in the stage game.

The expected payoff of a principal i playing the stage game against agent j at period t is given by $v_{(i,t)} = \alpha_{ij} g(a, a')$. If a principal i were to receive a payoff of v_i at each time period t , with a discount rate of $0 < \delta < 1$ to weight the value of future payoffs, i.e. a δ of 0 signifies no expectation of future interactions, the desirability of a given payoff stream starting a time t_s and lasting for t_e periods is given by its average discounted value of: $(1 - \delta) \sum_{t=t_s}^{t_s+t_e} \delta^t \alpha_{ij} v_{(i,t)}$.

Finally, since each play by a principal i realises an expected payoff (v_i) for the principal, we define V to contain the set of feasible and individually rational payoffs, such that $v \in V > g_m$. In essence, the set V is defined to contain values that are strictly greater than g_m , so that there is a clear incentive to abstain from being punished and pursue rewards.

3.2 Reputation and Perfect Information

In this section, we discuss the impact of perfect information on a society of agents that make use of images. Under the condition of perfect information, all members of the society observe the actions of each of its members and any feasible and individually rational outcome ($v \in V | v > g_m$) can be enforced with the application of simple social (collective) norms [4, 10, 18]. In the section that follows, we relate this proposition to the aim of reputation-based systems (under the definition of [8]) and show that the aim of a reputation-based system can be seen as endogenously creating the effect of perfect information.

Theorem 1. (*Folk theorem with perfect information*) *With perfect information, any feasible and rational outcome ($v \in V | v > g_m$), can be supported in the sub-game perfect equilibrium for large δ .*

Proof. Suppose that play operates under a trigger-strategy based social norm, dictating that: *if no one has deviated in the last T rounds, a principal 'i' play the strategy that*

¹ I.e., each agent tries to minimise the maximum payoff that its game partner can receive.

yields the payoff v_i . If, however, anyone deviates, the principal switches to the mini-max strategy that yields g_m . This means that while an agent is cooperative, the whole population is cooperative towards it and each other, i.e. all agents play for the individually rational outcome v_i . However, once an agent defects, all agents must respond by defecting against it and each other (playing for g_m) for $t_p - 1$ periods. In effect, the social norm realises a collective punishment on the whole community for the initial digression of a defector. Under this trigger-strategy social norm, if:

An agent i does not deviate from the social norm, it can expect to attain at the minimum an average discounted payoff of v_i .

An agent i does deviate from the social norm at time $t_d - 1$, and it is punished for the next $t_p - 1$ periods, it can expect to attain at most $\delta^{t_d-1}(1-\delta)\bar{g} + vp_i$. Where $vp_i = (\delta^{t_d} - \delta^{t_d+t_p})g_m + \delta^{t_d+t_p}v_i$, is the continuation payoff of the agent after the defection.

Since by definition, $v_i > 0 \geq g_m$, the first term of vp_i , (the punishment cost of $(\delta^{t_d} - \delta^{t_d+t_p})g_m$) is negative or equal to zero, the social norm is dominant ($v_i > vp_i$) iff:

$$\begin{cases} (\delta^{t_d} - \delta^{t_d+t_p})g_m < 0, \text{ or} \\ (\delta^{t_d} - \delta^{t_d+t_p})g_m = 0, \text{ and } \delta < 1, \text{ or} \\ (\delta^{t_d} - \delta^{t_d+t_p})g_m > 0, \text{ and } (1 - \delta^{t_d+t_p})v_i > (\delta^{t_d} - \delta^{t_d+t_p})g_m \end{cases} \quad (1)$$

The first two conditions of Equation 1 are given by the model assumptions, e.g. $g_m < 0$ and $\delta < 1$, while the last condition simply covers the case in which $g_m > 0$; insisting that if this is the case we require $v_i > g_m$. In effect, in the case where the punishment payoff is positive, the payoff for following the social norm *must* strictly be greater than the punishment payoff² to make the cooperative option attractive. Second, given a high enough discount factor to represent a patient agent, i.e. $\lim_{\delta \rightarrow 1}$, then $\delta^{t_d-1}(1-\delta)\bar{g} \approx 0 < v_i^3$, and the cooperative action dominates since $v_i > \delta^{t_d-1}(1-\delta)\bar{g} + vp_i$.

The trigger-based social norm dictates that the digression of some agent j be punished by the whole community. To be in-line with social norm, all agents are expected to play their mini-max strategies whenever they encounter j . However, a mini-max payoff of $g_m \leq 0$ is not necessarily justifiable for a rational agent to enforce. Therefore, without an explicit incentive to enforce the social norm, it is difficult to argue that agents will apply punishment simply for the sake of maintaining it.

To elicit the nature of the problem in more detail, let us examine it in the most extreme case. Assume that there exists an action $a_r \in A$ with a payoff g_r , such that the payoff from playing a_r has the following property: $v_i > g_r > g_m$ and, $g_r = (0, 0)$ (this is analogous to a *refusal* to take part in an interaction). Since, under this condition, g_r dominates g_m , a rational agent will always prefer to receive a payoff of g_r rather than the lower g_m . The result of this assumption is that there is no credibility in the threat of punishment. A deviating agent could expect to receive a payoff of $\delta^{t_d-1}(1-\delta)\bar{g} +$

² If $v_i = g_m$, then an agent is indifferent to the punishment.

³ This is true at both extremes of δ , if $\lim_{\delta \rightarrow 0}$, then, $\delta^{t_d-1}(1-\delta)\bar{g} \approx 0 < v_i$ - provided that $t_d > 2$, which in effect implies a preference of the agent for a defection that happens in the future, relative to the stable payoff of v_i .

$(\delta^{t_d} - \delta^{t_d+t_p})g_m + \delta^{t_d+t_p}v_i$. Knowing the dominance of g_r ($g_r > g_m$) enables an agent to reason that since all other rational agents prefer g_r to g_m , if it deviates, it can receive:

- At best $\delta^{t_d-1}(1 - \delta)\bar{g} + (\delta^{t_d} - \delta^{t_d+t_p})g_r + \delta^{t_d+t_p}v_i$, where it is punished with refusals to interact for t_p periods - with payoff of g_r in each of those periods, after which the game returns to normal and it is able to reap the continuation payoff of $\delta^{t_d+t_p}v_i$.
- At worst $\delta^{t_d-1}(1 - \delta)\bar{g}$, where it is punished with refusals to interact for the rest of the game.

Since $\bar{g} > v_i$, the best case outcome reaps higher reward than the cooperative agent case, i.e. $\delta^{t_d-1}(1 - \delta)\bar{g} + (\delta^{t_d} - \delta^{t_d+t_p})g_r + \delta^{t_d+t_p}v_i > v_i$. However, though the worst case scenario denies an agent the average discounted continuation payoff of $\delta^{t_d+t_p}v_i$, whether this is damaging depends on how long an agent expects the game to last after it cheats.

In effect, provided that the remaining time in the game is less than $\frac{\bar{g}}{v_i}$, it is always advantageous to defect. Therefore, to create a credible threat of punishment, the act of enforcing (punishing) must itself be enforced. To reintroduce the credibility of the threat of punishment, we must state that the social norm punish digressions both in the interaction and enforcement stages of the game. Therefore, under this setting, if:

An agent i does not deviate from the social norm, and carries out $t_p - 1$ periods of punishment as required, then the cost it incurs for enforcing the social norm is at most $(\delta^{t_s} - \delta^{t_p})g_m$, where t_s is the period of the first punishment. Further, since agents are matched under a uniform probability, the probability of a given agent enforcing all of t_p periods of punishment is $(\frac{1}{N-1})^{t_p}$, therefore the expected cost incurred from enforcing t_p periods of punishment is: $(\frac{1}{N-1})^{t_p}(\delta^{t_s} - \delta^{t_p})g_m + [(1 - (\frac{1}{N-1})^{t_p})(\delta^{t_s} - \delta^{t_p})g_m]$.

An agent i does deviate from the social norm, it will in turn be punished for t_p periods.

For one period of deviation at period t_s (with a payoff of g_r), the cost incurred for the deviation is at the very least: $\delta^{t_s}(1 - \delta)g_r + (\delta^{t_s} - \delta^{t_s+t_p})g_m$ ⁴.

Given such payoff profiles, provided that $N > 2$, then the cost of not deviating is strictly less than the cost of deviating since deviating on a punishment phase results in just restarting the process. Therefore, for a rational agent, it is much more beneficial to follow the social norm. In short, under a perfect information system, where the enforcement of social norms is incentivised, the short-term gain from deviating from the norm is outweighed by the long-term loss. \square

The proof in Theorem 1 asserts that the social norm is sub-game perfect, because it is independent of all previous history of play, and no one has an incentive to deviate. If agents deviate during the punishment stage, their action just leads to a restarting of the punishment.

⁴ Though of course this value is dependant on the distribution of agents that may defect on enforcing the punishment, since we cannot make any realistic assumptions about such a distribution, for now we ignore it and assume that agents are willing to enforce.

As discussed, the process of reputation management addresses the requirement on the availability of information by enabling agents to exchange their experiences of interaction. Agents may use these experiences to build up some expectation (represented as a prior) of each other, and to inform some conditional probability mechanism on what is currently known about a prospective interaction partner. To provide credibility, reputation management as often proposed (see [16] for review) makes an implicit and necessary assumption when performing an evaluation based on such a prior, namely, that it can be relied upon to act as a credible threat in deterring deviations from the norm.

If the requirements on information availability and credible threat hold, then the level of cooperation that a given agent can expect can be made directly dependent on the priors that are built about it. For example, agents whose conditional probability for cooperative behaviour leads to a lower than expected payoff, could expect to receive a relatively lower level of cooperation than would otherwise be the case.

4 Imperfect Information and Repeat Interactions

In a distributed computing environment, it is very difficult to sustain the assumptions that underpin the proof in Theorem 1, namely the capacity of all agents to monitor all interactions in the environment. Within economic literature, the problem of imperfect information is tackled by using a type tagging mechanism [4, 10, 12, 18]. The principle is simple: if we can provide a local information processing mechanism that is honest and tamper proof -in the form of a tag, then agents can update the tags after interactions and base their decisions on the content of tag alone. However, the requirements for such a mechanism match those of a PKI and are therefore unrealistic in an open distributed system.

In the presence of imperfect information, if the opportunity for repeat interaction exists, we can still achieve a long run cooperative equilibrium without resorting to trusted third parties or an exogenous tagging scheme. The simplest non tag-based schemes are founded on bilateral trust (reciprocity) [3] and rely on repeated interactions to build histories of the outcomes of their interactions with peers. However, in distributed and open network environments, bilateral trust mechanisms face two shortcomings; (i) there may be a limited scope for repeat interaction, and hence the information accumulated in a history may not be a good approximation of the real behaviour of an agent; (ii) as we shall see, reciprocity alone may not always provide a credible deterrent.

The requirement for information sharing and collective enforcement may be characterised by looking at the standard single shot prisoners' dilemma, in which there is no information sharing and the utility preference ordering favours the non-cooperative outcome i.e. given the actions *cooperate* or *defect*, $g(d, c) > g(c, c) > g(d, d) > g(c, d)$. Under these limiting assumptions, it is still possible to return to a cooperative equilibrium if we can incentivise the future payoffs of an agent in a game in terms its current actions. Indeed, this is precisely the aim of an image sharing reputation system.

From the perspective of maintaining play on the equilibrium path, a limited scope for repeat interaction introduces a change in the structure of a game. Unlike the case for just imperfect information where the aim is to make up for the information deficit,

to maintain cooperative behaviour as the dominant social norm (via a credible threat of punishment), a principal has to now also be concerned with the actions of others, i.e. whether other principals in the environment will punish on its behalf. If agents are able arbitrarily to defect on the punishment stage, then the threat against a defecting agent is weakened and in a long run game, there are agents that can sustain a payoff greater than their mini-max payoff while still defecting. The awareness of this uncertainty in the effectiveness of punishments re-introduces the backwards deduction argument for rational agents and leads to a breakdown of cooperation.

For this reason, there is a two-fold aim in introducing an aggregate opinion mechanism to support the decision making: (i) to compensate for the lack of complete information by informing principals about the full behavioural scope of their prospective partners and (ii) to support an endogenous enforcement mechanism from the application of collective action; so that there is a credible threat of punishment in the form of a loss in future utility. In both scopes, a game's participants are both the witnesses and enforcers of the game. Therefore, the reliability of the mechanism as whole relies on the ability to persuade the interacting agents into enforcing the social norms of the community - which may at times conflict with their own rational interests [15].

4.1 Incentives for Information Sharing

In the case where there are no universally trusted third parties and information is to be collected and distributed by the peers to whom it pertains, it is required that: (i) agents share the information that they individually collect so as to fill the information gap. (ii) This information is sufficiently trustworthy so that it is capable of influencing the decisions of others. As we shall see in this section, these two issues are not independent. To understand the incentives that guide the information sharing process, we consider the two possible cases under which information may be shared:

4.1.1 Case 1: There Are No Consequences Attached to the Utterances of Witnesses

Under this assumption, when witnesses share information, they are able to make claims with the knowledge that they will not face any penalties or reprisals if the information they provide turns out to be inaccurate. This situation is analogous to cheap-talk games as discussed in detail in [11], where the content of a message does not affect the future payoff of its sender.

The consequence of this assumption is that if all message providers have the same preferences over the actions of the message receiver, i.e. all senders would prefer a receiver to act cooperatively to their messages, then, regardless of what they believe or know, they will only provide information that is good for their cause. For example, suppose we have an environment made up of K types of agents $K : \{k_1, k_2, \dots, k_K\}$, and the feasible sets of actions $A : \{a_1, a_2, \dots\}$ and messages $M : \{m_1, m_2, \dots\}$ that they may exchange. Now suppose that, there exists a pure strategy equilibrium that can be used to identify the type of a message sender based on the message they send. Under this strategy, message senders of type k_i all send message m_i , and, in turn, if a message receiver receives m_i , it will assume that the message originated with an agent of type k_i and will play action a_i in response.

Suppose also that all agents know the pure strategy equilibrium of the game, and strictly prefer that a message receiver play a specific action, say a_j with them, rather than the action ordained by their type. Under this assumption, irrespective of their actual type all agents will strictly prefer to send message m_j knowing that the receiver will assume the sender to be of type k_j and play a_j in response.

Such a uniformity of preferences over the message space (a pooling equilibrium) will lead to: (i) a reduction in the value of the message m_j for agents of type k_j , and, (ii) it will reduce considerably the capacity of message receivers to differentiate between message sender types based on the messages they receive alone. Though the existence of a pooling equilibrium reduces the applicability of cost-free messaging, a principal can still *learn* to differentiate between message sender types through repeat interaction.

To analyse the merit of learning to differentiate between message senders, assume that the learning process is on-line and is divided into two phases: (i) an initial set of interactions that may be treated as the training set, that an agent uses to learn with, followed by (ii) a latter exploitation phase, in which an agent exploits the knowledge it has learnt.

For clarity, take the simplest case in which an agent's behaviour remains constant over time, i.e., the initial training examples are able fully to describe a problem space that does not change. Assume that (i) the learning phase lasts for t_l periods and a principal incurs an average cost of vl per-period during this phase and (ii) the exploitation phase lasts for t_e periods and a principal attains an average payoff of ve per-period during this phase. Given this case, a principal i 's average discounted payoff $v_i = (1 - \delta^{t_l})vl_i + (\delta^{t_l} - \delta^{t_l+t_e})ve_i$, and as $t_e \rightarrow \infty^5$, we have $\lim_{t_e \rightarrow \infty}(v_i) \approx (1 - \delta^{t_l})vl_i + \delta^{t_l}ve_i$

Proposition 1. *For learning to be worthwhile, a principal requires that its learning costs to be at least redeemed in its continuation payoff. More formally, it would require that: $(1 - \delta^{t_l})vl_i \leq (\delta^{t_l} - \delta^{t_l+t_e})ve_i$, thereby placing a constraint on the length of time that could be spent learning, or conversely the length of time the continuation period must last for the learning phase to be worthwhile. With regard to the tuples $\{t_l, vl\}$ and $\{t_e, ve\}$, such that $vl < 0 < ve$ and $0 < \left(\frac{(1-\delta^{t_l})vl_i}{\delta^{t_l}ve_i}\right) < 1$, we have the following lower bound for the length of the exploitation phase (t_e):*

$$t_e \geq \frac{\log\left(1 - \frac{(1-\delta^{t_l})vl_i}{\delta^{t_l}ve_i}\right)}{\log(\delta)} \quad (2)$$

The inequality in Equation 2 has the property that $\lim_{vl_i \rightarrow 0}(t_e) \geq 0$, therefore a smaller learning cost reduces the minimum required exploitation time for the agent⁶. Likewise, an agent can be seen to be more optimistic as the certainty in future interactions - a higher discount rate ($\lim_{\delta \rightarrow 1}$) makes it more willing to absorb a longer training period (t_l), or more willing to accept a larger learning cost (v_l).

⁵ For example, if the exploitation period lasts until the (unknown) length of a game.

⁶ We may also hypothesise that an agent aware of a smaller learning cost may be more accepting of the risks of engaging in interactions.

The lack of repeat interaction between principals and agents or analogously the uncertainty in the identity of the prospective players, can be made represented by a lower discount rate; since both result in reducing the capacity to correlate the outcomes of previous interactions. This has the effect of increasing the lower bound of the exploitation period and amplifying the problem of learning to trust.

We can see this in another way, if the cost of learning is greater than the continuation payoff, then the numerator in Equation 2 becomes undefined, hence t_e is also undefined. On the other hand, if the continuation payoff from learning is much greater than the cost of learning, then, $\log\left(1 - \frac{(1-\delta^t)vI_i}{\delta^t v e_i}\right) \rightarrow 1$ and t_e is not very sensitive to smaller changes in the discount rate.

4.1.2 Case 2: There Is Some Consequence Attached to the Utterances of Witnesses

Under this assumption, witnesses may be held to account by some mechanism for the utterances they make. For example, if the perceived action of an agent does not correspond with the testimony/recommendation of its witnesses, then the principal may cause some loss in future payoffs for those witnesses.

However, with the assumption of imperfect information, a principal cannot know the full interaction history of a prospective partners and, consequently, the identities of other agents that the partner has interacted with in the past. In this case, when requested to provide testimony for an agent and act as a witness, agents in the environment can plausibly deny any knowledge they may have. Therefore, the option of not reporting the outcome of previous interactions is open to a witness.

For the moment, assume that a principal is able to exact punishment on a witness for a misleading recommendation by again playing the mini-max strategy m and yielding a payoff of $(g_m \leq 0)$ for its self and the witness. Under this scheme, we have three possible outcomes for the payoff a witness attains with regard only to the information it shares:

1. If a witness's recommendation is in agreement with the outcome of an interaction for which it has testified, then the witness's average discounted payoff with regard to the the information it has shared is either:
 - (a) 0: if the witness is neither punished or rewarded for the information it provides and, therefore its overall payoff is independent of its testimonies and dependent only on its first person interactions.
 - (b) $\delta^{t_{rs}}(1 - \delta^{t_r})g_{rw}$: where g_{rw} is some reward that the witness receives for $t_r - 1$ periods starting at period t_{rs} ⁷ for the information that is has provided.
2. If the witness's testimony is judged as misleading and as a result it is punished for $t_p - 1$ periods starting at period t_s , then its average discounted payoff with regard to the its role as a witness is at the minimum $(\delta^{t_s} - \delta^{t_s+t_p})g_m$ ⁸.

⁷ In the case of a single shot reward, received at some period t_{rs} , we may substitute $\delta^{t_{rs}}(1 - \delta^{t_r+1})g_{rw}$ with $\delta^{t_{rs}}g_{rw}$.

⁸ A rational argument for the value of t_s is $t_d + 1$, where t_d is the period that the defection occurs. For example, given length of the game is unknown, a rational agent want to affect the punishment as soon as possible.

3. If again the witness's testimony is judged as misleading, but it is assumed to have colluded in the defection by deliberately providing a misleading recommendation, then for a defection that occurs at period t_d with a punishment phase, starting at period t_s and lasting for $t_p - 1$ periods, the witness's average discounted payoff with regard to the its role is at a maximum $\delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m$.

If after accepting the positive testimony of a witness on behalf of an agent, a principal is defected against by the agent, then to maintain a *credible threat* the principal must be able to exact a punishment that is at the very least equivalent to its loss in utility.

In this situation, the most limiting assumption that a principal can make with regard to the witness is that it was deliberately misled and the witness was party to a collusive act that resulted in the observed outcome. For example the witness received some part of the utility lost⁹. Under this assumption, given the following global preference ordering of the possible payoffs, $\bar{g} > g_{rw} \geq 0 \geq g_m$:

Proposition 2. *If no rewards are received by the witness for providing information, then to maintain a credible threat, the principal would require that: $0 \geq \delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m$, or simply, the cost of the punishment phase must outweigh the payoff from defecting, $\|(\delta^{t_s} - \delta^{t_s+t_p})g_m\| \geq \delta^{t_d}\bar{g}$, which, from the perspective of length of the punishment periods (t_p) requires:*

$$t_p \geq \frac{\log\left(1 - \frac{\delta^{t_d}\bar{g}}{\delta^{t_s}g_m}\right)}{\log(\delta)} \quad (3)$$

Proposition 3. *If rewards are received for information, the length of the reward period t_r is conditional on the start of the reward period. To ensure that receiving the reward and then not cheating is a dominant strategy, we must maintain the following inequality $\delta^{t_{rs}}(1 - \delta^{t_{rs}+t_r})g_{rw} \geq \delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m$, which from the perspective of the size of the reward payoff gr requires:*

$$g_{rw} \leq \frac{\delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m}{\delta^{t_{rs}}(1 - \delta^{t_{rs}+t_r})} \quad (4)$$

Or, which from the perspective of a punishment, requires that the cost of punishment be less than the gain from receiving a reward, then deviating, i.e. $(\delta^{t_s} - \delta^{t_s+t_p})g_m \leq -\delta^{t_{rs}}(1 - \delta^{t_r})g_{rw} - \delta^{t_d}\bar{g}$ ¹⁰. Therefore the length of the punishment period required to maintain a credible threat is:

$$t_p \geq \frac{\log\left(1 - \frac{-\delta^{t_{rs}}(1 - \delta^{t_r})g_{rw} - \delta^{t_d}\bar{g}}{\delta^{t_s}g_m}\right)}{\log(\delta)} \quad (5)$$

Given the threat of a loss of future payoff, we can see that a witnesses' disposition to share information is directly conditional on how that information will be interpreted by the receiver.

⁹ This particular example ignores the situation that the witness has been misled into providing a misleading testimony.

¹⁰ This is equivalent to stating that absolute value of the punishment cost is greater or equal to the continuation payoff for defecting, i.e. $\|(\delta^{t_s} - \delta^{t_s+t_p})g_m\| \geq \|-\delta^{t_{rs}}(1 - \delta^{t_r})g_{rw} - \delta^{t_d}\bar{g}\|$.

4.2 Incentives for Collective Punishment

So far, we have discussed the incentives for agents to support a cooperative community by contributing to reduce an information deficit. In this section we examine the other half of the problem: their incentives to act to deter what is classified as socially detrimental behaviour by a community of agents, i.e. to punish /enforce social norms on behalf of a community.

To do, so we will look at three related mechanisms to support collective action for the purpose of making the threat of punishment credible. Depending on the degree imperfect information and the likelihood of repeat interaction, we have the following mechanisms available; (i) collective punishment, (ii) personal punishment and (iii) personal enforcement.

4.2.1 Collective Punishment

Collective punishment is the process in which the whole community of agents reacts to a defection and all players are punished during the punishment period. As a social norm, it is typically presented in situations where there is maximal imperfect information and little or no opportunity for repeat interaction. Group sanctioning based social norms aim to reduce the information and state requirements of a system, and look to identify the types of cooperative equilibria that can exist in the most limiting circumstances, i.e. the contagious equilibrium.

These mechanisms are built on “public grim trigger strategies” (PGTS) [10, 18], in which the social norms dictate that agents play for v_i , so long as they have not been defected against, and switch to g_m either forever or for t_p periods (when combined with a public forgiveness mechanism) if they have been defected against. It should be obvious from the mechanism described for PGTS based social norms that, once defection occurs, it starts to spread epidemically.

Now, given the PGTS social norm which prescribes $t_p - 1$ periods of punishment for a defection that occurs in period t_d , and define γ_{ij} to be the probability of an agent i being matched to play a defecting agent j :

- If no agent defects, the average discounted payoff is v_i .
- If no agent has defected so far, and agent i defects, then the average discounted payoff is at most $(1 - \delta^{t_d})\bar{g} + (\delta^{t_d} - \delta^{t_d+t_p})[\gamma_{ij}g_m + (1 - \gamma_{ij})\bar{g}]$.
- Once a defection has occurred, the average discounted payoff during the punishment phase for an agent that chooses to defect is at most: $(1 - \delta^{t_p})[\gamma_{ij}g_m + (1 - \gamma_{ij})\bar{g}]$.
- The average discounted payoff for a cooperative agent during the punishment phase is at least: $(1 - \delta)\underline{g} + (\delta - \delta^{t_p})[\gamma_{ij}g_m + (1 - \gamma_{ij})v_i]$.

For the threat of punishment to be credible, we require the following payoff ordering $a > b > c > d$. This inequality is sustained by the both the diminishing returns offered by discounting future payoffs (δ) and epidemic growth rate of γ_{ij} -the probability of interacting with a defecting agent, once defection starts.

As a realistic and feasible strategy, the PGTS strategy is too fragile, it is indiscriminate in its punishment, too easily susceptible to noise or to deliberate manipulation by a malicious agent seeking to destroy a cooperative equilibrium and, finally, it cannot

always converge back to the cooperative equilibrium¹¹. Personal punishment schemes are proposed to overcome these issues.

4.2.2 Personal Punishment

Personal punishment is the process in which only the defector is punished by the group [4] and is typically presented in situations in which there is some public information, but again little to no opportunity for repeat interaction. Personal punishment schemes aim to overcome the main fragility of a community enforcement mechanism. However, there are a number of constraints attached to such mechanisms.

First given that some form of collective punishment is required to maintain a credible threat against defection when there is a limited opportunity for repeat interactions. To focus only on the incentive issues for enforcement, assume that (i) defections during the interaction phase can be monitored, but defection at the punishment stage cannot, and, (ii) the social norm prescribes that all agents that defect on either on the interaction or punishment stage be punished for t_p rounds.

Given this scenario, if an agent defects against a principal at period t_d , and is punished for $t_p - 1$ periods starting at t_s , then average discounted value of the payoff stream during this phase is $\delta^{t_d} \bar{g} + (\delta^{t_s} - \delta^{t_s+t_p}) g_m + \delta^{t_s+t_p} v_i$, while for cooperating, it expects a payoff of v_i . Since the punishment phase payoff is necessarily less than the payoff in a cooperative phase, to maintain a credible threat of punishment, a rational agent ought strictly to prefer to cooperate rather than to defect.

In this situation, looking at the continuation payoff alone is somewhat misleading. The problem becomes clearer if we look at situations in which there exists an asymmetry in the feasible payoff sets available to agents; for example, if different types of agents attain differing payoffs from the same outcome. According to Bó [4] the threat of personal punishment incentivised by the long term continuation payoff cannot be a credible threat in this situation.

To overcome this problem, we must change our initial assumption by monitoring and/or incentivising the enforcement behaviour of agents, or modify the payoff mechanism. Otherwise, if there is significant imperfect information, and we are led back to the PGTS as the only sub-game perfect equilibrium.

To address this problem, both [4, 12] present mechanisms that work by modifying the payoff landscape. Both schemes introduce short term incentives for enforcers. In the case of Bó [4], these are in the form of “forgiveness” offerings, i.e. $\exists a_f \in A$, such that $g(m, a_f) > g_m \geq g(a_f, m)$. By playing a_f , the agent being punished asks for forgiveness (with the threat of being mini-maxed if it does not) and the principal gains an incentive to enforce the social norm in the form of the payoff from playing $g(m, a_f)$.

In a similar fashion, Friedman and Resnick [12], use a variation of the PGTS named “paying your dues” (PYD), in which instead of the threat of unyielding punishment, agents can instead be indebted to the interaction process through an incentive scheme that rewards cooperative behaviour rather than punishing deviant behaviour.

¹¹ Strictly speaking, this is not the always true, PGTS based mechanism can be tuned to perform to certain conditions, and even die out, however, this is still fragile and difficult to control [2].

Therefore, given a fee of g_f , and an agent's individually rational payoff of v_i , for the fee paying scheme to be a dominant strategy, it is required that rational agent at least be able to recoup it in some 't' future interactions. Therefore given a cost of entry of g_f , iff:

An agent i does not deviate from the norm, u the most desirable outcome is that; after paying the fee the agent proceeds to be cooperative for a game that lasts for $t_e - 1$ periods. In this case, average discounted payoff is at least $g_f + \delta(1 - \delta^{t_e-1})v_i$.

An agent i does deviate from the norm at period t_d - having already paid the entree fee, given a punishment of $t_p - 1$ periods starting at t_s , the average discounted payoff is at least $g_f + \delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m + (\delta^{t_s+t_p} - \delta^{t_s+t_p+t_e})v_i$.

Proposition 4. *To enable the loss of the entry cost to act as a credible threat of punishment, we require that:*

$$g_f + (\delta - \delta^{t_e})v_i \geq g_f + \delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m + (\delta^{t_s+t_p} - \delta^{t_s+t_p+t_e})v_i \quad (6)$$

Necessitating the following constraints; first, given $0 < \delta < 1$ and $v_i > 0$, the lower bound of the length of the exploitation period (t_e) is:

$$t_e \geq \frac{\log\left(\frac{\delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m + \delta^{t_s+t_p}v_i - \delta v_i}{(\delta^{t_s+t_p} - 1)v_i}\right)}{\log(\delta)} \quad (7)$$

Second with regard to the credibility of the punishment payoff, given $0 < \delta < 1$ and $v_i > 0$, we require that:

$$(\delta^{t_s} - \delta^{t_s+t_p})g_m \leq -g_f - \delta^{t_d}\bar{g} - (\delta^{t_s+t_p} - \delta^{t_s+t_p+t_e})v_i \quad (8)$$

Therefore the lower bound on the length of a punishment period (t_p), is:

$$t_p \geq \frac{\log\left(\frac{-g_f - \delta^{t_d}\bar{g} - \delta^{t_s}g_m}{\delta^{t_s}(v_i - \delta^{t_e}v_i - g_m)}\right)}{\log(\delta)} \quad (9)$$

Unless these bounds can be satisfied, no rational agent has an incentive to pay a joining fee of g_f , because a rational agent would expect that either it cannot be policed or exploited relative to the cost of the investment.

Notice that in the PYD scheme, there exists either some universally trusted mechanism to verify the age of identities (the likelihood of repeat interaction under same identity), or individual repeat interaction can be relied upon to satisfy the requirements on t as shown. If there is no mechanism for verifying the age of identities, and a limited opportunity for repeat interaction, then the PYD mechanism loses its incentive appeal. In short, in the personal punishment scheme long-term continuation payoffs cannot be relied upon when there exists asymmetry in the payoff structures, therefore the act of punishing (enforcement) must also be incentivised.

4.2.3 Personal Enforcement

Personal enforcement in which only the agents that have been defected against carry out the punishment. This type of social norm is typically presented in situations in which reciprocity can be applied, i.e regardless of the level of imperfect information, there is significant opportunity for repeat interaction. Under this assumption, reciprocity is feasible, and bilateral trust mechanisms easily apply. Axelrod [3], provides comprehensive work on the feasibility and credibility of reciprocity. However, the assumption of repeat interactions narrows the applicability of reciprocity in this case.

5 Conclusion

We have highlighted the need for both information sharing and collective sanctioning to support reputation-management and discussed the problem of incentivising them. To be incentivised, both of these features require commitments, relative to the costs they incur and seem infeasible if such commitments cannot be fulfilled.

From a design perspective, we may use these results to reason about the applicability and requirements of reputation based systems, especially when the intended deployment environment requires a strong level of assurance, or when trying to understand how to bootstrap such a system. Further, if we can specify the requirements for reputation based systems, for example, the associated payoffs, we have the opportunity to treat the problem from a mechanism design perspective and to engineer more direct solutions that reflect the associated payoff structures, the degree of observability and likelihood of repeat interaction, rather than rely on the more vague incentive mechanism offered by information sharing and classification alone.

We have also shown that there is tangible value in both information of high quality and commitments to collective action; expressed as a factor of the payoffs realised from utilising them. This raises the possibility of using markets to manage both information sharing and collective action, i.e. the development of agents that trade information and act to police an environment for example through applying enforcement on behalf of their clients. We plan to explore these possibilities in future work.

Acknowledgement. The authors would like to thank the EC for funding under the ARAGORN project and the anonymous reviewers for their helpful comments.

References

1. Adomavicius, G., Tuzhilin, E.: Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 734–749 (2005)
2. Ahmed, M., Hailes, S.: Controlling contagious equilibrium. Technical report, Department of Computer Science, University College London (2009)
3. Axelrod, R.: *The Evolution of Cooperation*. Basic Books, New York (1984)
4. Bó, P.: Social norms, cooperation and inequality. *Journal of Economic Theory* 30(1), 89–105 (2007)
5. Boukerche, A., Li, X.: An agent-based trust and reputation management scheme for wireless sensor networks. In: *IEEE GLOBECOM* (2005)

6. Buchegger, S., Le Boudec, J.Y.: The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In: Proceedings of WiOpt 2003: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, Sophia-Antipolis, France (March 2003)
7. Buchegger, S., Le Boudec, J.Y.: A robust reputation system for peer-to-peer and mobile ad-hoc networks. In: Proceedings of P2PEcon 2004, June 2004. Harvard University, Cambridge (2004)
8. Conte, R., Paolucci, M.: Reputation in Artificial Societies: Social Beliefs for Social Order. Kluwer Academic Publishers, Dordrecht (2002)
9. Douceur, J.R.: The sybil attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)
10. Ellison, G.: Cooperation in the Prisoner's Dilemma with Anonymous Random Matching. Review of Economic Studies 61(3), 567–588 (1994)
11. Farrel, J., Rabin, M.: Cheap talk. Journal of Economic Perspectives 10, 103–118 (1996)
12. Friedman, E., Resnick, P.: The social cost of cheap pseudonyms. Journal of Economics and Management Strategy 10(2), 173–199 (2001)
13. Gambetta, D.: Can we trust trust? In: Gambetta, D. (ed.) Trust: Making and Breaking Cooperative Relations, electronic edn., ch. 4, pp. 49–72. Department of Sociology, University of Oxford, 1988/2000 (1988)
14. Ganeriwal, S., Srivastava, M.B.: Reputation-based framework for high integrity sensor networks. In: Proceedings of SASN 2004, Washington, DC, USA (October 2004)
15. Gibbons, R.: Trust in Social Structure: Hobbes and Coase Meet Repeated Games. In: Cook, K. (ed.) Trust in Society. Russel Sage Foundation, Thousand Oaks (2000)
16. Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems 43(2), 618–644 (2007)
17. Jurca, R., Faltings, B.: An incentive compatible reputation mechanism. In: Proceedings of the IEEE Conference on E-Commerce, pp. 285–292 (2003)
18. Kandori, M.: Social Norms and Community Enforcement. Review of Economic Studies 59(1), 63–80 (1992)
19. Kim, J., Bentley, P., Wallenta, C., Ahmed, M., Hailes, S.: Danger is Ubiquitous: Detecting Malicious Activities in Sensor Networks using the Dendritic Cell Algorithm. In: Bersini, H., Carneiro, J. (eds.) ICARIS 2006. LNCS, vol. 4163, pp. 390–403. Springer, Heidelberg (2006)
20. Kreps, D., Milgrom, P., Roberts, J., Wilson, R.: Rational Cooperation in the Finitely Repeated Prisoners Dilemma. Journal of Economic Theory 27, 245–252 (1982)
21. Levin, D.: Punishment in Selfish Wireless Networks: A Game Theoretic Analysis. In: First Workshop on the Economics of Networked Systems (2006)
22. Michiardi, P., Molva, R.: Core: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In: Communication and Multimedia Security (September 2002)
23. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting honest feedback in electronic markets. Technical report, SG Working Paper Series RWP02-039 (2002)
24. Quercia, D., Lad, M., Hailes, S., Capra, L., Bhatti, S.: STRUDEL: Supporting Trust in the Dynamic Establishment of peering coalitions. In: Proceedings of the 21st ACM Symposium on Applied Computing, Dijon, France (April 2006)
25. Luke Teacy, W.T., Patel, J., Jennings, N.R., Luck, M.: Travos: Trust and reputation in the context of inaccurate information sources. Journal of Autonomous Agents and Multi-Agent Systems 12 (2006)
26. Yu, H., Kaminsky, M., Gibbons, P.B., Flaxman, A.: Sybilguard: Defending against sybil attacks via social networks. In: ACM SIGCOMM (2006)

Appendix

Proof of Proposition 1

Given that the learning period and exploitation period are distinct and exploitation starts directly after learning. We may characterise principal i 's average discounted payoff as:

$$(1 - \delta) \left[\sum_{t=0}^{t_l-1} \delta^t v l_i + \sum_{t=t_l}^{t_l+t_e-1} \delta^t v e_i \right] \quad (10)$$

$$= (1 - \delta^{t_l}) v l_i + (\delta^{t_l} - \delta^{t_l+t_e}) v e_i$$

Now, if learning is to be worthwhile for a principal, it is required that $(\delta^{t_l} - \delta^{t_l+t_e}) v e_i \geq (1 - \delta^{t_l}) v l_i$. Given $0 < \delta < 1$, solving Equation 10, the lower bound of t_e is:

$$t_e \geq \frac{\log \left(1 - \frac{(1 - \delta^{t_l}) v l_i}{\delta^{t_l} v e_i} \right)}{\log(\delta)} \quad (11)$$

Proof of Proposition 3

The most limiting assumption that can be made by a principal that is defected against, after it has paid a reward for some information, is that the witness it used has colluded against it and, it has received two sets of rewards; first, for the information it has provided, then from the defection. Under this assumption, the average discounted payoff of the witness is:

$$(1 - \delta) \left[\sum_{t=t_{rs}}^{t_{rs}+t_r-1} \delta^t g_{rw} + \sum_{t=t_d}^{t_d} \delta^t \bar{g} \right] \quad (12)$$

$$= \delta^{t_{rs}} (1 - \delta^{t_r}) g_{rw} + \delta^{t_d} \bar{g}$$

Now, to provide a credible threat of punishment, it is required that the punishment costs remove any gains from the actions; therefore, given $0 < \delta < 1$ and $0 < g_m$, we therefore require that:

$$(\delta^{t_s} - \delta^{t_s+t_p}) g_m \leq -\delta^{t_{rs}} (1 - \delta^{t_r}) g_{rw} - \delta^{t_d} \bar{g} \quad (13)$$

Solving Equation 13, the lower bound of t_p is:

$$t_p \geq \frac{\log \left(1 - \frac{-\delta^{t_{rs}} (1 - \delta^{t_r}) g_{rw} - \delta^{t_d} \bar{g}}{\delta^{t_s} g_m} \right)}{\log(\delta)} \quad (14)$$

Proof of Proposition 4

First, for the cooperative outcome to dominate, we must ensure the continuation payoff of the cooperative outcome dominates the payoff from defecting, and then reaping a continuation payoff after punishment, giving us:

$$g_f + (\delta - \delta^{t_e})v_i \geq g_f + \delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m + (\delta^{t_s+t_p} - \delta^{t_s+t_p+t_e})v_i \quad (15)$$

Since by definition $v_i < \bar{g}$, we must ensure that the difference is made up for by the cooperative continuation payoff, hence, given $0 < \delta < 1$ and $v_i > 0$, from solving Equation 15 the lower bound of t_e is:

$$t_e \geq \frac{\log\left(\frac{\delta^{t_d}\bar{g} + (\delta^{t_s} - \delta^{t_s+t_p})g_m + \delta^{t_s+t_p}v_i - \delta v_i}{(\delta^{t_s+t_p} - 1)v_i}\right)}{\log(\delta)} \quad (16)$$

Second, to function as a credible threat, the punishment phase must result in an average discounted payoff that is less than the cost of entry plus the gain from interactions. If this requirement does not hold, then the continuation payoff after defecting will outweigh the punishment, and will therefore diminish the threat. Therefore, we have:

$$(\delta^{t_s} - \delta^{t_s+t_p})g_m \leq -g_f - \delta^{t_d}\bar{g} - (\delta^{t_s+t_p} - \delta^{t_s+t_p+t_e})v_i \quad (17)$$

Again, given $0 < \delta < 1$, $g_m < 0$ and $\bar{g} > v_i > 0$, the lower bound of t_p from Equation 17 is:

$$t_p \geq \frac{\log\left(\frac{-g_f - \delta^{t_d}\bar{g} - \delta^{t_s}g_m}{\delta^{t_s}(v_i - \delta^{t_e}v_i - g_m)}\right)}{\log(\delta)}. \quad (18)$$