

Designing Social Norm Based Incentive Schemes to Sustain Cooperation in a Large Community

Yu Zhang, Jaeok Park, and Mihaela van der Schaar

Electrical Engineering Department, University of California, Los Angeles
yuzhang@ucla.edu, {jaeok,mihaela}@ee.ucla.edu

Abstract. Sustaining cooperation among self-interested agents is critical for the proliferation of emerging networked communities, such as the communities formed by social networking services. Providing incentives for cooperation in networked communities is particularly challenging because of their unique features: a large population of anonymous agents interacting infrequently, having asymmetric interests, and dynamically joining and leaving the network; network operation errors; and low-cost identity whitewashing. In this paper, taking these features into consideration, we propose a framework for the design and analysis of a class of incentive schemes based on social norms. We first define the concept of sustainable social norm under which no agent has an incentive to deviate. We then formulate the problem of designing an optimal social norm, which selects a social norm that maximizes overall social welfare among all sustainable social norms. Using the proposed framework, we study the structure of optimal social norms and the impacts of punishment lengths and whitewashing on optimal social norms. Our results show that optimal social norms are capable of sustaining cooperation, with the amount of cooperation varying depending on the community characteristics.

Keywords: Incentive Schemes, Networked Communities, Reputation Schemes, Social Norms.

1 Introduction

Recent developments in technology have expanded the boundaries of communities in which individuals interact with each other. However, a large population and the anonymity of individuals in network-based communities make it difficult to sustain cooperative behavior among self-interested individuals, with the so-called free-riding behavior prevailing [1]. Hence, incentive schemes are needed to provide individuals with incentives for cooperation.

The literature has proposed various incentive schemes. The popular forms of incentive devices used in many incentive schemes are payment and differential service. Pricing schemes use payments to reward and punish individuals for their behavior, which in principle can lead self-interested individuals to achieve social optimum by internalizing their external effects (see, for example, [8]). However, it is often claimed that pricing schemes are impractical because they require an accounting infrastructure [2]. Differential service schemes, on the other hand, reward and punish individuals by

providing differential services depending on their behavior. Such incentive schemes are based on the principle of reciprocity and can be classified into personal reciprocation and social reciprocation. In personal reciprocation schemes, individuals can identify each other, and behavior toward an individual is based on their personal experience with the individual. Personal reciprocation is effective in sustaining cooperation in a small community where individuals interact frequently and can identify each other, but it loses its force in a large community where anonymous individuals interact infrequently. In social reciprocation schemes, individuals obtain some information about other individuals (for example, rating) and decide their behavior toward an individual based on their information about that individual. Hence, an individual can be rewarded or punished by other individuals in the community who have not had direct interaction with it. Since social reciprocation requires neither observable identities nor frequent interactions, it has a potential to form a basis of successful incentive schemes for network-based communities. As such, this paper is devoted to the study of incentive schemes based on social reciprocation.

Sustaining cooperation using social reciprocation has been investigated in the economics literature using the framework of anonymous random matching games and social norm [7] and [10]. Each individual is attached a label indicating its reputation, status, etc. which contains information about its past behavior, and individuals with different labels are treated differently by other individuals they interact with. However, [7] and [10] have focused on obtaining the Folk Theorem by characterizing the set of equilibrium payoffs that can be achieved by using a social norm based strategy when the discount factor is sufficiently close to 1. Our work, on the contrary, addresses the problem of designing a social norm given a discount factor and other parameters arising from practical considerations. Specifically, our work takes into account the following features of network-based communities.

- *Asymmetry of interests.* We allow the possibility of asymmetric interests by modelling the interaction between a pair of individuals as a gift-giving game, instead of a prisoner's dilemma game, which assumes mutual interests between a pair of individuals.
- *Report errors.* In a social norm based incentive scheme, it is possible that the reputation of an individual is updated incorrectly because of errors in the report of individuals. Our model incorporates the possibility of report errors, which allows us to analyze its impact on the design and performance, whereas most existing works on reputation schemes [4][6] adopt an idealized assumption that reputations are always updated correctly.
- *Dynamic change in the population.* The members of a community change over time as individuals gain or lose interest in the services provided by community members. We model this feature by having a constant fraction of individuals leave and join the community in every period to study the impact of population turnover on the design and performance.

The remainder of this paper is organized as follows. In Section 2, we describe the repeated matching game and incentive schemes based on a social norm, and then formulate the problem of designing an optimal social norm. In Section 4, we provide analytical results about optimal social norms. We conclude the paper in Section 5.

2 Model

We consider an infinite-horizon discrete time model with a continuum of agents [4]. In a period, each agent generates a service request [9], which is sent to another agent that can provide the requested service. Each agent is equally likely to receive the request from a particular agent, and the matching in each period is independent. In a pair of matched agents, the agent that requests for a service is called a *client* while the agent that receives a service request is called a *server*. The interaction between a pair of matched agents is modelled as a gift-giving game [5], where the server has the binary choice from the set $\mathcal{A} = \{F, D\}$ of whether to fulfil, denoted as F , or decline the request, denoted as D , while the client has no choice. If the server fulfills the client's request, the client receives a service benefit of $b > 0$ while the server suffers a service cost of $c > 0$. If the server declines the request, both agents receive zero payoff. An agent plays the gift-giving game repeatedly with changing partners until it leaves the community. We assume that at the end of each period a fraction $\alpha \in [0, 1]$ of agents in the current population leave and the same amount of new agents join the community. We refer to α as the *turnover rate* [4].

Social welfare in a time period is measured by the average payoff of the agents in that period. As we assume $b > c$, social welfare is maximized when all the servers choose action F in the gift-giving game they play, which yields payoff $b - c$ to every agent. On the contrary, action D is the dominant strategy for the server in the gift-giving game, which can be considered as the myopic equilibrium of the gift-giving game. When every server chooses its action to maximize its current payoff myopically, an inefficient outcome arises where every agent receives zero payoffs.

In order to improve the inefficiency of the myopic equilibrium, we use incentive schemes based on the idea of *social norms*. A social norm is defined as the rules that a group uses to regulate the behavior of members. We consider a social norm that consists of a *reputation scheme* and a *social strategy*, as in [7] and [10]. Formally, a reputation scheme determines the reputations of agents depending on their past actions as a server and is represented by two elements (Θ, τ) . Θ is the set of reputations that an agent can hold, and τ is the reputation update rule. After a server takes an action, the client sends a report about the action of the server to the third-party device or infrastructure that manages the reputations of agents, but the report is subject to errors with a small probability ε . That is, with probability ε , D is reported when the server takes action F , and vice versa. Assuming a binary set of reports, it is without loss of generality to restrict ε in $[0, 1/2]$. We consider a reputation update rule that updates the reputation of a server based only on the reputations of matched agents and the reported action of the server. Then, a reputation update rule can be represented by a mapping $\tau: \Theta \times \Theta \times \mathcal{A} \rightarrow \Theta$, where $\tau(\theta, \tilde{\theta}, a_R)$ is the new reputation for a server with current reputation θ when it is matched with a client with reputation $\tilde{\theta}$ and its action is reported as $a_R \in \mathcal{A}$. A social strategy prescribes the actions that servers should take depending on the

reputations of the matched agents and is represented by a mapping $\sigma : \Theta \times \Theta \rightarrow \mathcal{A}$, where $\sigma(\theta, \tilde{\theta})$ is the approved action for a server with reputation θ that is matched with a client with reputation $\tilde{\theta}$.

To simplify our analysis, we impose the following restrictions on reputation schemes we consider: (i) Θ is a nonempty finite set, i.e., $\Theta = \{0, 1, \dots, L\}$ for some nonnegative integer L ; (ii) τ is defined by

$$\tau(\theta, \tilde{\theta}, a_R) = \begin{cases} \min\{\theta + 1, L\} & \text{if } a_R = \sigma(\theta, \tilde{\theta}), \\ 0 & \text{if } a_R \neq \sigma(\theta, \tilde{\theta}). \end{cases} \tag{1}$$

Note that with the above restrictions a nonnegative integer L completely describes a reputation scheme, and thus a social norm can be represented by a pair $\kappa = (L, \sigma)$. We call the reputation scheme determined by L the *maximum punishment reputation scheme with punishment length L* . In the maximum punishment reputation scheme with punishment length L , there are total $L + 1$ reputations, and the initial reputation for new peers entering the network is given as L . If the reported action of the server is the same as that specified by the social strategy σ , the server's reputation is increased by 1 while not exceeding L . Otherwise, the server's reputation is set as 0.

Below we summarize the sequence of events in a time period.

- Agents generate service requests and are matched.
- Each server observes the reputation of its client and then determines its action.
- Each client reports the action of its server.
- The reputations of agents are updated, and each agent observes its new reputation for the next period.
- A fraction of agents leave the community, and the same amount of new agents join the community.

As time passes, the reputations of agents are updated and agents leave and join the network. Thus, the distribution of reputations evolves over time. In this paper, we use the stationary distribution in our analysis, which will be written $\{\eta_L(\theta)\}$, where $\eta_L(\theta)$ be the fraction of θ -agents in the total population at the beginning of an arbitrary period t and a θ -agent means an agent with reputation θ . $\{\eta_L(\theta)\}$ satisfies the following equality as

$$\begin{aligned} \eta_L(0) &= (1 - \alpha)\varepsilon, \\ \eta_L(\theta) &= (1 - \alpha)(1 - \varepsilon)\eta_L(\theta - 1) \quad \text{for } 1 \leq \theta \leq L - 1, \\ \eta_L(L) &= (1 - \alpha)(1 - \varepsilon)\{\eta_L(L) + \eta_L(L - 1)\} + \alpha. \end{aligned} \tag{2}$$

We now investigate the incentive of agents to follow a prescribed social strategy. Since we consider a non-cooperative scenario, we need to check whether an agent can improve its long-term payoff by a unilateral deviation.

Let $c_\sigma(\theta, \tilde{\theta})$ be the cost suffered by a server with reputation θ that is matched with a client with reputation $\tilde{\theta}$ and follows a social strategy σ . Similarly, let $b_\sigma(\theta, \tilde{\theta})$ be the benefit received by a client with reputation $\tilde{\theta}$ that is matched with a server with reputation θ following a social strategy σ . Since we consider uniform random matching, the expected period payoff of a θ -agent under social norm κ before it is matched is given by

$$v_\kappa(\theta) = \sum_{\tilde{\theta} \in \Theta} \eta_L(\tilde{\theta}) b_\sigma(\tilde{\theta}, \theta) - \sum_{\tilde{\theta} \in \Theta} \eta_L(\tilde{\theta}) c_\sigma(\theta, \tilde{\theta}). \quad (3)$$

To evaluate the long-term payoff of an agent, we use the discounted sum criterion in which the long-term payoff of an agent is given by the expected value of the sum of discounted period payoffs from the current period. Let $p_\kappa(\theta' | \theta)$ be the transition probability that a θ -agent becomes a θ' -agent in the next period under social norm κ , the long-term payoff of an agent from the current period (before it is matched) can be solved by the following recursive equations

$$v_\kappa^\infty(\theta) = v_\kappa(\theta) + \delta \sum_{\theta' \in \Theta} p_\kappa(\theta' | \theta) v_\kappa^\infty(\theta') \quad \text{for } \theta \in \Theta, \quad (4)$$

where $\delta = \beta(1 - \alpha)$ is the weight that an agent puts on its future payoff.

Now suppose that an agent deviates and uses a social strategy σ' under social norm κ . Since the deviation of a single agent does not affect the stationary distribution, the expected period payoff of a deviating θ -agent is given by

$$v_{\kappa, \sigma'}(\theta) = \sum_{\tilde{\theta} \in \Theta} \eta_L(\tilde{\theta}) b_\sigma(\tilde{\theta}, \theta) + \sum_{\tilde{\theta} \in \Theta} \eta_L(\tilde{\theta}) c_{\sigma'}(\theta, \tilde{\theta}). \quad (5)$$

Let $p_{\kappa, \sigma'}(\theta' | \theta, \tilde{\theta})$ be the transition probability that a θ -agent using social strategy σ' becomes a θ' -agent in the next period under social norm κ , when it is matched with a client with reputation $\tilde{\theta}$. The long-term payoff of a deviating agent from the current period (before it is matched) can be computed by solving

$$v_{\kappa, \sigma'}^\infty(\theta) = v_{\kappa, \sigma'}(\theta) + \delta \sum_{\theta' \in \Theta} p_{\kappa, \sigma'}(\theta' | \theta) v_{\kappa, \sigma'}^\infty(\theta') \quad \text{for } \theta \in \Theta. \quad (6)$$

In our model, a server decides whether to provide a service or not after it is matched with a client and observes the reputation of the client. Hence, we check the incentive for a server to follow a social strategy at the point when it knows the reputation of the client. Suppose that a server with reputation θ is matched with a client with reputation $\tilde{\theta}$. When the server follows the social strategy σ prescribed by social norm κ , it receives the long-term payoff $-c_\sigma(\theta, \tilde{\theta}) + \delta \sum_{\theta'} p_\kappa(\theta' | \theta) v_\kappa^\infty(\theta')$, excluding the possible benefit as a client. On the contrary, when the server deviates to a social strategy σ' , it receives the

long-term payoff $-c_{\sigma'}(\theta, \tilde{\theta}) + \delta \sum_{\theta'} p_{\kappa, \sigma'}(\theta' | \theta, \tilde{\theta}) v_{\kappa, \sigma'}^{\infty}(\theta')$, again excluding the possible benefit as a client. By comparing these two payoffs, we can check whether a θ -agent has an incentive to deviate to σ' when it is matched with a client with reputation $\tilde{\theta}$, and define a social norm κ is *sustainable* if no agent can gain from a unilateral deviation regardless of the reputation of the client it is matched with, i.e.

$$\begin{aligned} & -c_{\sigma}(\theta, \tilde{\theta}) + \delta \sum_{\theta'} p_{\kappa}(\theta' | \theta) v_{\kappa}^{\infty}(\theta') \geq \\ & -c_{\sigma'}(\theta, \tilde{\theta}) + \delta \sum_{\theta'} p_{\kappa, \sigma'}(\theta' | \theta, \tilde{\theta}) v_{\kappa, \sigma'}^{\infty}(\theta') \end{aligned} \tag{7}$$

for all σ' , for all $(\theta, \tilde{\theta})$. Thus, under a sustainable social norm, agents follow the prescribed social strategy in their self-interest. Checking whether a social norm is sustainable using the above definition requires computing deviation gains from all possible social strategies, whose computation complexity can be high for moderate values of L . By employing the criterion of unimprovability in Markov decision theory [11], we establish the one-shot deviation principle for sustainable social norms. For notation, let c_a be the cost suffered by a server that takes action a , and let $p_{\kappa, a}(\theta' | \theta, \tilde{\theta})$ be the transition probability that θ -agent becomes a θ' -agent in the next period under social norm κ when it takes action a to a client with reputation $\tilde{\theta}$. The values of $p_{\kappa, a}(\theta' | \theta, \tilde{\theta})$ can be obtained in a similar way to obtain $p_{\kappa, \sigma'}(\theta' | \theta, \tilde{\theta})$, by comparing a with $\sigma(\theta, \tilde{\theta})$.

Lemma 1 (One-shot Deviation Principle). A social norm κ is sustainable if and only if

$$c_{\sigma}(\theta, \tilde{\theta}) - c_a \leq \delta \left[\sum_{\theta'} \left\{ p_{\kappa}(\theta' | \theta) - p_{\kappa, a}(\theta' | \theta, \tilde{\theta}) \right\} v_{\kappa}^{\infty}(\theta') \right] \tag{8}$$

for all $a \in \mathcal{A}$, for all $(\theta, \tilde{\theta})$. ■

Lemma 1 shows that if an agent cannot gain by unilaterally deviating from σ only in the current period and following σ afterwards, it cannot gain by switching to any other social strategy σ' either, and vice versa. First, consider a pair of reputations $(\theta, \tilde{\theta})$ such that $\sigma(\theta, \tilde{\theta}) = F$. If the server with reputation θ serves the client, it suffers the service cost of c in the current period while its reputation in the next period becomes $\min\{\theta + 1, L\}$ with probability $(1 - \varepsilon)$ and 0 with probability ε . Thus, the expected long-term payoff of a θ -agent when it provides a service is given by

$$V_{\theta}(F; F) = -c + \delta \left[(1 - \varepsilon) v_{\kappa}^{\infty}(\min\{\theta + 1, L\}) + \varepsilon v_{\kappa}^{\infty}(0) \right]. \tag{9}$$

Similarly, if a θ -agent deviates and declines the service request, the expected long-term payoff of a θ -agent when it does not provide a service is given by

$$V_\theta(D; F) = \delta[(1 - \varepsilon)v_\kappa^\infty(0) + \varepsilon v_\kappa^\infty(\min\{\theta + 1, L\})]. \quad (10)$$

The incentive constraint that a θ -agent does not gain from a one-shot deviation is given by $V_\theta(F; F) \geq V_\theta(D; F)$, which can be expressed as,

$$\delta(1 - 2\varepsilon)[v_\kappa^\infty(\min\{\theta + 1, L\}) - v_\kappa^\infty(0)] \geq c. \quad (11)$$

Now consider a pair of reputations $(\theta, \tilde{\theta})$ such that $\sigma(\theta, \tilde{\theta}) = D$. Using a similar argument as above, we can show that the incentive constraint that a θ -agent does not gain from a one-shot deviation can be expressed as

$$\delta(1 - 2\varepsilon)[v_\kappa^\infty(\min\{\theta + 1, L\}) - v_\kappa^\infty(0)] \geq -c. \quad (12)$$

Note that (11) implies (12), and thus for θ such that $\sigma(\theta, \tilde{\theta}) = F$ for some $\tilde{\theta}$, we can check only the first incentive constraint (11). Therefore, a social norm κ is sustainable if and only if (11) holds for all θ such that $\sigma(\theta, \tilde{\theta}) = F$ for some $\tilde{\theta}$ and (12) holds for all θ such that $\sigma(\theta, \tilde{\theta}) = D$ for all $\tilde{\theta}$. The left-hand side of the incentive constraints (11) and (12) can be interpreted as the loss from punishment that social norm κ applies to a θ -agent for not following the social strategy. Therefore, in order to induce a θ -agent to provide a service to some clients, the left-hand side should be at least as large as the service cost c , which can be interpreted as the deviation gain. We use $\min_{\theta \in \Theta} \{\delta(1 - 2\varepsilon)[v_\kappa^\infty(\min\{\theta + 1, L\}) - v_\kappa^\infty(0)]\}$ to measure the strength of the *incentive for cooperation* under social norm κ , where cooperation means providing the requested service in our context.

Since we assume that the community operates at the stationary distribution of reputations, social welfare under social norm κ can be computed by

$$U_\kappa = \sum_\theta \eta_L(\theta) v_\kappa(\theta). \quad (13)$$

We assume that the community operator aims to choose a social norm that maximizes social welfare among sustainable social norms. Then the problem of designing a social norm can be formally expressed as

$$\begin{aligned} & \underset{(L, \sigma)}{\text{maximize}} && U_\kappa = \sum_\theta \eta_L(\theta) v_\kappa(\theta) \\ & \text{subject to} && \delta(1 - 2\varepsilon)[v_\kappa^\infty(\min\{\theta + 1, L\}) - v_\kappa^\infty(0)] \geq c, \quad \forall \theta \text{ such that } \exists \tilde{\theta} \text{ such that } \sigma(\theta, \tilde{\theta}) = F, \\ & && \delta(1 - 2\varepsilon)[v_\kappa^\infty(\min\{\theta + 1, L\}) - v_\kappa^\infty(0)] \geq -c, \quad \forall \theta \text{ such that } \sigma(\theta, \tilde{\theta}) = D \quad \forall \tilde{\theta}. \end{aligned} \quad (14)$$

A social norm that solves the problem (14) is called an *optimal social norm*.

3 Analysis of Optimal Social Norms

We first investigate whether there exists a sustainable social norm, i.e., whether the design problem (14) has a feasible solution. Fix the punishment length L and consider a social strategy σ_L^D where agents do not provide a service at all, i.e., for all $(\theta, \tilde{\theta})$. Since there is no service provided in the community when all the agents follow σ_L^D , we have $v_{(L, \sigma_L^D)}^\infty(\theta) = 0$ for all θ . Hence, the relevant incentive constraint (12) is satisfied for all θ , and the social norm (L, σ_L^D) is sustainable. This shows that the design problem (14) always has a feasible solution.

Assuming that an optimal social norm exists, let U^* be the optimal value of the design problem (14). In the following proposition, we study the properties of U^* .

Proposition 1. (i) $0 \leq U^* \leq b - c$; (ii) $U^* = 0$ if $\frac{c}{b} > \frac{\beta(1-\alpha)(1-2\varepsilon)}{1-\beta(1-\alpha)(2-3\varepsilon)}$; (iii) $U^* \geq [1 - (1-\alpha)\varepsilon](b - c)$ if $\frac{c}{b} \leq \beta(1-\alpha)(1-2\varepsilon)$; (iv) $U^* < b - c$ if $\varepsilon > 0$; (v) $U^* = b - c$ if $\varepsilon = 0$ and $\frac{c}{b} \leq \beta(1-\alpha)$; (vi) $U^* = b - c$ only if $\varepsilon = 0$ and $\frac{c}{b} \leq \frac{\beta(1-\alpha)}{1-\beta(1-\alpha)}$.

Proof: See Appendix A. ■

We obtain zero social welfare at myopic equilibrium, without using a social norm. Hence, we are interested in whether we can sustain a social norm in which agents cooperate in a positive proportion of matches. In other words, we look for conditions on the parameters $(b, c, \beta, \alpha, \varepsilon)$ that yield $U^* > 0$.

In order to obtain analytical results, we consider the design problem (14) with a fixed punishment length L , called DP_L . Note that DP_L has a feasible solution, namely σ_L^D , for any L and that there are a finite number (total $2^{(L+1)^2}$) of possible social strategies given L . Therefore, DP_L has an optimal solution for any L . We use U_L^* and σ_L^* to denote the optimal value and the optimal social strategy of DP_L , respectively. We first show that increasing the punishment length cannot decrease the optimal value.

Proposition 2. $U_L^* \geq U_{L'}^*$ for all L and L' such that $L \geq L'$.

Proof: Choose an arbitrary L . To prove the result, we will construct a social strategy σ_{L+1} using punishment length $L+1$ that is feasible and achieves U_L^* . Define σ_{L+1} by

$$\sigma_{L+1}(\theta, \tilde{\theta}) = \begin{cases} \sigma_L^*(\theta, \tilde{\theta}) & \text{for } \theta \leq L \text{ and } \tilde{\theta} \leq L, \\ \sigma_L^*(L, \tilde{\theta}) & \text{for } \theta = L + 1 \text{ and } \tilde{\theta} \leq L, \\ \sigma_L^*(\theta, L) & \text{for } \theta \leq L \text{ and } \tilde{\theta} = L + 1, \\ \sigma_L^*(L, L) & \text{for } \theta = L + 1 \text{ and } \tilde{\theta} = L + 1. \end{cases} \tag{15}$$

Let $\kappa = (L, \sigma_L^*)$ and $\kappa' = (L + 1, \sigma_{L+1})$. From (2), we have $\eta_{L+1}(\theta) = \eta_L(\theta)$ for $\theta = 0, \dots, L - 1$ and $\eta_{L+1}(L) + \eta_{L+1}(L + 1) = \eta_L(L)$. Using this and (3), it is straightforward to see that $v_{\kappa'}(\theta) = v_{\kappa}(\theta)$ for all $\theta = 0, \dots, L$ and $v_{\kappa'}(L + 1) = v_{\kappa}(L)$. Hence, we have that

$$\begin{aligned} U_{\kappa'} &= \sum_{\theta=0}^{L+1} \eta_{L+1}(\theta) v_{\kappa'}(\theta) = \sum_{\theta=0}^{L-1} \eta_{L+1}(\theta) v_{\kappa}(\theta) + \sum_{\theta=L}^{L+1} \eta_{L+1}(\theta) v_{\kappa}(\theta) \\ &= \sum_{\theta=0}^{L-1} \eta_L(\theta) v_{\kappa}(\theta) + \sum_{\theta=L}^{L+1} \eta_{L+1}(\theta) v_{\kappa}(L) \\ &= \sum_{\theta=0}^{L-1} \eta_L(\theta) v_{\kappa}(\theta) + \eta_L(L) v_{\kappa}(L) = U_{\kappa} = U_L^*. \end{aligned} \tag{16}$$

Using (25), we can show that $v_{\kappa'}^{\infty}(\theta) - v_{\kappa'}^{\infty}(0) = v_{\kappa}^{\infty}(\theta) - v_{\kappa}^{\infty}(0)$ for all $\theta = 1, \dots, L$ and $v_{\kappa'}^{\infty}(L + 1) - v_{\kappa'}^{\infty}(0) = v_{\kappa}^{\infty}(L) - v_{\kappa}^{\infty}(0)$. By the definition of σ_{L+1} , the right-hand side of the relevant incentive constraint (i.e., c or $-c$) for each $\theta = 0, \dots, L$ is the same both under σ_L^* and under σ_{L+1} . Also, under σ_{L+1} , the right-hand side of the relevant incentive constraint for $\theta = L + 1$ is the same as that for $\theta = L$. Therefore, σ_{L+1} satisfies the relevant incentive constraints for all $\theta = 0, \dots, L + 1$. ■

Proposition 2 shows that U_L^* is non-decreasing in L . Since $U_L^* \leq b - c$, we have $U^* = \lim_{L \rightarrow \infty} U_L^* = \sup_L U_L^*$. It may be the case that the incentive constraints eventually prevent the optimal value from increasing with L so that the supremum is attained by some finite L . If the supremum is not attained, the protocol designer can set an upper bound on L based on practical consideration. Now we analyze the structure of optimal social strategies given a punishment length.

Proposition 3. Suppose that $\varepsilon > 0$ and $\alpha < 1$. (i) If $\sigma_L^*(0, \hat{\theta}) = F$ for some $\hat{\theta}$, then $\sigma_L^*(0, \tilde{\theta}) = F$ for all $\tilde{\theta} \geq \min\left\{\ln \frac{c}{b} / \ln \beta, L\right\}$; (ii) If $\theta \in \{1, \dots, L - 1\}$ satisfies $\theta \geq L - \left(\ln \frac{c}{b} - \ln Y(\alpha, \varepsilon, L)\right) / \ln \beta$, where

$$Y(\alpha, \varepsilon, L) = \frac{(1-\alpha)^{L+1}(1-\varepsilon)^L \varepsilon - (1-\alpha)^{L+2}(1-\varepsilon)^{L+1} \varepsilon}{(1-\alpha)^{L+1}(1-\varepsilon)^L \varepsilon + \alpha} \quad (17)$$

and $\sigma_L^*(\theta, \hat{\theta}) = F$ for some $\hat{\theta}$, then $\sigma_L^*(\theta, L) = F$; (iii) If $\sigma_L^*(L, \hat{\theta}) = F$ for some $\hat{\theta}$, then $\sigma_L^*(L, L) = F$.

Proof: To facilitate the proof, we define $u_\kappa^\infty(\theta)$ by $u_\kappa^\infty(\theta) = \sum_{l=0}^{\infty} \gamma^l v_\kappa(\min\{\theta+l, L\})$ for $\theta = 0, \dots, L$. Then, by (25), we have $v_\kappa^\infty(\theta) - v_\kappa^\infty(0) = u_\kappa^\infty(\theta) - u_\kappa^\infty(0)$ for all $\theta = 1, \dots, L$.

Suppose that $\sigma_L^*(0, \hat{\theta}) = F$ for some $\hat{\theta}$. Then the relevant incentive constraint for a 0-agent is $\delta(1-2\varepsilon)[u_\kappa^\infty(1) - u_\kappa^\infty(0)] \geq c$. Suppose that $\sigma_L^*(0, \bar{\theta}) = D$ for some $\bar{\theta} \in \{1, \dots, L-1\}$ such that $\bar{\theta} \geq \ln \frac{c}{b} / \ln \beta$. Consider a social strategy σ_L' defined by

$$\sigma_L'(\theta, \tilde{\theta}) = \begin{cases} \sigma_L^*(\theta, \tilde{\theta}) & \text{for } (\theta, \tilde{\theta}) \neq (0, \bar{\theta}), \\ F & \text{for } (\theta, \tilde{\theta}) = (0, \bar{\theta}). \end{cases} \quad (18)$$

Let $\kappa = (L, \sigma_L^*)$ and $\kappa' = (L, \sigma_L')$. Note that $v_{\kappa'}(0) - v_\kappa(0) = -\eta_\tau(\bar{\theta})c < 0$ and $v_{\kappa'}(\bar{\theta}) - v_\kappa(\bar{\theta}) = \eta_\tau(0)b > 0$ since $\varepsilon > 0$ and $\alpha < 1$. Thus, $U_{\kappa'} - U_\kappa = \eta_L(0)\eta_L(\bar{\theta})(b-c) > 0$. Also,

$$u_{\kappa'}^\infty(\theta) - u_\kappa^\infty(\theta) = \begin{cases} [v_{\kappa'}(0) - v_\kappa(0)] + \gamma^{\bar{\theta}} [v_{\kappa'}(\bar{\theta}) - v_\kappa(\bar{\theta})] \\ \quad = (1-\alpha)^{\bar{\theta}+1}(1-\varepsilon)^{\bar{\theta}} \varepsilon [\beta^{\bar{\theta}} b - c] & \text{for } \theta = 0, \\ \gamma^{\bar{\theta}-\theta} [v_{\kappa'}(\bar{\theta}) - v_\kappa(\bar{\theta})] & \text{for } \theta = 1, \dots, \bar{\theta}, \\ 0 & \text{for } \theta = \bar{\theta} + 1, \dots, L. \end{cases} \quad (19)$$

Since $\bar{\theta} \geq \ln \frac{c}{b} / \ln \beta$, we have $u_{\kappa'}^\infty(0) - u_\kappa^\infty(0) \leq 0$. Thus, $u_{\kappa'}^\infty(\theta) - u_\kappa^\infty(0) \geq u_\kappa^\infty(\theta) - u_\kappa^\infty(0)$ for all $\theta = 1, \dots, L$. Hence, σ_L' satisfies the incentive constraints of DP_L , which contradicts the optimality of σ_L^* . This proves that $\sigma_L^*(0, \tilde{\theta}) = F$ for all $\tilde{\theta} \geq \ln \frac{c}{b} / \ln \beta$. Similar approaches can be used to prove $\sigma_L^*(0, L) = F$, (ii), and (iii). \blacksquare

We can represent a social strategy σ_L as an $(L+1) \times (L+1)$ matrix whose (i, j) -entry is given by $\sigma_L(i-1, j-1)$. Proposition 3 provides some structures of an optimal social strategy σ_L^* in the first row and the last column of the matrix representation, but it does not fully characterize the solution of DP_L . Here we aim to obtain the solution of DP_L for $L = 1$ and 2 and analyze how it changes with the parameters. We first begin with the case of two reputations, i.e., $L = 1$. In this case, if $\sigma_1(\theta, \tilde{\theta}) = F$ for some $(\theta, \tilde{\theta})$, the relevant incentive constraint to sustain $\kappa = (1, \sigma_1)$ is $\delta(1 - 2\varepsilon)[v_\kappa^\infty(1) - v_\kappa^\infty(0)] \geq c$. By Proposition 3(i) and (iii), if $\sigma_1^*(\theta, \tilde{\theta}) = F$ for some $(\theta, \tilde{\theta})$, then $\sigma_1^*(0, 1) = \sigma_1^*(1, 1) = F$, provided that $\varepsilon > 0$ and $\alpha < 1$. Hence, among the total of 16 possible social strategies, only four can be optimal social strategies. These four social strategies are

$$\sigma_1^1 = \begin{bmatrix} D & F \\ F & F \end{bmatrix}, \sigma_1^2 = \begin{bmatrix} F & F \\ D & F \end{bmatrix}, \sigma_1^3 = \sigma_1^{D0} = \begin{bmatrix} D & F \\ D & F \end{bmatrix}, \sigma_1^4 = \sigma_1^D = \begin{bmatrix} D & D \\ D & D \end{bmatrix}. \quad (20)$$

The following proposition specifies the optimal social strategy given the parameters.

Proposition 4. Suppose that $0 < (1 - \alpha)\varepsilon < 1/2$. Then

$$\sigma_1^* = \begin{cases} \sigma_1^1 & \text{if } 0 < \frac{c}{b} \leq \frac{\beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon}{1 + \beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon}, \\ \sigma_1^2 & \text{if } \frac{\beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon}{1 + \beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon} < \frac{c}{b} \leq \frac{\beta(1 - \alpha)(1 - 2\varepsilon)[1 - (1 - \alpha)\varepsilon]}{1 - \beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon}, \\ \sigma_1^3 & \text{if } \frac{\beta(1 - \alpha)(1 - 2\varepsilon)[1 - (1 - \alpha)\varepsilon]}{1 - \beta(1 - \alpha)^2(1 - 2\varepsilon)\varepsilon} < \frac{c}{b} \leq \beta(1 - \alpha)(1 - 2\varepsilon), \\ \sigma_1^4 & \text{if } \beta(1 - \alpha)(1 - 2\varepsilon) < \frac{c}{b} < 1. \end{cases} \quad (21)$$

Proof: Let $\kappa^i = (1, \sigma_1^i)$, for $i = 1, 2, 3, 4$. We obtain that

$$\begin{aligned} U_{\kappa^1} &= (1 - \eta_1(0)^2)(b - c), & U_{\kappa^2} &= (1 - \eta_1(0)\eta_1(1))(b - c), \\ U_{\kappa^3} &= (1 - \eta_1(0))(b - c), & U_{\kappa^4} &= 0. \end{aligned} \quad (22)$$

Since $0 < (1 - \alpha)\varepsilon < 1/2$, we have $\eta_1(0) < \eta_1(1)$. Thus, we have $U_{\kappa^1} > U_{\kappa^2} > U_{\kappa^3} > U_{\kappa^4}$. Also, we obtain that

$$\begin{aligned} v_{\kappa^1}^\infty(1) - v_{\kappa^1}^\infty(0) &= \eta_1(0)(b - c), & v_{\kappa^2}^\infty(1) - v_{\kappa^2}^\infty(0) &= b - \eta_1(0)(b - c), \\ v_{\kappa^3}^\infty(1) - v_{\kappa^3}^\infty(0) &= b, & v_{\kappa^4}^\infty(1) - v_{\kappa^4}^\infty(0) &= 0. \end{aligned} \quad (23)$$

Thus, we have $v_{\kappa^3}^\infty(1) - v_{\kappa^3}^\infty(0) > v_{\kappa^2}^\infty(1) - v_{\kappa^2}^\infty(0) > v_{\kappa^1}^\infty(1) - v_{\kappa^1}^\infty(0) > v_{\kappa^4}^\infty(1) - v_{\kappa^4}^\infty(0)$. By choosing the social strategy that yields the highest social welfare among feasible ones, we obtain the result. ■

Proposition 4 shows that the optimal social strategy is determined by the ratio of the service cost and benefit, c/b . When c/b is sufficiently small, the social strategy σ_1^1 can be sustained, yielding the highest social welfare among the four candidate social strategies. As c/b increases, the optimal social strategy changes from σ_1^1 to σ_1^2 to σ_1^3 and eventually to σ_1^4 . Figure 1 shows the optimal social strategies with $L = 1$ as c varies. The parameters we use to obtain the results in the figures of this paper are set as follows unless otherwise stated: $\beta = 0.8$, $\alpha = 0.1$, $\varepsilon = 0.2$, and $b = 10$. Figure 1 (a) plots the incentive for cooperation of the four social strategies. We can find the region of c in which each strategy is sustained by comparing the incentive for cooperation with the service cost c for σ_1^1 , σ_1^2 , and σ_1^3 , and with $-c$ for σ_1^4 . The solid portion of the lines indicates that the strategy is sustained while the dashed portion indicates that the strategy is not sustained. Figure 1 (b) plots the social welfare of the four candidate strategies, with solid and dashed portions having the same meanings. The triangle-marked line represents the optimal value, which takes the maximum of the social welfare of all sustained strategies.

Next, we analyze the case of three reputations, i.e. $L = 2$. In Figure 2, we show the optimal value and the optimal social strategy of DP_2 as we vary c . The optimal social strategy σ_2^* changes in the following order before becoming σ_2^D as c increases:

$$\begin{aligned} \sigma_2^1 &= \begin{bmatrix} F & F & F \\ D & F & F \\ F & F & F \end{bmatrix}, \sigma_2^2 = \begin{bmatrix} D & F & F \\ F & F & F \\ F & F & F \end{bmatrix}, \sigma_2^3 = \begin{bmatrix} D & F & F \\ D & F & F \\ F & F & F \end{bmatrix}, \\ \sigma_2^4 &= \begin{bmatrix} F & F & F \\ F & F & F \\ D & F & F \end{bmatrix}, \sigma_2^5 = \begin{bmatrix} F & F & F \\ D & F & F \\ D & F & F \end{bmatrix}, \sigma_2^6 = \begin{bmatrix} D & F & F \\ F & F & F \\ D & F & F \end{bmatrix}, \sigma_2^7 = \begin{bmatrix} D & F & F \\ D & F & F \\ D & F & F \end{bmatrix}. \end{aligned} \tag{24}$$

Note that $\sigma_2^1 = \sigma_2^B$ for small c and $\sigma_2^7 = \sigma_2^{D0}$ for large c (but not too large to sustain cooperation), which are consistent with the discussion about Proposition 5. For the intermediate values of c , only the elements in the first column change in order to increase the incentive for cooperation. We find that the order of the optimal social strategies between $\sigma_2^1 = \sigma_2^B$ and $\sigma_2^7 = \sigma_2^{D0}$ depends on the community's parameters $(b, \beta, \alpha, \varepsilon)$.

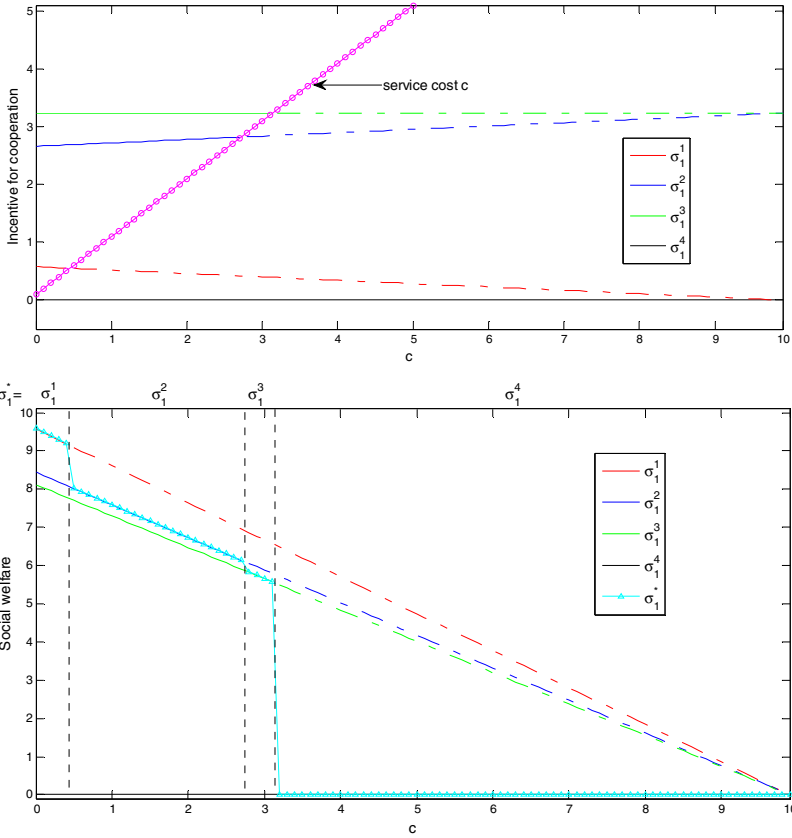


Fig. 1. Performance of the four candidate social strategies when $L = 1$: (a) incentive for cooperation, and (b) social welfare and the optimal social strategy




Fig. 2. Optimal social welfare and the optimal social strategy of DP_2

4 Conclusions

In this paper, we used the idea of “social norm” to establish a rigorous mathematical framework to analyze the incentive mechanisms based on indirect reciprocity in network-based communities. We proposed the idea of “sustainable social norm” in which no agent has the incentive to deliberately deviate. Based on this, we analyzed the optimal social norm problem and quantified the efficiency loss in the social norm as a trade-off to the incentive given to agents to follow it. We also showed that the optimal social strategy has some common properties which are independent from the number of reputations in the social norm. Our current analysis can be extended in different directions, among which we mention two. First, instead of implementing the maximum punishment reputation scheme, more complicated reputation schemes can be used where the agent’s reputation does not fall to 0 directly upon deviation. Second, we grant the agents newly entering the network the highest reputation L in this paper. Nevertheless, this will provide agents with low reputations in the network the incentive to repeatedly leave and rejoin the network in the aim of obtaining a higher reputation and thus better long-term payoff, a phenomenon commonly known as “whitewash”. Hence, more sophisticated rules of reputation assignment can be designed to discourage the “whitewash” behavior.

References

1. Adar, E., Huberman, B.A.: Free Riding on Gnutella. *First Monday* 5(10) (October 2000)
2. Buragohain, C., Agrawal, D., Suri, S.: A Game-theoretic Framework for Incentives in P2P Systems. In: *Proc. Int. Conf. Agent-to-Agent Computing*, pp. 48–56 (September 2003)
3. Cohen, B.: Incentives Building Robustness in BitTorrent. In: *Proc. P2P Econ. Workshop*, Berkeley, CA (2003)
4. Feldman, M., Lai, K., Stoica, I., Chuang, J.: Robust Incentive Techniques for Agent-to-Agent Networks. In: *Proc. of the 5th ACM Conf. on Elec. Commerce, Session 4*, pp. 102–111 (2004)
5. Johnson, P., Levine, D., Pesendorfer, W.: Evolution and Information in a Gift-Giving Game. *J. Econ. Theory* 100(1), 1–21 (2001)
6. Kamvar, S., Schlosser, M.T., Molina, H.G.: The Eigentrust Algorithm for Reputation Management in P2P Networks. In: *Proc. 12th Int’l Conf. on World Wide Web*, pp. 640–651 (2003)
7. Kandori, M.: Social Norms and Community Enforcement. *Rev. Economic Studies* 59(1), 63–80 (1992)
8. MacKie-Mason, J.K., Varian, H.R.: Pricing congestible network resources. *IEEE J. Sel. Areas Commun.* 13(7), 1141–1149 (1995)
9. Massoulie, L., Vojnovic, M.: Coupon Replication Systems. *IEEE/ACM Trans. on Networking* 16(3), 603–616 (2005)
10. Okuno-Fujiwara, M., Postlewaite, A.: Social norms and random matching games. *Games Econ. Behavior* 9(1), 79–109 (1995)
11. Whittle, P.: *Optimization Over Time*. Wiley, New York (1983)

5 Appendix A

Proof of Proposition 1:

(i) $U^* \geq 0$ follows by noting that (L, σ_L^D) is feasible. Note that the objective function can be rewritten as $U_\kappa = (b - c) \sum_{\theta, \tilde{\theta}} \eta_L(\theta) \eta_L(\tilde{\theta}) I(\sigma(\theta, \tilde{\theta}) = F)$, where I is an indicator function. Hence, $U_\kappa \leq b - c$ for all κ , which implies $U^* \leq b - c$.

(ii) By (4), we can express $v_\kappa^\infty(1) - v_\kappa^\infty(0)$ as

$$\begin{aligned} & v_\kappa^\infty(1) - v_\kappa^\infty(0) \\ &= v_\kappa(1) + \delta[(1 - \varepsilon)v_\kappa^\infty(2) + \varepsilon v_\kappa^\infty(0)] - v_\kappa(0) - \delta[(1 - \varepsilon)v_\kappa^\infty(1) + \varepsilon v_\kappa^\infty(0)] \\ &= v_\kappa(1) - v_\kappa(0) + \delta(1 - \varepsilon)[v_\kappa^\infty(2) - v_\kappa^\infty(1)]. \end{aligned} \tag{25}$$

Similarly, we have

$$\begin{aligned} v_\kappa^\infty(2) - v_\kappa^\infty(1) &= v_\kappa(2) - v_\kappa(1) + \delta(1 - \varepsilon)[v_\kappa^\infty(3) - v_\kappa^\infty(2)], \\ & \vdots \\ v_\kappa^\infty(L - 1) - v_\kappa^\infty(L - 2) &= v_\kappa(L - 1) - v_\kappa(L - 2) + \delta(1 - \varepsilon)[v_\kappa^\infty(L) - v_\kappa^\infty(L - 1)], \\ v_\kappa^\infty(L) - v_\kappa^\infty(L - 1) &= v_\kappa(L) - v_\kappa(L - 1). \end{aligned} \tag{26}$$

In general, for $\theta = 1, \dots, L$,

$$v_\kappa^\infty(\theta) - v_\kappa^\infty(\theta - 1) = \sum_{l=0}^{L-\theta} \gamma^l [v_\kappa(\theta + l) - v_\kappa(\theta + l - 1)], \tag{27}$$

where we define $\gamma = \delta(1 - \varepsilon)$. Thus, we obtain

$$\begin{aligned} & v_\kappa^\infty(\theta) - v_\kappa^\infty(0) \\ &= v_\kappa(\theta) - v_\kappa(0) + \gamma[v_\kappa(\theta + 1) - v_\kappa(1)] + \dots + \gamma^{L-\theta}[v_\kappa(L) - v_\kappa(L - \theta)] \\ & \quad + \gamma^{L-\theta+1}[v_\kappa(L) - v_\kappa(L - \theta + 1)] + \dots + \gamma^{L-1}[v_\kappa(L) - v_\kappa(L - 1)] \\ &= \sum_{l=0}^{L-1} \gamma^l [v_\kappa(\min\{\theta + l, L\}) - v_\kappa(l)], \end{aligned} \tag{28}$$

for $\theta = 1, \dots, L$. Since $-c \leq v_\kappa(\theta) \leq b$ for all θ , we have $v_\kappa(\theta) - v_\kappa(\tilde{\theta}) \leq b + c$ for all $(\theta, \tilde{\theta})$. Hence, by (25),

$$v_\kappa^\infty(\theta) - v_\kappa^\infty(0) \leq \frac{1 - \gamma^L}{1 - \gamma}(b + c) \leq \frac{b + c}{1 - \gamma} \tag{29}$$

for all $\theta = 1, \dots, L$, for all $\kappa = (L, \sigma)$. Therefore, if $\delta(1 - 2\varepsilon)[(b + c)/(1 - \gamma)] < c$, or equivalently, $c/b > [\beta(1 - \alpha)(1 - 2\varepsilon)]/[1 - \beta(1 - \alpha)(2 - 3\varepsilon)]$, then the incentive constraint (11) cannot be satisfied for any θ , for any social norm (L, σ) . This implies that any social strategy σ such that $\sigma(\theta, \tilde{\theta}) = F$ for some $(\theta, \tilde{\theta})$ is not feasible, and thus $U^* = 0$.

(iii) For any L , define a social strategy σ_L^{D0} by $\sigma_L^{D0}(\theta, \tilde{\theta}) = D$ for $\tilde{\theta} = 0$ and $\sigma_L^{D0}(\theta, \tilde{\theta}) = F$ for all $\tilde{\theta} > 0$, for all θ . In other words, with σ_L^{D0} each agent declines the service request of 0-agents while providing a service to other agents. Consider a social norm $\kappa = (1, \sigma_1^{D0})$. Then $v_\kappa(0) = -\eta_1(1)c$ and $v_\kappa(1) = b - \eta_1(1)c$. Hence, $U_\kappa = [1 - (1 - \alpha)\varepsilon](b - c)$ and $v_\kappa^\infty(1) - v_\kappa^\infty(0) = b$, and thus the incentive constraint $\delta(1 - 2\varepsilon)(v_\kappa^\infty(1) - v_\kappa^\infty(0)) \geq c$ is satisfied by the hypothesis $c/b \leq \beta(1 - \alpha)(1 - 2\varepsilon)$. Since there exists a feasible solution that achieves $U_\kappa = [1 - (1 - \alpha)\varepsilon](b - c)$, we have $U^* \geq [1 - (1 - \alpha)\varepsilon](b - c)$.

(iv) Suppose, on the contrary to the conclusion, that $U^* = b - c$. If $\alpha = 1$, then (11) cannot be satisfied for any θ , for any κ , which implies $U^* = 0$. Hence, it must be the case that $\alpha < 1$. Let $\kappa^* = (L^*, \sigma^*)$ be an optimal social norm that achieves $U^* = b - c$. Since $U^* = U_{\kappa^*} = (b - c) \sum_{\theta, \tilde{\theta}} \eta_{L^*}(\theta) \eta_{L^*}(\tilde{\theta}) I(\sigma^*(\theta, \tilde{\theta}) = F)$, σ^* should have $\sigma^*(\theta, \tilde{\theta}) = F$ for all $(\theta, \tilde{\theta})$. However, under this social strategy, all the agents are treated equally, and thus $v_{\kappa^*}^\infty(0) = \dots = v_{\kappa^*}^\infty(L^*)$. Then σ^* cannot satisfy the relevant incentive constraint (11) for all θ since the left-hand side of (11) is zero, which contradicts the optimality of (L^*, σ^*) .

(v) The result can be obtained by combining (i) and (iii).

(vi) Suppose that $U^* = b - c$, and let (L, σ) be an optimal social norm that achieves $U^* = b - c$. By (iv), we obtain $\varepsilon = 0$. Then $\eta_L(\theta) = 0$ for all $0 \leq \theta \leq L - 1$ and $\eta_L(L) = 1$. Hence, σ should have $\sigma(L, L) = F$ in order to attain $U^* = b - c$. Since $v_\kappa(L) = b - c$ and $v_\kappa(\theta) \geq -c$ for all $0 \leq \theta \leq L - 1$, we have $v_\kappa^\infty(L) - v_\kappa^\infty(0) \leq b/(1 - \gamma)$ by (25). If $\delta b/(1 - \delta) < c$, then the incentive constraint for L -agents, $\delta[v_\kappa^\infty(L) - v_\kappa^\infty(0)] \geq c$, cannot be satisfied. Therefore, we obtain $c/b \leq \delta/(1 - \delta)$.