

[Marco Janssen](#) (2006)

## Evolution of Cooperation when Feedback to Reputation Scores is Voluntary

*Journal of Artificial Societies and Social Simulation* vol. 9, no. 1  
<<http://jasss.soc.surrey.ac.uk/9/1/17.html>>

For information about citing this article, click [here](#)

Received: 13-Mar-2005 Accepted: 24-Aug-2005 Published: 31-Jan-2006



### Abstract

Reputation systems are used to facilitate interaction between strangers in one-shot social dilemmas, like transactions in e-commerce. The functioning of various reputation systems depend on voluntary feedback derived from the participants in those social dilemmas. In this paper a model is presented under which frequencies of providing feedback to positive and negative experiences in reputation systems explain observed levels of cooperation. The results from simulations show that it is not likely that reputation scores alone will lead to high levels of cooperation.

#### Keywords:

Trust, Reputation, One-Shot Prisoner Dilemma, Voluntary Feedback, Symbols

### Introduction

#### 1.1

Why do strangers sometimes cooperate in one-shot social dilemma situations? It is well understood that cooperation can evolve due to kin selection, ([Hamilton 1964](#)), or due to direct reciprocity in repeated interactions ([Trivers 1971](#); [Axelrod 1984](#)). The theories of indirect reciprocity and costly signaling show how cooperation in larger groups can emerge when the cooperators can build a reputation ([Nowak and Sigmund 1998](#); [Gintis et al. 2001](#)). Theories of indirect reciprocity become very important in the application in e-commerce. Strangers interact on-line and exchange large amounts of resources. The focus of the research on e-commerce and cooperation, is on trust, reputation and auction markets ([Malaga 2001](#); [Ba and Pavlou 2002](#); [Grabner-Kräuter and Kaluscha 2003](#); [Dellarocas 2003a](#); [Bolton et al. 2004](#)).

#### 1.2

Reputation systems like those from eBay are analyzed in terms of their effect on the value of the transactions in auctions. The effect of reputation is found to be mixed ([Stanifird 2001](#); [Ba and Pavlou 2002](#); [Grabner-Kräuter and Kaluscha 2003](#); [Resnick and Zeckhauser 2003](#); [Resnick et al. 2002](#); [Dellarocas 2003b](#)).

#### 1.3

In the study of reputation systems for e-commerce, the focus is on designing mechanisms that are effective in representing trust. Various problems are identified, including the inaccuracy of reflecting reputation when positive and negative feedbacks are aggregated to one score, barriers of entering the system by those who have not yet build up reputation, the lack of incentives to provide feedback, and the unlimited memory of various reputation systems ([Malaga 2001](#)). Despite the limitations of reputation systems, scholars conclude that these reputation systems seem to work ([Malaga 2001](#)). Among the transactions within eBay less than 1% of the feedback is negative ([Resnick et al. 2002](#)). However, only half of the buyers and sellers provide feedback ([Resnick et al. 2002](#)).

#### 1.4

Since not all participants provide feedback, it is not clear whether the statistics of much less than 1% negative feedback is representative. These statistics might suggest that participants are more likely to provide a positive feedback than a negative feedback, when they are confronted with positive or negative experiences. There are some possible explanations for this. Since negative rating carry a heavy weight, leaving negative feedback may cause retaliation ([Ba and Pavlou 2002](#); [Resnick and Zeckhauser 2003](#)). Members may be reluctant to leave negative feedback, fearing the action may endanger their own feedback profile ([Ba and Pavlou 2002](#)). On the other hand sellers actively solicit for positive feedbacks after a successful transaction, and positive feedback might be rewarded with a positive feedback.

#### 1.5

Reputation systems with voluntary feedback are public goods. Everybody benefits from the information provided to the reputation system, but each individual needs to invest some effort to provide feedback. Selfish rational actors may not

invest in public goods, but experimental studies show that investments in public goods are significant above zero (Ledyard 1995). If feedback is voluntary, and not all participants provide feedback, which frequency of providing honest feedback is required in order to make the reputation system functional? This question will be analyzed by using an agent-based model, where agents play one-shot prisoner dilemma games, and provide input to the reputation system with a certain probability. It will be shown that reputation scores alone are not likely to explain why systems like eBay function. Why do strangers, who interact only once, cooperate if the reputation system can not explain this? I will show that it is likely that participants in these social dilemmas use a variety of sources of information besides reputation scores in order to estimate the trustworthiness of their opponents. As shown by Janssen (in review), cooperation in one-shot prisoner dilemma games may evolve when agents are able to learn to recognize how symbols of agents relate to their trustworthiness. With respect to online-auctions, these symbols are derived by communication about the products by email or phone. Model analysis shows how symbols additional to the reputation score affect the level of cooperation. Note that the model in this paper is not aimed to a specific application but is used to address the general question of the impact of voluntary feedback for reputation on cooperation.

1.6

In the next section a model will be presented that combines ideas from partner selection, reputation scores, and the recognition of trustworthiness via symbols displayed. In section 3 the experimental setup will be discussed and the results are presented in section 4. Section 5 concludes on the implications of the findings.

 **The Model**

2.1

The model is an adapted version of Janssen (in review) and consists of a population of  $n$  players, who randomly play one-shot two-person Prisoner's Dilemma games. In that study Janssen shows that a population of agents playing one-shot PD games can evolve into a population of agents who are conditionally cooperative, leading to high levels of cooperation, if they are able to learn who to trust. The success of learning depends on the amount of information available, the number of symbols, which needs to be 20 or higher to derive cooperation. We built here on the insights from the study of Janssen (in review) by including a very specific symbol, a reputation score. The model is inspired on the puzzle why eBay works although there are design problems with the reputation system (Malaga 2001). I will make an abstraction of the eBay system by assuming one-shot Prisoner's Dilemma games, the same type of agents (no distinction of sellers and buyers), and agents having equal rate of being matched up (randomly) with another agent (although they may decide not to play). I do not include explicit auctions, neither different type of products. In sum, the model is an abstraction of dilemmas subjects experience on electronic commerce systems like eBay. The following components are discussed: the reputation system, which strategy the players use, what types of symbols the players display, how they learn to recognize trustworthiness of these symbols, and the entry and exit rules.

**The Game**

2.2

Each individual has three possible actions: cooperate (C), defect (D), or withdraw (W). If both players cooperate, they each get a payoff of R (reward for cooperation). If both players defect, they each get a payoff of P (punishment for defecting). If player A defects and B cooperates, A gets a payoff of T (temptation to defect), and B gets S (sucker's payoff). If at least one of the players withdraws from the game, both players get a payoff of E (exit payoff). The resulting payoffs table is given in Table 1.

**Table 1:** Pay-off table of the Prisoner's Dilemma with the option to withdraw from the game

		Player B		
		Cooperate	Defect	Withdraw
Player A	Cooperate	R,R	S,T	E,E
	Defect	T,S	P,P	E,E
	Withdraw	E,E	E,E	E,E

In line with Tullock (1985) and Vanberg and Congelton (1992), I assume that the costs and benefits of exit are such that the expected payoffs from choosing not to play are higher than those resulting from mutual defection, but lower than those expected from mutual cooperation. The Prisoner's Dilemma is defined when  $T > R > E > P > S$  and  $2R > T + S$ . In this situation the best option for any one move is to withdraw from the game. If one expects that the other agent will cooperate, the best option is to defect. If one expects that the other agent will defect, the best option is to withdraw. Each player comes to the same conclusion, so they both withdraw and end up with payoffs that are much lower than if they both trust that the other will cooperate. The pay-off matrix for the game in this article is defined using  $T = 2, R = 1, E = 0, P = -1,$  and  $S = -2,$  which is in line with Tullock (1985).

2.3

Experimental research has shown that the material payoff does not have to equal the utility payoff experienced by the players (see, for example, Ahn et al., (2001; 2003)), Not every subject shows selfish behavior. In fact, the majority of non-selfish players seem to be conditionally cooperative and cooperate only when they know that the other will probably cooperate. To include the heterogeneity of motivations the utility function is used as formulated in Table 2, where the material payoffs can be adjusted by individual motivations. The  $\alpha$  value can be regarded as the strength of an individual's aversion to exploiting others, and  $\beta$  can be regarded as an individual's degree of altruistic tendency. These  $\alpha$  and  $\beta$  values determine the strategies of the agents to cooperate or not. Furthermore, the simulated process

may affect the distribution of  $\alpha$  and  $\beta$  values within the population. Both aspects will be explained below in more detail.

**Table 2:** Utility pay-off table of the Prisoner's Dilemma with the option to withdraw from the game

		Player B		
		Cooperate	Defect	Withdraw
Player A	Cooperate	R,R	$S + \beta_A, T - \alpha_B$	E,E
	Defect	$T - \alpha_A, S + \beta_B$	P,P	E,E
	Withdraw	E,E	E,E	E,E

An agent has four elements: (1) the reputation score, (2) the set of symbols that it displays, (3) the strategy that it uses to decide whether to trust or not another agent, and (4) the strategy it uses in Prisoner's Dilemma games.

## Reputation

### 2.4

Many types of reputation systems exist. Inspired by the query why the eBay reputation systems seem to work although it is not theoretically sound (Malaga 2001), we define the reputation score  $RS$  in this paper in line with the eBay reputation system. An agent receives +1, 0 or -1 as feedback for the action they took in the game. As discussed in the introduction, feedback is voluntary and not all agents may provide feedback. The feedback frequency is assumed to depend on probabilities  $p_D$  and  $p_C$ . With probability  $p_D$  an agent gives (negative) feedback when the opponent defected, and with probability  $p_C$  an agent gives (positive) feedback when the opponent cooperated. In the initial experiments, the probability of feedback is unrelated to the feedback from the response of the opponent, and unrelated to the  $\alpha$  and  $\beta$  values of the agents. This is quite a simplification, but it enables me to start to explore the consequences if not all agents provide feedback. Future work may include endogenous feedback responses and decisions to retaliate.

### 2.5

The reputation score the summation of all feedback it receives over a fixed amount of time an agent played or was drawn with the opportunity to play a game. This period is denoted by  $I_m$ . The reputation score is represented as a symbol of the agents. For that use it is scaled to the maximum score by dividing it by  $I_m$ , such that the value is between -1 and 1.

## Symbols

### 2.6

Besides a reputation score, agents can have  $s$  other symbols, which are represented in the following way. The symbols can have value between 0 and 1. The values of the symbols are drawn in such a way that a positive, negative or no correlation exists with the value of  $\alpha_i$  of agent  $i$ . For each symbol  $x_j$  a random number  $\omega_j$  is drawn from a uniform distribution. This value of  $\omega$  is used to create a correlation with the type of preferences of the agents, which the agents need to learn during the interactions. The value of the symbol  $x_j$  is now defined as

$$x_j = \left( \frac{\alpha_i}{\alpha_{\max}} \cdot \omega_j + (1 - \omega_j) \cdot \left(1 - \frac{\alpha_i}{\alpha_{\max}}\right) \right)^2 \quad (1)$$

where  $\alpha_{\max}$  is used to scale the symbol value between 0 and 1. In order to represent the assumed nonlinear relationship between the preferences of the agent and the symbols, the equation is squared.

### 2.7

The agents are not aware of the imposed correlations of the symbols and types of agents. They will learn possible relationships during repeated play of one-shot PD games. The symbols can be used to estimate the trustworthiness of partners in the games.

## Trust

### 2.8

The rule an agent uses to decide to trust the other agent, and thus be willing to play a Prisoner's Dilemma game, is represented as a simple updating of relative weight to different symbols. In fact, the model is a single-layer neural network (Mehrota et al. 1997). The agents learn the relation between trustworthiness and  $s+1$  symbols, which have values between -1 and 1. A weighted sum  $M$  of these inputs is calculated using the following equation:

$$M = w_0 + \sum_{i=1}^{s+1} w_i x_i, \quad (2)$$

where  $w_0$  is the bias,  $w_j$  is the weight of the  $i$ th input, and  $x_j$  is the  $i$ th input. Initially, all weights are zero, but during the simulation the network is trained, when new information is derived, by updating the weights as described below in equation (4).

## 2.9

The relationship between expected trustworthiness and the weighted sum of symbol inputs  $M$  is assumed to depend on a sigmoid function. This function is used in neural networks as a threshold function to translate the inputs into one output, and it determines the probability  $Pr[Tr]$  that the agent will choose to trust its prospective partner:

$$Pr[Tr] = \frac{1}{1 + e^{-M}} \quad (3)$$

## 2.10

The higher the value of  $M$ , the higher the probability will be. The probability of not trusting the other agent is  $1 - Pr[Tr]$ . Initial weights are assumed to be randomly drawn from a uniform distribution between zero and one. For the bias  $w_0$  an initial value is assumed to be drawn from a uniform distribution between minus one and zero since a default low trust is assumed when agents express no symbols and have no reputations.

## 2.11

If a game is played, each agent receives feedback,  $F$ , on the experience. This feedback is simply whether the partner cooperated or not. If the partner cooperated ( $F = 1$ ), the agent adjusts the weights associated with the other agent's symbols upward, so that it will be more likely to trust that agent, and others displaying similar symbols, in the future. On the other hand, if the partner defected ( $F = 0$ ), the agent will adjust the same weights downward, so that it will be less likely to trust that agent and others with similar symbols. The equation to adjust the weights is as follows:

$$\Delta w_i = \lambda \cdot (F - Pr[Tr]) \cdot x_i, \quad (4)$$

where  $\Delta w_i$  is the adjustment to the  $i$ th weight,  $\lambda$  is the learning rate,  $F$  is the feedback,  $F - Pr[Tr]$  is the difference between the agent's level of trust in the other agent and the observed trustworthiness of the other agent, and  $x_i$  is the other agent's  $i$ th symbol. In effect, if the other agent displays the  $i$ th symbol, the corresponding weight is updated by an amount proportional to the difference between the observed trustworthiness of an agent and the trust placed in that agent. The weights of symbols associated with positive experiences increase, while the weights of those associated with negative experiences decrease, reducing discrepancies between the amount of trust placed in an agent and that agent's trustworthiness.

## 2.12

The initial values of  $\alpha$  are drawn from a uniform distribution between 0 and 3, and for  $\beta$  between 0 and  $\alpha$ , by which only initial conditions are accepted where  $\beta \leq \alpha$ .

## Strategies

## 2.13

When neither agent withdraws from playing a game they have to decide to cooperate or to defect. The agents' decisions to cooperate or defect are based on expected utilities. I assume that the players will estimate the expected utility for cooperation,  $E[U(C)]$  or defection,  $E[U(D)]$ . The expected utility is determined by assuming that the level of expected trust of an agent in its opponent, defined in (3), represents the probability that the opponent will cooperate:

$$E[U(C)] = Pr[Tr] \cdot R + (1 - Pr[Tr]) \cdot (S + \beta_i) \quad (5)$$

$$E[U(D)] = Pr[Tr] \cdot (T - \alpha_i) + (1 - Pr[Tr]) \cdot P \quad (6)$$

Given the two estimates of expected utility, the player is confronted with a discrete choice problem which addressed with a logit function. The probability to cooperate,  $Pr[C]$ , depends on the expected utilities and the parameter  $\gamma$ , which represents how sensitive the player is to differences in the estimates. The higher the value of  $\gamma$ , the more sensitive the probability to cooperate is to differences between the estimated utilities.

$$Pr[C] = \frac{e^{\gamma \cdot E[U(C)]}}{e^{\gamma \cdot E[U(C)]} + e^{\gamma \cdot E[U(D)]}} \quad (7)$$

Logit models are used by other scholars to explain observed behavior in one-shot games such as the logit equilibrium approach by Anderson et al. ([forthcoming](#)). Although the functional relation is similar, their approach differs from the one used here since they assume an equilibrium and perfect information of the actions and motivations of other players. Moreover, in the games in this paper agents do not play anonymously but can observe symbols of others in

order to estimate the behavior of the opponent.

## Exit And Entry Rules

### 2.14

The population size remains constant during the simulation, but agents exit the system and are replaced by random new agents. This leads to an evolution of the types of agents in the population. Agents can leave the game if the reputation score, RS, reaches the value of the maximum negative reputation score, MNRS. We explored the consequences of various reasonable values, but the qualitative results are not sensitive to our assumption. Since our motivation for the analysis is triggered by the puzzle why eBay's reputation system works, we decided to use the threshold value of  $-4$ , as used in the eBay reputation system. In the case that accumulated income of the agents is not positive, the agent is also assumed to leave the system, since it does not experience a positive incentive to continue. We think it is a reasonable rational decision to exit the system, when an agent is not able to derive positive income. Note that this hinders the introduction of cooperative agents, since they are the ones who derive negative monetary payoffs.

### 2.15

An agent that has left the population is replaced by a new agent with random drawn values for  $\alpha_i$ ,  $\beta_i$ , and  $w_i$ .



## Experimental Design

### 3.1

In the model experiments the effect of the memory length, symbols, and the probabilities of providing feedback will be analyzed. Parameter values for the default case are presented in Table 3. This default case will be used as a reference for testing the sensitivity for a number of assumptions.

### 3.2

Each set of parameter values is used to simulate 100 runs, where each run consists of  $10^6$  rounds. This number of rounds is chosen to derive a stable distribution of population characteristics. Taking into account an initialization period and focusing on evolved equilibria, statistics for the last  $10^5$  rounds is recorded. Since agents enter and leave the system, they do play less than  $10^6$  games on average.

**Table 3:** List of parameters and their default values

Parameter	Value
Population size (n)	100
Number of symbols (s)	0
Conditional cooperation parameter $\alpha$	[ $\beta$ , 3]
Conditional cooperation parameter $\beta$	[0, $\alpha$ ]
Learning rate ( $\lambda$ )	0.1
Number of round	$10^6$
Steepness ( $\gamma$ )	4
Length of memory ( $l_m$ )	100
Probability of getting feedback when cooperate.	$p_C$
Probability of getting feedback when defect.	$p_D$

In the first set of experiments the agents are always assumed to provide feedback ( $p_C=1$  and  $p_D=1$ ), only include the reputation score. We also vary the length of memory  $l_m$  among the experiments. The memory length is argued to be one of the problematic factors of reputation systems ([Malaga 2001](#)). A short memory does not capture sufficient information to distinguish different types of behavior, while a long memory may hinder to identify a change in behavior. Note that memory of previous decisions is usually lumped in reputation scores in the number of positive and negative experiences during the last  $l_m$  games. Therefore the topic of interest is the minimum length of memory before the reputation system becomes effective, and whether a long memory reduces the effectiveness of the reputation system.

### 3.3

In the second set of experiments, the consequences are explored of different probability levels of providing feedback when only reputation scores are used. The interest is whether minimum levels of feedback are required in order to have an effective reputation system.

### 3.4

In the third set of experiments we include the possibility of agents to retaliate. We are interested in how retaliation affects the level of cooperation derived with different probability levels of voluntary feedback. Agents make their decisions in a random order, and when the first agent provides a negative feedback, the other agent will do so, independent of the decision the first agent made in the one-shot game.

### 3.5

In the fourth set of experiments the variation of probability of providing feedback is extended by including additional symbols  $x_i$ . The goal of this experiment is to assess the impact of inclusion of more symbols on the level of cooperation.

### 3.6

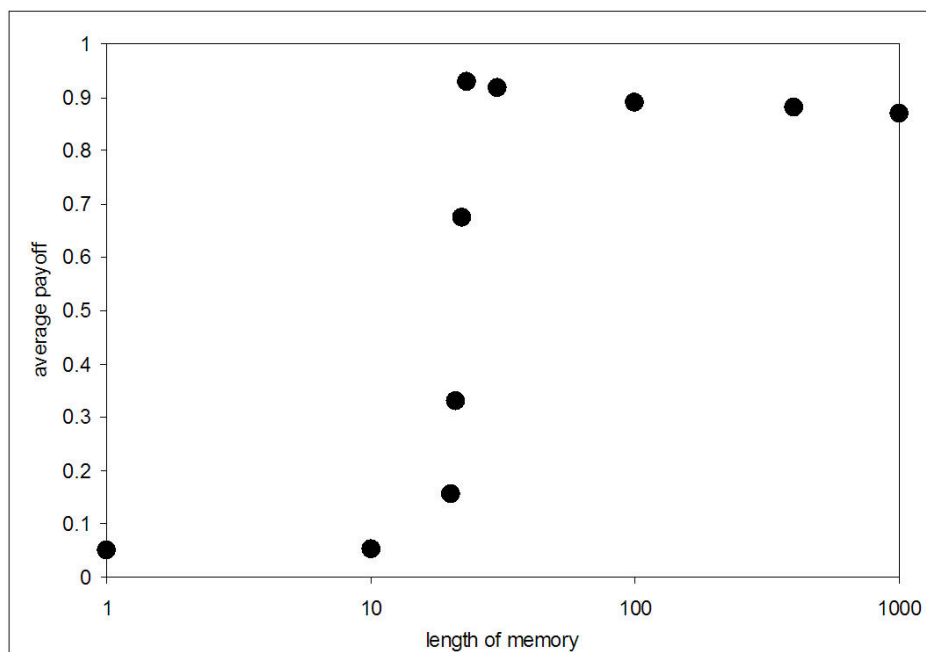
The last set of experiments tests the consequence if an agent can act somewhat strategically. Selfish agents may not defect always if they understand the importance of reputation and expect to derive more returns in future rounds by cooperating. This is implemented by assuming that agents with lower values of  $\alpha$  weight more the expected utility derived from the current round, with the expected utility for the next round. Since defection or cooperation affect their reputation score, they can estimate the expected utility in the next round given the action in the current round.



## Results

### 4.1

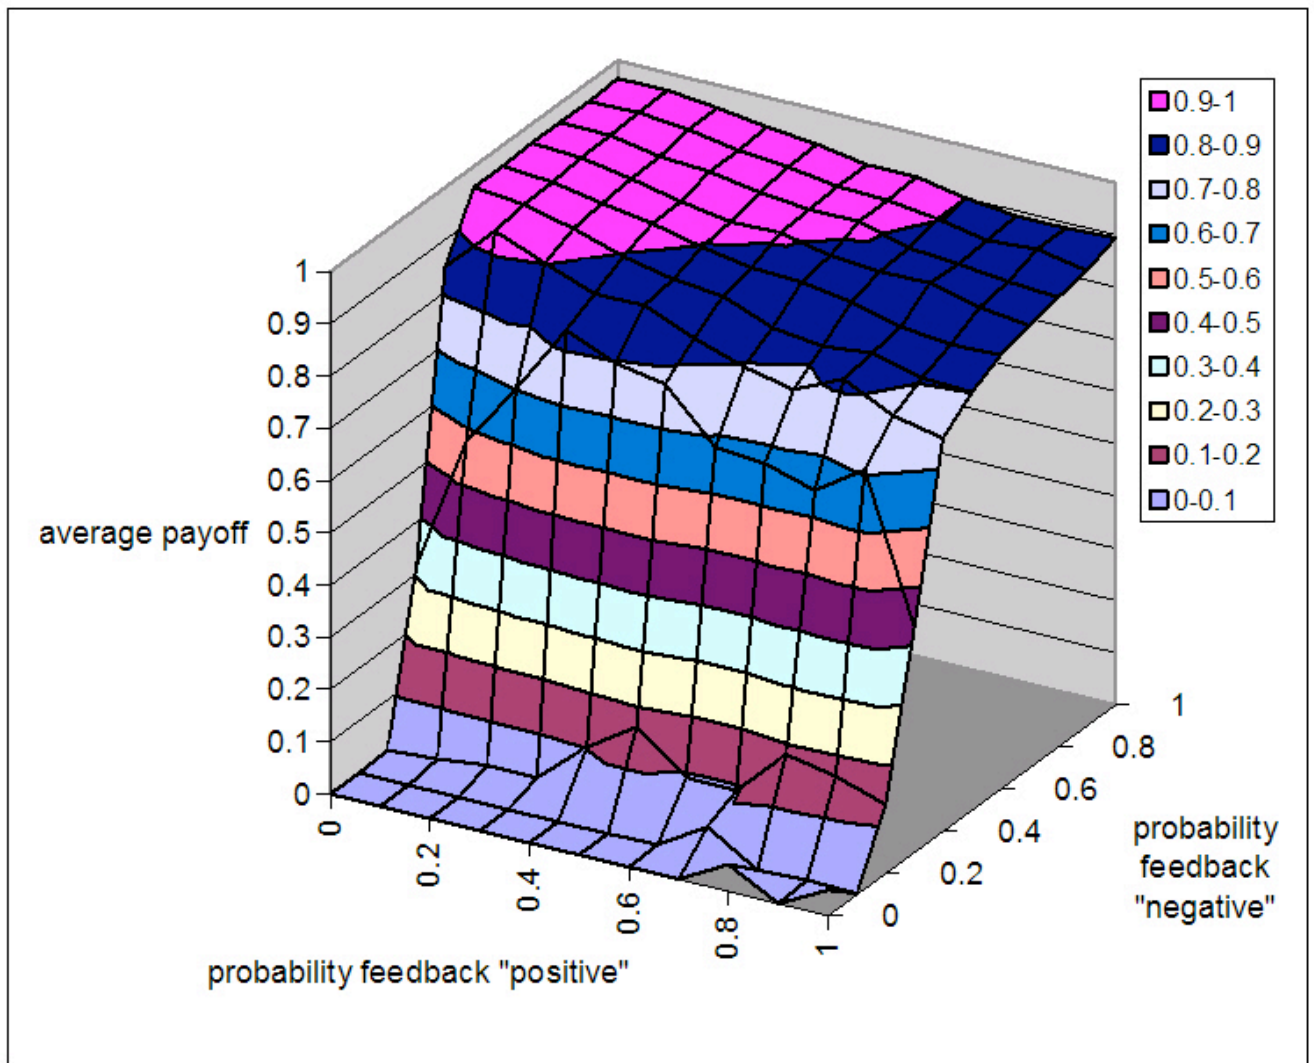
The length of memory affects the ability of a reputation system to include sufficient relevant information about the reputation of your opponents in the one-shot PD games. Therefore, the level of cooperation is low when the memory length is smaller than 25 interactions (Figure 1). When memory length is small it is difficult to distinguish cooperative agents from less trustworthy agents. Therefore, agents learn not to play, or to play randomly, leading to an average payoff near zero. Around a memory length of 25 the agents learn to identify who they can trust, leading to mutual cooperation of trustworthy agents, and to the exit of less trustworthy agents. Note that these 25 interactions also include cases where drawn agents did decide not to play the game itself. Increasing the length of memory leads to a very modest decline. The decline is caused by a longer existence of agents with modest  $\alpha$  and  $\beta$  values within the population, since occurrence of negative feedback is more likely to be compensated by positive feedback. A longer length of memory makes the reputation system a little bit less effective, but agents with low values of  $\alpha$  and  $\beta$ , will be detected rapidly anyway if agents always provide feedback.



**Figure 1.** Average payoff per agent per game for different length of history. The dots represent the average payoffs over  $10^5$  interactions over 25 simulations

### 4.2


When agents do not always give feedback, but with probabilities  $p_C$  and  $p_D$ , the level of cooperation varies significantly among the various probabilities (Figure 2). The highest level of cooperation is derived when agents give only and always feedback if their opponent defected. In that case selfish agents will be identified the fastest and removed from the system. If the probability of feedback to positive experiencing increases, agents who sometimes cooperate, but also frequently defect, are maintained longer into the system, and reduce the level of cooperation. Interestingly, if the probability of providing feedback to negative experiences decreases below 0.3, no cooperation is derived. Uncooperative agents are identified too slow and moved too late to make the reputation score sufficient. Thus voluntary reputation systems do not always work in one shot social dilemma games, especially when agents do not provide feedback to negative experiences with a high probability.



**Figure 2.** The average payoff per agent per game for different levels of  $p_C$  and  $p_D$ , when agents express only reputation score. The historical memory length is 100 interactions

#### 4.3

Agents make their decisions in random order. When agents retaliate, the second agent gives a negative feedback when the first agent did so. These retaliations lead to a lack of cooperation when agents do not give positive feedback and only high levels of feedback in negative cases (Figure 3 for  $p_C = 0$  and  $p_D = 1$ ). The reason for this is that with  $p_C = 0$  agents only give negative feedback, even when they had positive experience but retaliate the feedback of the first agent. As a result agents in a population starting with a random distribution of other-regarding preferences do not benefit strongly from being altruistic (and still getting negative feedbacks). This leads to a no cooperation level in Figure 3.




**Figure 3.** The average payoff per agent per game for different levels of  $p_C$  and  $p_D$ , when agents express only reputation score and retaliate. The historical memory length is 100 interactions


#### 4.4

Suppose 10 additional symbols are included. In that case agents derive more information and are able to recognize trustworthy agents even when the probability of feedback to negative experiences is low (Figure 4). Agents learn possible correlations between symbols and types of agents. eBay users, for example, exchange emails, phone each other, communicate with each other on the item, and can provide additional technical information on the items to be traded, such that the players involved can improve their judgment on the trustworthiness of the other player. As would be expected, larger numbers of symbols lead to higher levels of cooperation (Figure 5).






**Figure 4.** The average payoff per agent per game for different levels of  $p_C$  and  $p_D$ , when agents express reputation scores and 10 additional symbols. The historical memory length is 100 interactions



**Figure 5.** The average payoff per agent per game when  $p_C = 1$  and  $p_D = 0$ , for different numbers of additional symbols. The historical memory length is 100 interactions

4.5

Agents may express strategic behavior, by taking into account the effect of cooperation and defection on the expected future returns (Figure 6). In the experiment only the reputation symbol is included. Inclusion of strategic behavior lowers somewhat the levels of the average payoff, but does affect the results from Figure 2 qualitatively in a minor way. There is a higher threshold of providing feedback to defectors before the reputation system will be effective. The reason for this is that for higher probabilities of feedback to cooperation, strategic behavior will make selfish agents somewhat more cooperative. Therefore it is even more important that defections are reported to make the reputation system effective.



**Figure 6.** The average payoff per agents per game for different levels of  $p_C$  and  $p_D$ , when agents express reputation scores and act strategically. The historical memory length is 100 interactions

4.6

Other scholars have looked at the consequences of missing feedback. Dellarocas (2003b) analyzed online feedback mechanism design issues in trading environments with opportunistic sellers. Dellarocas shows that if missing feedback is dealt with as positive feedback in the reputation system, the system reaches a high level of cooperation. However, in the case not everybody provides feedback and missing feedback is not assumed to be positive feedback, no high levels of cooperation will be derived if the probability of positive feedback is higher than the probability of negative feedback.

## Conclusions

5.1

A simple agent based model is used to explore the possible consequences of varying assumptions on voluntary feedback to the reputation system. The model is an abstraction of electronic commerce systems, where agents play one-shot PD games, and does not include differences between sellers and buyers, and does not include specific auctions. The goal was not to mimic a specific empirical situation, but provide a more general starting point for analyzing the effectiveness of reputation systems with voluntary feedback.

5.2

The current model distinguishes different type of agents with regard to providing feedback, but the agents do not rationalize when to give feedback. This could be an interesting extension of the model. For example, retaliation or exchange of positive evaluations might become evolving strategies when feedback decisions are endogenized.

5.3

When contributions to reputation systems are voluntary, reputation systems may not function especially when the level of feedback to negative experiences is lower than the level of positive experiences. This might be the case in electronic commerce where sellers actively try to collect positive responses to increase their reputation score. Given the assumptions of the model, the model analysis suggests that reputation scores alone can not explain the high level of

cooperation in various examples of electronic commerce, like eBay. A possible explanation is the use of multiple other types of information by the agents to detect trustworthiness of the agents. More understanding is needed about what alternative kinds of information by the participants that can explain high levels of cooperation in systems of one-shot interactions.



## Acknowledgements

The author thanks TK Ahn, Nigel Gilbert and two anonymous reviewers for useful comments on an earlier version of this paper. Support from the National Science Foundation (SES0083511) is gratefully acknowledged.



## References

- AHN T-K, Ostrom E, Schmidt D, Shupp R and Walker J (2001), Cooperation in PD games: Fear, greed, and history of play. *Public Choice*, 106, pp. 137–155.
- AHN T.K, Ostrom E and Walker J (2003), Incorporating motivational heterogeneity into game theoretic Models of collective action. *Public Choice*, 117. pp. 295–314.
- ANDERSON S P, Goeree J K and Holt C A (Forthcoming), Logit equilibrium models of anomalous behavior: What to do when the Nash equilibrium says one thing and the data something else. In C. Plott & V. Smith (Eds.) *Handbook of Experimental Economic Results*. New York: Elsevier.
- AXELROD R (1984), *The Evolution of Cooperation*. New York: Basic Books.
- BA S and Pavlou P A (2002), Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior, *MIS Quarterly* 26(3). pp. 1–26.
- BOLTON G E, Katok E, and Ockenfels (2004) How effective are electronic reputation mechanisms? An experimental investigation. *Management Science* 50(11). Pp. 1587–1602
- DELLAROCAS C (2003a), The digitization of worth of mouth: promise and challenges of online feedback mechanisms, *Management Science* 49(10). pp. 147–1424.
- DELLAROCAS C (2003b), *Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard*, Working paper 4297–03, MIT Sloan School of Management.
- GINTIS H, Smith E and Bowles S (2001), Costly signaling and cooperation. *Journal of Theoretical Biology*, 21. pp. 103–119.
- GRABNER-KRÄUTER S and Kaluscha E A (2003), Empirical research in on-line trust: a review and critical assessment, *International Journal of Human-Computer Studies* 58. pp. 783–812.
- HAMILTON W D (1964), Genetical evolution of social behavior I and II. *Journal of Theoretical Biology*, 7. pp. 1–52.
- JANSSEN M A (in review), Evolution of Cooperation in a One-Shot Prisoner's Dilemma Based on Recognition of Trustworthy and Untrustworthy Agents.
- LEDYARD J O (1995), Public Goods: A Survey of Experimental Research," in *Handbook of Experimental Economics*, J. Kagel and A. Roth (eds.), Princeton University Press.
- MALAGA R A (2001), Web-based reputation management systems: problems, and suggested solutions, *Electronic Commerce Research* 1. pp. 403–417.
- MEHROTA K, Mohan C K and Ranka S (1997), *Elements of Artificial Neural Networks*. Cambridge, MA: MIT Press.
- NOWAK M A and Sigmund, K (1998), Evolution of indirect reciprocity by image scoring. *Nature*, 393. pp. 573–577.
- RESNICK P and Zeckhauser R (2003), Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. M.R. Baye (ed.). *The Economic of the Internet and E-Commerce*. Advances in Applied Microeconomics, 11. JAI Press, Greenwich CT.
- RESNICK P, Zeckhauser R, Swanson J and Lockwood K (2002), The value of reputation on eBay: A controlled experiments. Working paper, University of Michigan, Ann Arbor, MI.
- STANIFIRD S S (2001), Reputation and e-commerce: eBay auctions and the asymmetrical impacts of positive and negative ratings, *Journal of Management* 27. pp. 279–295.
- TRIVERS R (1971), The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46. pp. 35–57
- TULLOCK G (1985), Adam Smith and the Prisoner's Dilemma. *Quarterly Journal of Economics*, 100. pp. 1073–1081.
- VANBERG VJ and Congelton R D (1992), Rationality, morality and exit. *American Political Science Review*, 86. pp. 418–431.

[Return to Contents of this issue](#)

© [Copyright Journal of Artificial Societies and Social Simulation, \[2006\]](#)

