ELSEVIER

Effects of a reputation feedback system on an online consumer-to-consumer auction market

Jian Yang [a,1], Xiaorui Hu [b,2], Han Zhang [c,*]

[a] *Department of Industrial and Manufacturing Engineering, New Jersey Institute of Technology Newark, NJ 07102, United States*
[b] *John Cook School of Business St. Louis University, St. Louis, MO 63108, United States*
[c] *College of Management, Georgia Institute of Technology, Atlanta, GA 30332, United States*

Abstract

This research establishes a dynamic game-theoretic model that interprets the mechanism of reputation feedback systems in online consumer-to-consumer (C2C) auction markets. Based on the model, a numerical study is conducted to reveal the effects of feedback systems on auction markets. The study shows that the existence of feedback systems greatly improves the performance of online C2C auction markets: buyers are more willing to trade and gain more benefit from the transactions; sellers' honest behavior is encouraged, as honest sellers' gains are increased and dishonest sellers' gains are reduced. It also offers practical insights on the design of a feedback system: rewarding an honestly-behaving seller is less effective on promoting market performance than punishing a cheating seller.
© 2007 Elsevier B.V. All rights reserved.

1. Introduction

We study, from both theoretical and numerical angles, reputation feedback systems (hence-forward feedback systems) in online consumer-to-consumer (C2C) auction markets. Our purpose is to show that reasonably-designed feedback systems can promote trust and mitigate fraud, and help ensure the healthy development of the online markets. We also intend our numerical study to offer managerial insights on the market impacts of feedback systems' design features.

While online C2C auction markets have been growing rapidly in recent years, fraud in these markets is also on the rise. The average loss per claim in online auction frauds jumped from $895 in 2004 to $1917 in 2005 [14]. In online markets, interacting with strangers is inevitable, and most transactions between buyers and sellers are one-time deals. Meanwhile, auction sites serve as market makers for buyers and sellers to meet, but claim no liability for any fraudulent transactions. For example, eBay claims that they "have no control over the quality, safety or legality of the items advertised, the truth or accuracy of the listings" [10]. Thus, online auction participants have to face a market where goods are purchased before one can assure the quality. Therefore, promoting trust between strangers and reducing the uncertainty and risk for online traders are

* Corresponding author. Tel.: +1 404 894 4373.
E-mail addresses: yang@adm.njit.edu (J. Yang), hux2@slu.edu (X. Hu), han.zhang@mgt.gatech.edu (H. Zhang).
[1] Tel.: +1 973 596 3658.
[2] Tel.: +1 314 977 3849.

the critical issues facing current online C2C auction markets.

Various mechanisms have been designed and utilized in C2C auction markets to promote trust and reduce risk. Specific online payment systems (e.g., PayPal) have been implemented to provide secure and instantaneous online transactions for small online merchants and auction buyers. Online auction sites (e.g., eBay) have also been offering limited insurance or guarantees to protect auction participants and create a safer environment to trade. Feedback systems have also been offered by most of the C2C auction sites to reduce online traders' uncertainties and the risks associated with online trading. For instance, eBay's "Feedback Forum" is a form of community enforcement. In eBay, after each trade, both buyers and sellers are encouraged to leave comments about their trading partners based on their experience. Comments about traders are kept under each trader's profile, and can be accessed by everyone who visits eBay. This way, the system tries to deter dishonest behavior by conveying facts and opinions about past trades.

Does the feedback system work as advertised? There is substantial research that says it does. Kollock [15] conceptually summarizes online reputation systems and concludes that their effectiveness to manage the risks of unsecured trades seems to be impressive. Resnick et al. [20] review the online reputation systems and argue that the reputation systems appear to perform reasonably well despite their theoretical and practical difficulties. Resnick and Zeckhauser [19] empirically examine a large data set from eBay and claim that the reputation system appears to be effective. Ba and Pavlou [1] empirically explore the extent to which trust can be induced by proper feedback mechanisms in electronic markets and find that feedback systems can generate price premiums for reputable sellers. Resnick et al. [21] conduct a controlled experiment on eBay to assess returns to reputation.

In the related economics literature on repeated games involving reputation, most papers deal with a situation where players are not strangers to each other and a player's reputation is equivalent to the entire history of his actions. The reader may refer to Kreps and Wilson [16], Milgrom and Roberts [17], Fudenberg and Levine [11,12], Cripps and Thomas [6], Celentani et al. [5], Battigali and Watson [3], etc. for a glimpse of this body of literature. In an online C2C auction market, players are strangers to each other, and without a proper mechanism being installed, they know very little about each other when they trade. Moreover, it is unrealistic for any mechanism to require that every trader's entire history be remembered (stored). A feedback system, on the other hand, offers a platform where merely a

function, usually many-to-one, of every trader's entire history needs to be stored.

Dellarocas and Bakos have recently studied feedback systems from a theoretical perspective. Dellarocas [7] provides an overview of relevant past research on reputation mechanisms based on the repeated-game setting, and points out several future research directions concerning online feedback systems. Bakos and Dellarocas [2] study a trading system involving a single seller who repeatedly trades with buyers. The authors show that even the most primitive feedback system, as incarnated by the binary score, can provide an economically more efficient solution than the threat of litigation. Dellarocas [9] examines a similar model with a more sophisticated scoring system. The author is able to obtain a closed-form solution to the problem, which shows that sustainable cooperations between buyers and the seller are achievable as long as the return–cost ratio to the seller is high enough. Dellarocas [8] shows that in a certain environment, the combination of listing fees and binary-score feedbacks can induce sellers to announce the true quality of their products and at the same time maximize the average social welfare. From another angle, Miller, Resnick, and Zeckhauser [18] examine the elicitation of proper buyers' feedback writing behavior that makes a feedback system function.

As Bakos and Dellarocas [2] and Dellarocas [9] have done, this paper studies the merits of online feedback systems in a repeated-game setting where buyers' reputation score updating behavior is made exogenous, and sellers are assumed to be of different types, with each type pertaining to a specific tendency towards cheating. It differs from the two aforementioned papers most saliently in two aspects: 1) multiple seller types are considered, so that the score associated with a seller infers his type in addition to his future behavior. Because of this, features of Bayesian learning and the marriage of perceived and actual seller distributions appear in the formulation; and 2) it contains a general, as opposed to stylized, feedback system and speculates that a set of stochastic ordering relationships in its evolution is what makes the system work.

Specifically, we propose a dynamic game-theoretic framework to model the mechanism of a feedback system in an online C2C auction market. In the framework, we assume that all buyers are honest while sellers are of different types with different propensities for cheating. The feedback system is comprised of scores associated with sellers, which are updated by their respective trading partners in ways that are dependent on the treatment the partners have received. As the trading game is played over a sufficiently long period of time, buyers will form associations between sellers' scores and their types on the

basis of Bayesian learning from past experiences. For example, a seller with more cheating history (higher score) is considered more prone to cheating than one with less cheating history, and buyers will treat sellers with different scores differently. Knowing the different possible treatments from buyers, sellers will weigh their decisions about cheating/playing honestly based not only on their immediate one-time gain but also on their future business opportunities, which are affected by how their scores are updated by their respective trading partners.

Following the theoretical analysis, a numerical study is presented. The numerical results verify the benefits introduced by the feedback system. Our results also show that even with feedback systems in place, dishonest behavior from traders with excellent reputation ratings occurs, but at the same time, feedback systems do enhance the overall level of honesty in online auction markets.

The rest of this paper is structured as follows: in Section 2 we establish the dynamic game-theoretic model; in Section 3 we discuss the anticipated properties of various market performances; in Section 4 we conduct a numerical study based on the model and analyze the findings from the study; and in Section 5 we conclude the paper. We have relegated materials of secondary importance to our appendices. The latter will be available upon request.

2. Problem formulation

For the reader to better understand our formulations, we have compiled a list of symbols used in the general model in Appendix A.

2.1. The stage-game setting

Each seller is characterized by a prone-to-cheating factor $z \geq 0$, which reflects his average gain while cheating in a transaction. Buyers are all of the same type in terms of their payoff functions and they never cheat. Each transaction is characterized by a size parameter $x \geq 0$, which can be viewed as the trading surplus from the trade. A more realistic model may allow buyers to cheat as well. However, current online C2C auction markets allow buyers less chance to cheat, as they are required to pay before goods are shipped. Thus our model offers a reasonable approximation. Miller, Resnick, and Zeckhauser [18] offer similar arguments on why a seller's reputation is more important than a buyer's.

In a size-x transaction that involves a type-z seller, if the seller plays honestly, his average gain will be $g_{SX}^{H}(x)$, and the buyer's (his trading partner's) average gain will be $g_{BX}^{H}(x)$. However, if the seller cheats, his average gain

will be $g_{SXZ}^{C}(x, z)$ while the buyer's average gain will be $g_{BXZ}^{C}(x, z)$. Both gains depend on the level of the cheating which is captured by the factor z. Reasonably, we should have the following assumptions.

A seller with a larger z gains more from cheating:

$$g_{SXZ}^{C}(x, z + \Delta z) \geq g_{SXZ}^{C}(x, z) \text{ for } \Delta z \geq 0. \tag{1}$$

A seller always benefits from trading, and in a single stage, he gains more from cheating than behaving honestly:

$$g_{SXZ}^{C}(x, 0) \geq g_{SX}^{H}(x) > 0. \tag{2}$$

A buyer hurts more when a seller with a larger z cheats:

$$g_{BXZ}^{C}(x, z + \Delta z) \leq g_{BXZ}^{C}(x, z) \text{ for } \Delta z \geq 0. \tag{3}$$

A buyer benefits from trading when his trading partner behaves honestly and hurts from trading when his trading partner cheats:

$$g_{BX}^{H}(x) \geq 0 > g_{BXZ}^{C}(x, z) \text{ when } z \text{ is sufficiently large.} \tag{4}$$

Other than the above four assumptions, we allow our gain functions to be of any particular form. Also, they can be realized by the online auction process as well as other trading mechanisms.

Each transaction involves two stages. In the first stage, both the seller and the buyer observe the transaction size x, and the buyer decides whether to proceed with the transaction or to withdraw from it. If the former is the answer, then there is a second stage. Otherwise, both parties receive zero gains. If the buyer decides to proceed, then in the second stage, the seller decides whether to play honestly or to cheat, and subsequently both parties receive their corresponding gains.

Suppose in each transaction, a buyer faces a randomly drawn seller whose type he does not know. Then by the four assumptions, we can easily see that as long as the random distribution of sellers is sufficiently tilted toward the more dishonest types (larger z's), any Nash equilibrium of the stage game of any transaction size results in no trading. On the other hand, folk theorems in the economics literature suggest that when the stage games are repeatedly played, players might use credible threats and rewards to induce each other to adopt otherwise unacceptable strategies in their stage-game plays. However, these theorems require that either the same set of players keep on playing with each other forever, or any long-run player's past behavior is either

publicly known or sufficiently discernible (see e.g., [13], Theorem 5.4 on Page 157, Theorem 5.10 on Page 171, and Theorem 5.11 on Page 196). There is no known folk theorem that readily applies to the situation where any two players have no chance of playing more than once with each other and any player's past history is in no way recorded.

However, in a bare-bone online C2C auction market, buyers have to repeatedly deal with random sellers who have different propensities for cheating and whose true types and past histories are unknown to the buyers. Hence by the above observation, without any proper mechanism, it seems unlikely that the market participants can entice and discipline themselves into mutually beneficial tradings. Therefore, the implementation of feedback systems can be very crucial in such markets.

Probably the ideal feedback mechanism is the one that records each seller's entire history of past behavior. The economics literature on reputation effects often assumes that this mechanism is implementable (e.g., [11]). Due to limited space (memory), however, this ideal version usually cannot be implemented in reality. On the other hand, real online feedback systems try to use limited space to store information which to some extent reflects the past behavior of sellers.

2.2. The reputation feedback system

We view a reputation feedback system as a system where every seller is associated with a score, which is constantly being updated by each of his trading partners in response to the treatment the latter has received. Upon making the assumption that buyers behave fairly rationally, we expect the score to serve as a mirror of a seller's past behavior and to indicate his true type. In order for all sellers' scores to reveal their types, which are ranked from the morally strongest to the morally weakest, we need the scores to fall into an ordered set, wherein one score is always lower than the other given any two scores in the set. By further assuming that a cheated buyer tends to increase his partner's score while a well treated buyer tends to reduce his partner's score, we expect that once trading has gone on long enough to allow all behavior and beliefs to fall into steady states, a seller who is prone to cheating will more likely have a high score.

For example, eBay's feedback system allows buyers to write comments about their trading partners. There are several ways to translate the comments that are stored in a seller's account into a sortable score. We may categorize the comments stored in a seller's account into good comments and bad comments. Then we may let the score of a seller whose account contains n_G number of good comments and n_B number of bad comments be $(-n_G, n_B)$, and let any two arbitrary scores be ordered lexicographically with either $-n_G$ or n_B being the dominant component. Or, we may assign to each comment a real value and make the summation of all values assigned to the comments stored in a seller's account as this seller's score. Here, two scores may be ranked as two real numbers.

In this paper, we make the buyers' feedback writing behavior exogenous so that it is described by certain reasonable random variables. Research has found that the reasonable behavior on the buyers' part can be elicited (see e.g., [18]). In our feedback model, each seller is associated with a score, say w, out of a totally ordered set $\mathcal{W}$. After each trade, w will be updated by the buyer to a random new score $(W^H|w)$ or $(W^C|w)$ depending on whether the buyer has been treated honestly or has been cheated. For brevity, from now on we assume that $\mathcal{W}$ is a subset of nonnegative real numbers. The random variables $(W^H|w)$ and $(W^C|w)$ need to meet two types of requirements.

First, all of them should be increasing in w in some stochastic sense, so that the new score memorizes a seller's past behavior.

Second, at the same w, in some stochastic sense $(W^H|w)$ should be smaller than the degenerate random variable $W \equiv w$ and $(W^C|w)$ should be larger than the degenerate random variable $W \equiv w$, so that a positive correspondence will likely be formed between a seller's type and his score after enough transactions have taken place.

There are several possible ways to mathematically express that one random variable, say W^1, is smaller than another, say W^2: in the likelihood ratio (lr) sense, in the stochastic (st) sense, or in the average sense [22].

From now on, we use $f_{WW}^H(w,w')$ and $f_{WW}^C(w,w')$ to respectively denote the distributions of $(W^H|w)$ and $(W^C|w)$ when they are continuous random variables, and use $p_{WW}^H(w,w')$ and $p_{WW}^C(w,w')$ to respectively denote the probability masses of them.

Here is an example of a stylized feedback system. Each seller's account stores up to $\bar{W}$ comments. A score-w seller is a seller whose account contains $\bar{W}-w$ good comments and w bad comments. If the seller's trading partner has been honestly treated, then with probability p_H the partner will overwrite a randomly-chosen existing comment with a good comment and with probability $1-p_H$, he will do nothing. If the seller's trading partner has been cheated, then with probability p_C, the partner will overwrite a randomly chosen existing comment with a bad comment and with probability $1-p_C$, he will do nothing.

After translating such a description of a buyer's score-updating behavior into the corresponding $p_{\mathrm{WW}}^{\mathrm{H}}(w,w')$'s and $p_{\mathrm{WW}}^{\mathrm{C}}(w,w')$'s, we have

$$p_{\mathrm{WW}}^{\mathrm{H}}(0,0) = 1,$$
$$p_{\mathrm{WW}}^{\mathrm{H}}(w, w-1) = \frac{w}{\overline{W}}p_{\mathrm{H}} \text{ and } p_{\mathrm{WW}}^{\mathrm{H}}(w, w)$$
$$= 1 - \frac{w}{\overline{W}}p_{\mathrm{H}} \text{ for } w = 1, 2, \ldots, \overline{W},$$

$$p_{\mathrm{WW}}^{\mathrm{C}}(w, w) = 1 - \frac{\overline{W}-w}{\overline{W}}p_{\mathrm{C}} \text{ and } p_{\mathrm{WW}}^{\mathrm{C}}(w, w+1)$$
$$= \frac{\overline{W}-w}{\overline{W}}p_{\mathrm{C}} \text{ for } w = 0, 1, \ldots, \overline{W}-1,$$

$$p_{\mathrm{WW}}^{\mathrm{C}}(\overline{W}|\overline{W}) = 1, \tag{5}$$

and that these probabilities at all other unmentioned elements are 0. For this example, we found that the aforementioned two requirements are met in the strongest lr-sense.

Since the situation without a feedback system can be viewed as when the cardinality of $\mathcal{W}$ is merely 1 or when

$$(W^{\mathrm{H}}|w) = (W^{\mathrm{C}}|w) \equiv w,$$

we do not have to separately delineate the decision-making processes under this situation. In later sections, we will use a "0" superscript to signify values pertaining to the without-feedback situation.

2.3. The repeated game setting

Our game is an infinitely repeated game with nature, buyers, and sellers as players. Buyers form the set $[0, 1]$ and sellers form the set $[0, +\infty)$. Each seller is associated with a score $w \in \mathcal{W}$. Nature continuously draws pairs of sellers and buyers to play the stage game introduced earlier in such a way that every buyer $y \in [0, 1]$ gets to play an infinite and countable number of times; every seller $z \in [0, +\infty)$ gets to play an infinite and countable number of times; to a seller in each of his stage game, the type of the paired buyer is known; and to a buyer in each of his stage game, the paired seller seems to be randomly drawn from $[0, +\infty)$ according to distribution $f_{\mathrm{Z}}(z)$. Also, we let $\beta_{\mathrm{B}} \in (0, 1)$ and $\beta_{\mathrm{S}} \in (0, 1)$ be each buyer's and seller's discount factor per stage, respectively.

In every stage game, after the involved players have been decided, nature proceeds to randomly draw the transaction size x from distribution $f_{\mathrm{X}}(x)$ on $[0, +\infty)$. Then, the selected buyer moves by deciding whether to trade with the currently paired seller whose type he does not know and whose score he knows. If the buyer decides

to proceed, the seller gets to move by deciding whether to play honestly or to cheat, and the buyer moves last to update the seller's score according to the random score-updating rule specified in the last subsection. Otherwise, none of the players move. Nature's payoff is not our concern and it has one strategy after all. For sellers and buyers, their payoffs in the current stage are the same as those of the stage game introduced earlier. We introduce different buyer types only to facilitate the aforementioned play frequencies and the actual buyer payoffs do not differentiate over different types.

We use $H_{\mathrm{Bn}}(y)$ to denote the set of all possible histories that buyer y can have at his nth play which consist of his past observations of transaction sizes and opponents' scores along with the past treatments he received. Any of buyer y's mixed strategy is specified by an infinite array of functions: $B(y) = \{p_{\mathrm{Bn}}^{\mathrm{P}}(y, h_{\mathrm{Bn}}(y))|n = 1, 2, \ldots, h_{\mathrm{Bn}}(y) \in H_{Bn}(y)\}$, where $p_{\mathrm{Bn}}^{\mathrm{P}}(y, h_{\mathrm{Bn}}(y))$ stands for the probability that buyer y will proceed to trade in his nth stage game when his history is $h_{\mathrm{Bn}}(y)$. Similarly, we use $H_{\mathrm{Sn}}(z)$ to denote the set of all possible histories that seller z can have at his nth play which consist of his observations of transaction sizes, his own scores, and his opponent's types along with the past treatments he received. Any of seller z's mixed strategy is specified by an array of functions $S(z) = \{p_{\mathrm{Sn}}^{\mathrm{H}}(z, h_{\mathrm{Sn}}(z))|n = 1, 2, \ldots, h_{\mathrm{Sn}}(z) \in H_{\mathrm{Sn}}(z)\}$, where $p_{\mathrm{Sn}}^{\mathrm{H}}(z, h_{\mathrm{Sn}}(z))$ stands for the probability that seller z, when his history is $h_{\mathrm{Sn}}(z)$, will play honestly if given the chance to play in his nth stage game.

Given any initial seller score profile, nature's strategy, buyers' strategy profile $\mathcal{B} = \{B(y)|y \in [0, 1]\}$, and sellers' strategy profile $\mathcal{S} = \{S(z)|z \in [0, +\infty)\}$, buyer y's total discounted expected payoff is

$$
\begin{aligned}
&\overline{g}_{\mathrm{B}}(y, \mathcal{B}, \mathcal{S}) \\
&= \sum_{n=1}^{+\infty} \beta_{\mathrm{B}}^{n-1} E_{h_{\mathrm{Bn}}(y) \in H_{\mathrm{Bn}}(y)} [p_{\mathrm{Bn}}^{\mathrm{P}}(y, h_{\mathrm{Bn}}(y)) \\
&\quad \times \sum_{m=1}^{+\infty} \int_0^{+\infty} \mathrm{d}z E_{h_{\mathrm{Sm}}(z) \in H_{\mathrm{Sm}}(z)} \\
&\quad [f_{\mathrm{Sm|Bn}}(z, h_{\mathrm{Sm}}(z) \mid y, h_{\mathrm{Bn}}(y)) \\
&\quad \times (p_{\mathrm{Sm}}^{\mathrm{H}}(z, h_{\mathrm{Sm}}(z))g_{\mathrm{BX}}^{\mathrm{H}}(x_{\mathrm{Bn}}(y)) \\
&\quad + (1 - p_{\mathrm{Sm}}^{\mathrm{H}}(z, h_{\mathrm{Sm}}(z)))g_{\mathrm{BXZ}}^{\mathrm{C}}(x_{\mathrm{Bn}}(y), z))]], \tag{6}
\end{aligned}
$$

where $E_{a \in A}[\cdot]$ denotes the expectation over a random variable a in its support A, $f_{\mathrm{Sm|Bn}}(z, h_{\mathrm{Sm}}(z)|y, h_{\mathrm{Bn}}(y))$ is the conditional probability density at seller z's mth game with history $h_{\mathrm{Sm}}(z)$ given that he is the opponent of buyer y with history $h_{\mathrm{Bn}}(y)$ in the buyer's nth game

$(\sum_{m=1}^{+\infty} \int_0^{+\infty} \mathrm{d}z E_{h_{Sm}(z) \in H_{Sm}(z)} [f_{Sm|Bn} (z, h_{Sm}(z) | y,$ $h_{Bn}(y))] = 1)$, which is part of $h_{Bn}(y)$. Similarly, seller z's total discounted expected payoff is

$$
\begin{aligned}
\overline{g}_S & (z, \mathcal{B}, \mathcal{S}) \\
= & \sum_{n=1}^{+\infty} \beta_S^{n-1} E_{h_{Sn}(z) \in H_{Sn}(z)} [(p_{Sn}^H(z, h_{Sn}(z)) g_{SX}^H(x_{Sn}(z)) \\
& + (1 - p_{Sn}^H(z, h_{Sn}(z))) g_{SXZ}^C(x_{Sn}(z), z)) \\
& \times \sum_{m=1}^{+\infty} E_{h_{Bm}(y_{Sn}(z)) \in H_{Bm}(y_{Sn}(z))} \\
& [f_{Bm|Sn}(h_{Bm}(y_{Sn}(z))|z, h_{Sn}(z)) \\
& \times p_{Bm}^P(y_{Sn}(z), h_{Bm}(y_{Sn}(z)))]],
\end{aligned}
\tag{7}
$$

where $x_{Sn}(z)$ and $y_{Sn}(z)$ are respectively the transaction size and buyer type in seller z's nth stage game, which are parts of history $h_{Sn}(z)$, and $f_{Bm|Sn}(h_{Bm}(y_{Sn}(z))|z,$ $h_{Sn}(z))$ is the conditional probability density at buyer $y_{Sn}(z)$'s mth game with history $h_{Bm}(y_{Sn}(z))$ given that he is the opponent of seller z with history $h_{Sn}(z)$ in the seller's nth game, which is learned in the Bayesian fashion by the seller.

Given any initial seller score distribution profile and nature's strategy, a strategy profile $(\mathcal{B}^*, \mathcal{S}^*)$ will be a Nash Equilibrium when for any $y \in [0, 1)$, $B^*(y)$ maximizes buyer y's total discounted expected payoff given that other players all play according to $\mathcal{B}^*$ and $\mathcal{S}^*$, and when for any $z \in [0, 1)$, $S^*(z)$ maximizes seller z's total discounted expected payoff given that other players all play according to $\mathcal{B}^*$ and $\mathcal{S}^*$. The profile will further be perfect Bayesian if the above is true when everything starts from any possible play history.

Note that there might be multiple nature strategies that fit the above descriptions. Also, we have not specified in the above how much each player knows about the marginal seller score distribution $f_W(w)$ (An eBay buyer may only pay attention to his current seller's comments or may compare his seller's comments with many other sellers', and an eBay seller may or may not be aware of his fellow sellers' comments). So the above repeated game will not be exactly specified until all these details have been settled. However, we need not consider these details if our focus is on the steady-state behavior of the game.

2.4. Analysis of the repeated game

Since it seems very unlikely to identify Nash or perfect Bayesian strategy profiles for the repeated game, we now proceed to analyze the game when static strategies lead it to a steady state. Once in this state, buyers of different types will have statistically the same experience and therefore are no longer required to be distinguished from each other, and the seller's score distribution profile $f_{W|Z}(w|z)$ will become static and publicly known. Each buyer will use this distribution and the current-opponent-score portion of his history to infer about his current opponent. Also, since a buyer's current action will not impact the profile and his own expected future payoff, the buyer will only be concerned with his stage-game payoff in each transaction. Consequently, sellers will learn nothing about their opponents using their own histories. In addition, different nature strategies will not lead to different long-run expected payoffs.

Now we define $\overline{g_{SZW}}(z, w)$ to be the long-run expected gain for a type-z score-w seller (seller z if he happens to have score w) right before a transaction starts. Then, in a size-x transaction, if the seller has the chance to play and plays honestly, his long-run expected gain will be

$$
\begin{aligned}
\overline{g_{SXZW}^H}(x, z, w) \\
= g_{SX}^H(x) + \beta_S \int_{w' \in \mathcal{W}} \overline{g_{SZW}}(z, w') f_{WW}^H(w, w') \mathrm{d}w',
\end{aligned}
\tag{8}
$$

while if he has the chance to play and cheats, his long-run expected gain will be

$$
\begin{aligned}
\overline{g_{SXZW}^C}(x, z, w) \\
= g_{SXZ}^C(x, z) + \beta_S \int_{w' \in \mathcal{W}} \overline{g_{SZW}}(z, w') f_{WW}^C(w, w') \mathrm{d}w'.
\end{aligned}
\tag{9}
$$

A type-z score-w seller will play honestly with a certain probability $p_{SXZW}^H(x, z, w)$ and cheat with probability $1 - p_{SXZW}^H(x, z, w)$. We may assume that

$$
p_{SXZW}^H(x, z, w) = P_S(\overline{g_{SXZW}^H}(x, z, w) - \overline{g_{SXZW}^C}(x, z, w))
\tag{10}
$$

for a certain function $P_S(u)$ which satisfies $0 \le P_S(u) \le 1$ and

$$
P_S(u + \Delta u) \ge P_S(u) \text{ for } \Delta u \ge 0.
\tag{11}
$$

In each transaction, a buyer believes that the distribution of a seller's score is some $f_W(w)$ and that the conditional distribution of a seller's type with respect to his score is some $f_{Z|W}(z|w)$. So in a size-x

transaction, a buyer will on average gain $g_{BXW}^P(x,w)$ if he proceeds with the transaction, where

$$
\begin{aligned}
g_{BXW}^P&(x,w)\\
&= \int_{z=0}^{+\infty} [p_{SXZW}^H(x,z,w)g_{BX}^H(x)\\
&\quad + (1 - p_{SXZW}^H(x,z,w))g_{BXZ}^C(x,z)]f_{Z|W}(z\mid w)\mathrm{d}z.
\end{aligned}
\tag{12}
$$

There is a certain $p_{BXW}^P(x,w)$ chance that the buyer will proceed with the transaction and therefore there is a $1 - p_{BXW}^P(x,w)$ chance that he will walk away from it. We may assume that

$$
p_{BXW}^P(x,w) = P_B(g_{BXW}^P(x,w))
\tag{13}
$$

for a certain function $P_B(u)$ which satisfies $0 \le P_B(u) \le 1$ and

$$
P_B(u + \Delta u) \ge P_B(u) \text{ for } \Delta u \ge 0.
\tag{14}
$$

For instance, we may use positive real numbers Δ_B and Δ_S to represent respectively the fuzziness of the buyer's and the seller's decision-making processes under the current situation. On one hand, their existence is necessary for the stability of the results; on the other hand, they reflect the difficulty of behaving exactly rationally on the part of the decision makers inside very complex systems. To model the fuzziness, we introduce function $\Lambda(a,b,c)$ so that it equals the medium value of -1, $(a-b)/c$, and $+1$. For instance, we let $\Lambda(4,2,3)=2/3$ and $\Lambda(4,0,3)=1$. Then, we may let

$$
P_S(u) = \frac{1}{2} + \frac{1}{2}\Lambda(u, 0, \Delta_S),
\tag{15}
$$

and

$$
P_B(u) = \frac{1}{2} + \frac{1}{2}\Lambda(u, 0, \Delta_B).
\tag{16}
$$

From the way buyers and sellers make their decisions, we also obtain the transitional kernel $T_{WW|Z}(w,w'|z)$ for $f_{W|Z}(w|z)$:

$$
\begin{aligned}
T_{WW|Z}&(w, w' \mid z)\\
&= \int_{x=0}^{+\infty} [p_{SXZW}^H(x,z,w)f_{WW}^H(w,w')\\
&\quad + (1 - p_{SXZW}^H(x,z,w))f_{WW}^C(w,w')]f_X(x)\mathrm{d}x.
\end{aligned}
\tag{17}
$$

Then, we have

$$
\int_{w' \in \mathcal{W}} f_{W|Z}(w' \mid z)T_{WW|Z}(w',w \mid z)\mathrm{d}w' = f_{W|Z}(w \mid z).
\tag{18}
$$

Now comes the marriage of perceived and real seller distributions. In reality, the score distribution $f_W(w)$ and the seller type distribution conditioned on score $f_{Z|W}(z|w)$ are determined through the following expressions by $f_Z(z)$ and $f_{W|Z}(w|z)$:

$$
f_W(w) = \int_{z=0}^{+\infty} f_{W|Z}(w \mid z)f_Z(z)\mathrm{d}z
\tag{19}
$$

and

$$
f_{Z|W}(z \mid w) = \frac{f_{W|Z}(w \mid z)f_Z(z)}{f_W(w)}.
\tag{20}
$$

The expected gain $\overline{g_{SZW}}(z,w)$ for a type-z score-w seller right before a transaction starts is determined by the following implicit equation:

$$
\begin{aligned}
\overline{g_{SZW}}&(z,w)\\
&= \int_{x=0}^{+\infty} \{ p_{BXW}^P(x,w)[p_{SXZW}^H(x,z,w)\overline{g_{SXZW}^H}(x,z,w)\\
&\quad + (1 - p_{SXZW}^H(x,z,w))\overline{g_{SXZW}^C}(x,z,w)]\\
&\quad + \beta_S(1 - p_{BXW}^P(x,w))\overline{g_{SZW}}(z,w)\}f_X(x)\mathrm{d}x.
\end{aligned}
\tag{21}
$$

Note that the last term on the right-hand side corresponds to the outcome in which the buyer withdraws from the current transaction and therefore involves $\overline{g_{SZW}}(z,w)$ itself (for the next transaction).

Once the equations in this subsection have been solved, all market performances can be calculated. We list the expressions of the most relevant performances in Appendix B.

3. Some discussion

We have established a repeated-game-theoretical model that has the potential to interpret the role of a feedback system in an online market where transactions take place between strangers. Key elements of the model include: the peculiar forms of the stage-game gain functions which condone incentives for sellers to cheat and buyers to shun away from trading; the feedback updating mechanism, exogenously and probabilistically reflecting buyers' feedback writing behavior, which leads each seller's score to two opposite directions depending on his behavior; and the infinite-horizon, multi-seller, multi-buyer, repeated-game setting where every player seeks to maximize his total discounted expected payoff.

We have not been able to make theoretical claims out of the general model. The major difficulties we face involve multiple dimensions and implicit equations. In the finite-horizon version of the game, each player's strategy is implicitly contingent on the seller score

distribution profile, whose number of dimensions is even uncountable; while in the steady-state analysis of the game, we encounter equations where the same functions are involved on both sides.

Naturally, we should expect the following properties from the repeated game model:

$$\overline{g_{\mathrm{B}}} \geq \overline{g_{\mathrm{B}}^{0}}, \tag{22}$$

and

$$\overline{g_{\mathrm{SZ}}}\,(z + \Delta z) - \overline{g_{\mathrm{SZ}}}\,(z) \leq \overline{g_{\mathrm{SZ}}^{0}}\,(z + \Delta z) \tag{23}$$
$$- \overline{g_{\mathrm{SZ}}^{0}}\,(z) \ \text{for} \ \Delta z \geq 0.$$

That is, the feedback system benefits buyers and honest sellers.

In Section 2.2, we have speculated the advent of the tendency for a more honest seller to have a lower score. This tendency will serve as one link in a cycle of benevolent effects that a feedback system can bring to the online markets. When this tendency does materialize, a buyer will, through his personal experience, form a belief that a lower score announces a seller less prone to cheating, and therefore he will be more willing to trade with the owner of such a score. Because of this, a seller with a lower score will have more chances for trading and will probably fare better in the long run. Being aware of this, even when cheating brings more short-term gain, a seller will be more reluctant to do so fearing that an unfavorable updating of his score by the antagonized buyer will jeopardize his future. However, the temptation for short-term gain is still stronger among the morally weaker sellers than among the morally stronger ones. Therefore in the long run, we will still observe the trend that those sellers more prone to cheating do cheat more often and receive higher scores than those less prone to cheating. Now, we have come back to the starting link of the cycle, which means that this chain of effects is self-sustainable.

Therefore, we should expect $(W|z)$ to be increasing in z in some sense and $(Z|w)$ to be increasing in w in some sense which prove that the first link in the above cycle does emerge. We should also expect to have $p_{\mathrm{SXZ}}^{\mathrm{H}}(x, z) \geq p_{\mathrm{SXZ}}^{\mathrm{H0}}(x, z)$ or to have at least $p_{\mathrm{SX}}^{\mathrm{H}}(x) \geq p_{\mathrm{SX}}^{\mathrm{H0}}(x)$ and $p_{\mathrm{SZ}}^{\mathrm{H}}(z) \geq p_{\mathrm{SZ}}^{\mathrm{H0}}(z)$ at the aggregate levels so that a seller will be more reluctant to cheat with a feedback system than without a feedback system.

For the without-feedback situation, it can be shown easily that $p_{\mathrm{SXZ}}^{\mathrm{H0}}(x, z)$ decreases in both x and z due to Inequality Eq. (3) (refer to Eqs. (10) and (15)). Now for the with-feedback situation, we should still expect $p_{\mathrm{SXZ}}^{\mathrm{H}}(x, z)$

to decrease in z so that sellers with higher z will still cheat more than sellers with lower z.

In Appendix C, we give a complete analysis of a simple example. In the example, there are only two types of sellers as well as two types of scores; transactions are all of the same size; and, Eqs. (15) and (16) are adopted as the expressions for $P_{\mathrm{S}}(u)$ and $P_{\mathrm{B}}(u)$, respectively, and it is assumed that $\Delta_{\mathrm{B}} = \Delta_{\mathrm{S}} = 0$. Through the analysis, the benefit of the feedback system is clearly demonstrated.

4. A numerical study

With the help of computers, we conduct a numerical study on the effects brought forth by the stylized feedback system introduced in Section 2.2. We adopt Eqs. (15) and (16) as the basis for the players' decision-making mechanism. Here, all distributions are assumed to be discrete. We assume that transaction sizes x can be 1, 2, ..., $\overline{X}$ and the probability for the realization of any x is $1/\overline{X}$. We assume that the seller types z can be 0, 1, ..., $\overline{Z}-1$ and the distribution of z follows the geometric distribution with a tail cutoff. That is, there is a parameter $R_{\mathrm{Z}} \in [0, 1)$ such that the probability for the realization of z is $(R_{\mathrm{Z}})^{z}(1-R_{\mathrm{Z}})$ for $z=0$, 1, ..., $\overline{Z}-2$ and is $(R_{\mathrm{Z}})^{\overline{Z}-1}$ for $z=\overline{Z}-1$. We let the gain functions be such that $g_{\mathrm{SX}}^{\mathrm{H}}(x)=g_{\mathrm{BX}}^{\mathrm{H}}(x)=x$, $g_{\mathrm{BXZ}}^{\mathrm{C}}(x, z)=x(1-z-2z/(\overline{Z}-1))+C\,(xz)^{\alpha}$, and $g_{\mathrm{SXZ}}^{\mathrm{C}}(x,\ z)=x(1+z)-D(xz)^{\alpha}$ for some positive C and D with $C<D$ and some $\alpha \in [0, 1]$. The rationale for such gain functions is given in Appendix D.

In the first round of the experiment, we let $\overline{X}=10$, $\overline{Z}=10$, $R_{\mathrm{Z}}=0.5$, $C=1.0$, $D=1.5$, $\alpha=0.5$, $\beta_{\mathrm{B}}=0.9$, $\beta_{\mathrm{S}}=0.99$, $\overline{W}=10$, $p_{\mathrm{H}}=0.5$, and $p_{\mathrm{C}}=0.8$. We let $\Delta_{\mathrm{B}}^{0}=\Delta_{\mathrm{S}}^{0}=\Delta^{0}=0.1$ to capture the uncertainty the decision makers face in the without-feedback situation. On the other hand, the iterative process for the with-feedback situation will not converge unless Δ_{B} and Δ_{S} are both large enough. If they are to be the same, the smallest multiple of 0.1 that they need to be is 0.7. In this round, we let $\Delta_{\mathrm{B}}=\Delta_{\mathrm{S}}=0.7$. We present our computational findings in the following.

Finding #1: Buyers are better off with the existence of a reputation feedback system in an online C2C auction market. We obtain $\overline{g_{\mathrm{B}}^{0}} \simeq 3.53$ while $\overline{g_{\mathrm{B}}} \simeq 47.02$. Hence, Inequality Eq. (22) is indeed satisfied. This is consistent with the finding of Bolton, Katok, and Ockenfels [4]. At the same time, we find that *on average a seller is worse off with a feedback system*, in that $\overline{g_{\mathrm{S}}^{0}} \simeq 687.99$ while $\overline{g_{\mathrm{S}}} \simeq 563.15$. Also, the existence of a feedback system makes cheating less attractive. To see how different sellers fare differently, we draw a figure (Fig. 1) of curve $\overline{g_{\mathrm{SZ}}}(z)$ vs. curve $\overline{g_{\mathrm{SZ}}^{0}}(z)$.

From Fig. 1, we see that $\overline{g_{\mathrm{SZ}}}(z)$ rises slower than $\overline{g_{\mathrm{SZ}}^{0}}(z)$. Hence, Inequality Eq. (23) is satisfied. Moreover,

we see in this example that the most honest sellers, the type-0 sellers, gain more when a feedback system is introduced: $\overline{g_{SZ}}(0) \simeq 539.56 > \overline{g_{SZ}^0}(0) \simeq 439.10$.

Finding #2: Buyers are more willing to trade with lower-scored sellers, i.e., more reputable sellers are more likely to sell their items. We draw the curve of $p_{BW}^P(w)$ in Fig. 2. We see that $p_{BW}^P(w)$ has a distinct decreasing trend over w. This observation is in agreement with the findings of Bolton, Katok, and Ockenfels [4] and Resnick and Zeckhauser [19]. The case where $p_{BW}^P(0) < p_{BW}^P(1)$ might be interpreted as sellers become complacent and less willing to play honestly once their scores have reached the best possible level. Hence, a buyer has to be more cautious while dealing with such sellers.

Finding #3: Buyers are more willing to trade when the auction market has a reputation feedback system than without. We note that $p_B^P \simeq 0.973 > p_B^{P0} \simeq 0.886$. This is also verified by Bolton, Katok, and Ockenfels [4].

Finding #4: Sellers are more willing to trade honestly when a reputation feedback system is in place. We find that $p_{SXZ}^H(x, z)$ almost always dominates $p_{SXZ}^{H0}(x, z)$ but for a few x–z pairs. In Figs. 3 and 4, we draw respectively the curves $p_{SX}^H(x)$ vs. $p_{SX}^{H0}(x)$ and $p_{SZ}^H(z)$ vs. $p_{SZ}^{H0}(z)$. Clearly, we have both $p_{SX}^H(x) \geq p_{SX}^{H0}(x)$ for any x and $p_{SZ}^H(z) \geq p_{SZ}^{H0}(z)$ for any z. Again, this result is consistent with the finding of Bolton, Katok, and Ockenfels [4]. Note also that $p_S^H \simeq 0.927 > p_S^{H0} \simeq 0.321$. The fact that a seller becomes more complacent and starts to cheat more once his score reaches 0 is again confirmed by the fact that $p_{SW}^H(0) \simeq 0.196$ is much smaller than $p_{SW}^H(w)$ for all other w's, which are all very close to 1.0.

From Fig. 3, we also see that sellers cheat more with larger sales, while the trend is less apparent when a feedback system is in place. Note also that $p_{SZ}^H(z)$ is almost always decreasing in z. So sellers more prone to cheating will still cheat more with a feedback system.
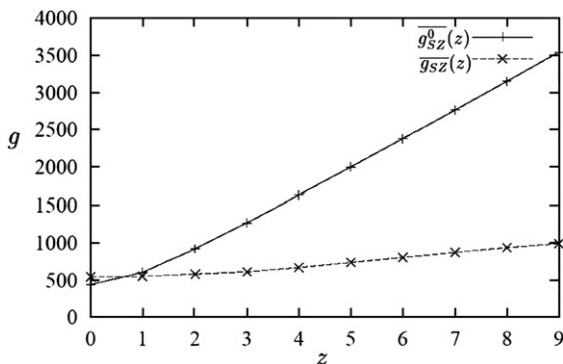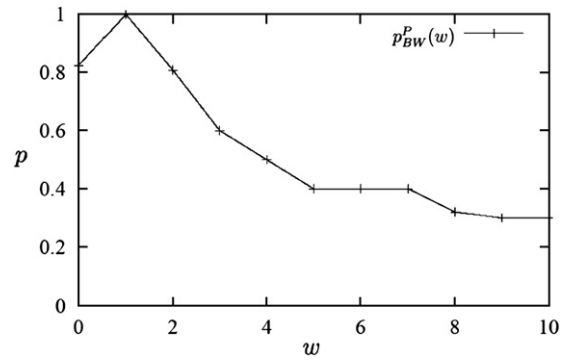

Fig. 2. $p_{BW}^P(w)$.

In our next round of experiments, we consider the important market performance indicators, i.e., g_B, $\overline{g_S}$, $\overline{g_{SZ}}(0)$, p_B^P, and p_S^H, while allowing $\overline{W}$, p_H, p_C, Δ_B and Δ_S to vary and fixing other parameters. In Table 1, we present the results when $\overline{W}$ is fixed at 10, p_H is at either 0.5 or 0.0, and Δ_B and Δ_S are fixed to be 0.9.

From Table 1, regardless of the p_H value, we see that there is a certain level p_{C0} of p_C (between 0.3727 and 0.3728 when $p_H = 0.5$ and between 0.3281 and 0.3282 when $p_H = 0.0$) such that the market performance indicators which favor the buyers, g_B and p_S^H, keep improving as p_C increases from 0 to p_{C0}, and there are no discernible trends for them as p_C increases beyond p_{C0}. Also, these performance indicators are at their best when p_C is just below p_{C0}. Therefore, the intuition that a higher level of buyer unforgiveness induces more honest behavior and more gains for the buyers is only true when the level does not go beyond a certain point.

On the other hand, the average seller's gain $\overline{g_S}$ decreases when p_C increases from 0 to p_{C0}, encounters an upward jump, and then mildly fluctuates when p_C increases beyond p_{C0}. Since the gain of an honest seller $\overline{g_{SZ}}(0)$ is almost unaffected by the change in p_C, the gains of less honest sellers will change in similar fashion
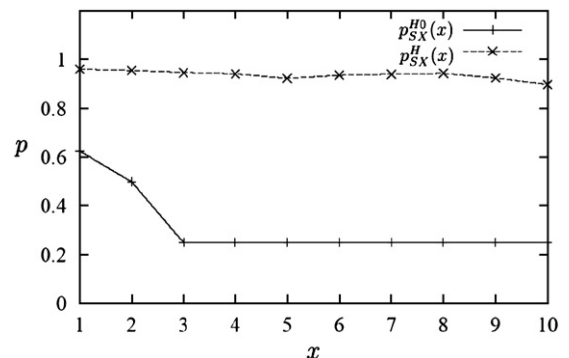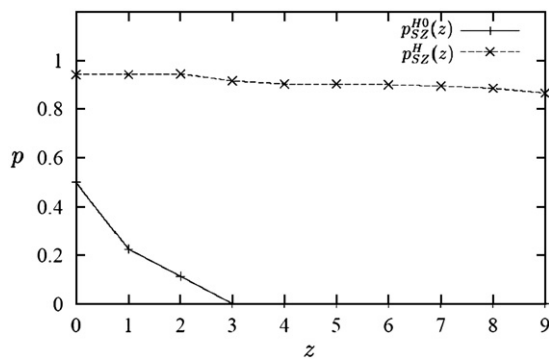

Fig. 1. $\overline{g_{SZ}}(z)$ vs. $\overline{g_{SZ}^0}(z)$.

Fig. 3. $p_{SX}^H(x)$ vs. $p_{SX}^{H0}(x)$.

Fig. 4. $p_{SZ}^{H}(z)$ vs. $p_{SZ}^{H0}(z)$.

Table 2
Results when $\bar{W}=10$, $p_C=0.3$, and $\Delta_B=\Delta_S=0.9$

p_H	g_B	$\overline{g_S}$	$\overline{g_{SZ}}(0)$	p_B^P	p_S^H
0.0	51.22	512.16	543.18	0.931	1.000
0.2	47.36	549.94	532.27	0.963	0.928
0.3	44.28	567.95	535.17	0.962	0.896
0.4	53.69	549.00	550.00	0.992	0.997
0.5	51.60	551.48	550.00	0.984	0.989
0.6	49.82	553.44	550.00	0.974	0.983
0.7	48.94	556.45	550.00	0.972	0.979
0.8	47.78	558.78	550.00	0.971	0.972
0.9	47.19	561.92	550.00	0.970	0.969
1.0	46.89	565.25	550.00	0.971	0.967

as $\overline{g_S}$. Comparing results for the two p_H values, we see that the market does not perform better when p_H changes from 0 to 0.5, and that sellers become more honest when p_H takes the smaller value, i.e., rewarding honest behavior does not seem to benefit the buyers and honest sellers.

Next, we fix p_C at 0.3, keep all other parameters the same as before, and allow p_H to vary. The results are contained in Table 2.

Finding #5: Rewarding an honestly-behaving seller is less effective on promoting market performances than punishing a cheating seller. Note that we have not included the results for $p_H=0.1$ in Table 2 because it requires much larger Δ_B and Δ_S for the iterative procedure to converge. We see from the table that the market performances fluctuate mildly as p_H increases.

Table 1
Results when $\bar{W}=10$, $p_H=0.5$ or 0.0, and $\Delta_B=\Delta_S=0.9$

p_H	p_C	g_B	$\overline{g_S}$	$\overline{g_{SZ}}(0)$	p_B^P	p_S^H
0.5	0.1	39.57	580.08	550.00	0.947	0.920
	0.2	47.88	556.32	550.00	0.969	0.971
	0.3	51.60	551.48	550.00	0.984	0.989
	0.3727	54.06	549.70	550.00	0.995	0.998
	0.3728	43.28	608.31	550.00	0.985	0.885
	0.4	43.79	605.06	550.00	0.985	0.891
	0.6	46.10	591.37	550.00	0.988	0.917
	0.8	45.74	587.78	550.00	0.984	0.919
	1.0	43.80	598.46	549.92	0.983	0.868
0.0	0.1	49.15	491.50	540.63	0.894	0.998
	0.2	49.86	498.58	538.14	0.907	1.000
	0.3	51.22	512.16	543.18	0.931	1.000
	0.3281	50.89	508.93	538.33	0.925	1.000
	0.3282	54.76	547.58	549.99	0.996	1.000
	0.4	54.92	549.18	549.98	0.999	1.000
	0.6	54.56	545.60	549.50	0.992	1.000
	0.8	54.36	543.57	549.39	0.988	1.000
	1.0	53.20	532.02	546.13	0.967	1.000

Last, we fix p_H at 0.0, p_C at 0.3, and Δ_B and Δ_S at 0.9, and allow $\bar{W}$ to vary. The results are recorded in Table 3.

Note that market performance at $\bar{W}=0$ is even worse than when there is no feedback system. But $\bar{W}=0$ is indeed equivalent to there being no feedback system. The difference is that we have allowed more fuzziness here than when we obtained the results for the case without a feedback system.

Finding #6: The existence of even the most primitive reputation feedback system makes a market perform much better. From the results, we see that market performance improves dramatically from $\bar{W}=0$ to $\bar{W}=1$. As $\bar{W}$ increases further, we see that market performance keeps on improving, though not quite as impressively, and after $\bar{W}$ reaches about 14, the performance starts to fluctuate mildly.

In his single-seller model, Dellarocas [9] finds that market efficiency cannot be improved by simply expanding each seller's feedback account. So the above empirical result seemingly contradicts his finding. However, in understanding the difference in these outcomes, we should take into account the difference in our respective assumptions about the number of seller types and their repercussions. In Dellarocas' single-seller model, a buyer uses his trading partner's score only to infer the partner's future behavior; while in our multi-type model, a buyer more importantly uses the score to infer the partner's type. A larger score space may help perform the inference better. Meanwhile, our outcome is also related to our specific assumption about the seller type distribution. A more realistic model may even allow sellers of different types to enter and exit online markets at different rates based on the different rewards from the markets, and let the seller type distribution to gradually converge to an equilibrium. With such a system, it may be that only the most honest sellers eventually remain in the market, not much seller-type inference is needed

Table 3
Results when $p_H=0.0$, $p_C=0.3$, and $\Delta_B=\Delta_S=0.9$

$\overline{W}$	g_B	$\overline{g_S}$	$\overline{g_{SZ}}(0)$	p_B^P	p_S^H
0	2.73	537.85	343.18	0.693	0.310
1	37.73	381.85	456.61	0.712	0.981
2	47.41	474.70	526.30	0.865	0.998
3	49.58	495.75	538.20	0.901	0.999
4	49.53	495.26	539.45	0.900	0.998
5	50.54	505.44	541.44	0.919	1.000
6	50.61	506.14	540.73	0.920	1.000
8	51.12	511.18	543.08	0.929	1.000
10	51.22	512.16	543.18	0.931	1.000
12	54.88	548.76	549.98	0.998	1.000
14	54.89	548.94	549.98	0.998	1.000
16	54.88	548.79	549.96	0.998	1.000
18	54.81	548.05	549.91	0.996	1.000
20	50.97	509.67	538.99	0.927	1.000
25	53.85	538.52	548.63	0.979	1.000
30	53.16	531.58	548.27	0.967	1.000

from the scores, and the score space needs only to be minimal.

Throughout the study, we find that $(W|z)$ is always st-increasing in z and $(Z|w)$ is always st-increasing in w.

5. Concluding remarks

In this paper, we have established a dynamic game-theoretic model for the mechanism of reputation feedback systems in online C2C auction markets. We have also conducted a numerical study based on the game-theoretic model.

Our results confirm the following widely-held and yet previously unproved beliefs that the existence of a feedback system improves the well-being of buyers and enhances their willingness to trade, and it deters dishonest behavior from sellers. Thus with the feedback system, the online C2C auction market as a whole becomes more healthy and attractive to buyers and honest sellers alike. Our findings also confirm that a feedback system does establish a positive correlation between a seller's tendency to cheat and his reputation score; that is, the higher the seller's propensity for cheating, the more likely he will have a high reputation score.

There are interesting findings that may bear implications to real practice: A buyer's willingness to punish ill-behaved sellers has a greater impact on the market performances than their willingness to reward well-behaved sellers. Also, the mere existence of a feedback system, no matter how simple it is, helps to improve the market performance.

We note that there are many viable ways to mathematically model online markets in which feedback systems work, and ours is merely one of them. For instance, in Dellarocas [8], sellers differ not in their cheating propensities but in their efficiencies of delivering quality goods, and also there exists a fee and reward mechanism in the market. Besides leading to the universal conclusion that feedback systems improve performances of markets, the different models inevitably reach conclusions that differ in such subtleties as what the optimal size of a feedback account for each seller should be. We believe, however, that it is out of the scope of the current paper to investigate the extent to which these conclusions hinge upon the specific assumptions as opposed to common features of feedback systems. Future research on feedback systems may shed more light on this issue.

One limitation of this paper is its lack of an explicit theoretical derivation. Even though we believe that there exists some correspondence between the orderings of the buyer score-updating behavior and the orderings of buyer beliefs, due to the complex nature of the underlying problem, we could not tackle it with an explicit solution mathematically. Another limitation is that we rely on the existence of the fuzziness in the decision-making processes to guarantee the repeated game's convergence to a steady state. Although the fuzziness itself has a fairly satisfactory interpretation in real situations, its convergence-guaranteeing threshold has been unpredictable to us.

Further numerical experiments can be conducted based on our existing model. Feedback systems other than the good–bad-comment system can be tested. Different payoff structures and various distribution assumptions suitable for different markets can also be applied and tested under the current model structure.

Appendix A. Symbols Used in the General Model

β_S discount factor per transaction for a seller;

β_B discount factor per transaction for a buyer;

z a seller's prone-to-cheating factor;

$f_Z(z)$ the distribution function of sellers' types;

x a transaction's size;

$f_X(x)$ the distribution function of transaction sizes;

$g_{SX}^H(x)$ a seller's gain in a size-x transaction when he plays honestly;

$g_{BX}^H(x)$ a buyer's gain in a size-x transaction when his trading partner plays honestly;

$g_{SXZ}^C(x, z)$ a type-z seller's gain in a size-x transaction when he cheats;

$g_{BXZ}^C(x, z)$ a buyer's gain in a size-x transaction when his trading partner is of type-z and cheats;

$\mathcal{W}$ the set of all possible scores;

W a generic random score;

w a generic score realization;

$(W^H|w)$ the random score updated from the initial score w by a buyer whose trading partner has played honestly;

$(W^C|w)$ the random score updated from the initial score w by a buyer whose trading partner has cheated;

$f_{WW}^H(w,w')$ the distribution function for $(W^H|w)$;

$f_{WW}^C(w,w')$ the distribution function for $(W^C|w)$;

$h_{Bn}(y)$ a generic history of buyer y in his nth stage game;

$H_{Bn}(y)$ the set of histories of buyer y in his nth stage game;

$h_{Sn}(z)$ a generic history of seller z in his nth stage game;

$H_{Sn}(z)$ the set of histories of seller z in his nth stage game;

$p_{Bn}^P(y,h_{Bn}(y))$ buyer y's probability of proceeding with a trade in his nth stage game when he is with history $h_{Bn}(y)$;

$p_{Sn}^H(z, h_{Sn}(z))$ seller z's probability of playing honestly when given the chance to play in his nth game and when he is with history $h_{Sn}(z)$;

$\mathcal{B}$ the buyer strategy profile;

$\mathcal{S}$ the seller strategy profile;

$\overline{g_B}(y,\mathcal{B},\mathcal{S})$ buyer y's total discounted expected payoff when buyer strategy profile is $\mathcal{B}$ and seller strategy profile is $\mathcal{S}$;

$\overline{g_S}(z,\mathcal{B},\mathcal{S})$ seller z's total discounted expected payoff when buyer strategy profile is $\mathcal{B}$ and seller strategy profile is $\mathcal{S}$;

$f_W(w)$ the aggregated distribution of scores;

$f_{Z|W}(z|w)$ the conditional distribution of seller types while the seller's score is given;

$\overline{g_{SZW}}(z,w)$ the long-run expected gain for a type-z score-w seller right before a transaction starts;

$\overline{g_{SXZW}^H}(x,z,w)$ the long-run expected gain for a type-z score-w seller in a size-x transaction if he opts to play honestly;

$\overline{g_{SXZW}^C}(x,z,w)$ the long-run expected gain for a type-z score-w seller in a size-x transaction if he opts to cheat;

$p_{SXZW}^H(x, z, w)$ the probability that a type-z score-w seller will play honestly in a size-x transaction;

$P_S(u)$ the function being used to calculate $p_{SXZW}^H(x, z, w)$ from $g_{SXZW}^H(x,z,w)$ and $g_{SXZW}^C(x,z,w)$;

$g_{BXW}^P(x,w)$ a buyer's average gain while dealing with a score-w seller in a size-x transaction if the buyer proceeds with the transaction;

$p_{BXW}^P(x,w)$ the probability that a buyer will proceed with a size-x transaction involving a score-w seller;

$P_B(u)$ the function being used to calculate $p_{BXW}^P(x,w)$ from $g_{BXW}^P(x,w)$;

$T_{WW|Z}(w,w'|z)$ the score transition kernel for a type-z seller after each transaction;

$f_{W|Z}(w|z)$ the conditional distribution of seller scores while the seller's type is given.

Appendices B to D will be available from the authors upon request.

References

[1] S. Ba, P. Pavlou, Evidence of the effect of trust building technology in electronic markets: price premiums and buyer behavior, MIS Quarterly 26 (3) (2002) 243–268.

[2] Y. Bakos, C. Dellarocas, Cooperation without Enforcement? A comparative analysis of litigation and online reputation quality assurance mechanism, Proceedings of the 23rd International Conference on Information Systems, Barcelona, Spain, December 15–18, 2002.

[3] P. Battigali, J. Watson, On "reputation" refinements with heterogeneous beliefs, Econometrica 65 (2) (1997) 369–374.

[4] G. Bolton, E. Katok, A. Ockenfels, How effective are online reputation mechanisms? An experimental investigation, Management Science 50 (11) (2004) 1587–1602.

[5] M. Celentani, D. Fudenberg, D.K. Levine, W. Pesendorfer, Maintaining a reputation against a long-lived opponent, Econometrica 64 (3) (1996) 691–704.

[6] M.W. Cripps, J.P. Thomas, Reputation and commitment in two-person repeated games without discounting, Econometrica 63 (6) (1995) 1401–1419.

[7] C. Dellarocas, The digitization of word-of-mouth: promise and challenges of online reputation mechanisms, Management Science 49 (10) (2003) 1407–1424.

[8] C. Dellarocas, Efficiency through feedback-contingent fees and rewards in auction marketplaces with adverse selection and moral hazard, Proceedings of the 4th ACM Conference on Electronic Commerce, June 9–13, 2003. Association for Computing Machinery, San Diego, California, USA.

[9] C. Dellarocas, Reputation mechanism design in online trading environments with pure moral hazard, Information Systems Research 16 (2) (2005) 209–230.

[10] eBay's User Agreement, http://pages.ebay.com/help/policies/user-agreement.html. Available: (2003).

[11] D. Fudenberg, D.K. Levine, Reputation and equilibrium selection in games with a patient player, Econometrica 57 (4) (1989) 759–778.

[12] D. Fudenberg, D.K. Levine, Maintaining a reputation when strategies are imperfectly observed, Review of Economic Studies 59 (3) (1992) 561–579.

[13] D. Fudenberg, J. Tirole, Game Theory, MIT Press, Cambridge, Massachusetts, 1991.

[14] Internet Fraud Watch, http://www.fraud.org/2005_Internet_-Fraud_Report.pdf. Available: (2005).

[15] P. Kollock, The production of trust in online markets, in: E.J. Lawler, M. Macy, S. Thyne, H.A. Walker (Eds.), Advances in Group Processes, vol. 16, JAI Press, Greenwich, Connecticut, 1999, pp. 99–123.

[16] D.M. Kreps, R. Wilson, Reputation and imperfect information, Journal of Economic Theory 27 (2) (1982) 253–279.

[17] P. Milgrom, J. Roberts, Predation, reputation, and entry deterrence, Journal of Economic Theory 27 (2) (1982) 280–312.

[18] N. Miller, P. Resnick, R. Zeckhauser, Eliciting informative feedback: the peer-prediction method, Management Science 51 (9) (2005) 1359–1373.

[19] P. Resnick, R. Zeckhauser, Trust among strangers in internet transactions: empirical analysis of eBay's reputation system, in: M.R. Baye (Ed.), The Economics of the Internet and E-Commerce, Advances in Applied Microeconomics, vol. 11, Elsevier Science, Amsterdam, Holland, 2002.

[20] P. Resnick, R. Zeckhauser, E. Friedman, K. Kuwabara, Reputation systems: facilitating trust in internet interactions, Communications of the ACM 43 (12) (2000) 45–48.

[21] P. Resnick, R. Zeckhauser, J. Swanson, K. Lockwood, The value of reputation on eBay: a controlled experiment, Experimental Economics 9 (2) (2006) 79–101.

[22] M. Shaked, J.G. Shanthikumar, Stochastic Orders and their Applications, Academic Press, Boston, Massachusetts, 1994.

Jian Yang obtained his Ph.D. in Management Science from the University of Texas at Austin in 2000. He is currently an associate professor in the Department of Industrial and Manufacturing Engineering at New Jersey Institute of Technology. His research interest is in the application of operations research techniques to real-life problems emerging from supply chain management, logistics, and electronic commerce. His articles have appeared in *Operations Research, Naval Research Logistics, IIE Transactions, European Journal of Operational Research*, and others.

Xiaorui Hu is an associate professor of Management Information Systems at the John Cook School of Business, Saint Louis University. She received her Ph.D. in Economics from the University of Texas at Austin. Her research focuses on trust related issues in electronic commerce, business-to-business markets, telecommunication market, and culture impact on international business. She has published in *Information Systems Research, Journal of Organizational Computing and Electronic Commerce, IEEE Computer*, and other academic journals.

Han Zhang is an associate professor of Information Technology Management at the College of Management, Georgia Institute of Technology. He received his Ph.D. in Information Systems from the University of Texas at Austin in 2000. His research focuses on online trust and reputation related issues, trusted third parties or intermediaries, the evolution of electronic markets, and online payment methods. He has published in *Information Systems Research, Journal of Management Information Systems, Decision Support Systems*, and other academic journals.