

Efficiency and Robustness of eBay-like Online Feedback Mechanisms in Environments with Moral Hazard

Chrysanthos Dellarocas *

23rd January 2003

Abstract

This paper offers a systematic exploration of online feedback mechanism design issues in trading environments with imperfect monitoring of a seller's effort level and two possible transaction outcomes (corresponding to "high" and "low" quality respectively), one of which has no value to buyers. The objective of feedback mechanisms in such settings is to induce sellers to exert high effort and, therefore, to maximize the probability of high quality outcomes. I study a practically significant family of mechanisms that resemble the one used by online auction house eBay. These *eBay-like* mechanisms solicit ratings of transaction outcomes as either *positive* or *negative* and publish the sums of ratings posted by buyers on a seller during the N most recent periods. My analysis finds that eBay-like mechanisms can induce high average levels of cooperation that remain stable over time. Surprisingly, their efficiency cannot be improved by summarizing larger numbers of ratings or by publishing a seller's detailed feedback history. I further examine the robustness of these mechanisms to incorrect or incomplete feedback as well as to strategic changes of online identities. The theoretical outcomes predicted by this paper are consistent with empirical observations and offer new, theory-backed, explanations to hitherto poorly understood phenomena such as the remarkably low fraction of negative feedback on eBay. Finally, they provide concrete suggestions on how eBay's current mechanism can be improved.

1 Introduction

Online reputation mechanisms (Dellarocas, 2002; Resnick, et. al., 2000) are emerging as a promising approach for building trust and inducing cooperation in online trading environments where more

*Address: MIT Sloan School of Management, 50 Memorial Drive, Room E53-315, Cambridge MA 02142-1347, e-mail: dell@mit.edu.

established methods of social control (such as state regulation or the threat of litigation) are often difficult or too costly to implement. Their growing popularity motivates rigorous research in better understanding how these systems affect behavior in communities where they are introduced and how they can be designed to achieve maximally efficient outcomes.

The use of word-of-mouth networks as a basis for social control is arguably as old as society itself (Greif, 1993; Klein, 1997; Milgrom, North and Weingast, 1990). Nevertheless, the advent of the Internet is adding some significant new dimensions to this age-old concept (Dellarocas, 2002). Most important among them is the ability to systematically design and control such networks through the use of properly architected information systems (feedback mediators). Feedback mediators specify who can participate in feedback communities, what type of information is solicited from participants, how it is aggregated and what type of information is made available to them about other community members. They enable online community operators to exercise precise control over a number of parameters that are difficult or impossible to influence in brick-and-mortar settings. For example, feedback mediators can replace detailed feedback histories with a wide variety of summary statistics, apply filtering algorithms to eliminate outlier or suspect ratings, control the initial state of feedback profiles of new members, etc.

Through the use of information technology, what had traditionally fallen within the realm of the social sciences is, to a large extent, being transformed to an engineering design problem. The potential to engineer social outcomes through the introduction of carefully crafted information systems is opening a new chapter on the frontiers of information systems research. Understanding the full space of design possibilities and determining what designs work best in what settings is one of the most interesting research challenges introduced by these new systems.

This paper offers a systematic exploration of online feedback mechanism design issues in settings with two seller effort levels, imperfect monitoring of a seller's action, and two possible outcomes, corresponding to "high" and "low" quality respectively. Under the assumption that "low" quality outcomes have no value to buyers, sellers will always promise to deliver high quality. However, they will also be tempted to go back on their promise and exert low effort once they have received payment from the buyer. The objective of a feedback mechanism in such settings is to *discipline* sellers, that is, to induce them to exert high effort and, therefore, to maximize the probability of the high quality outcomes promised to buyers. Although stylized, my setting captures the essential properties of a large number of real-life trading environments, ranging from online purchasing to professional services.

I show that the ability of online feedback mediators to control the amount and type of feedback

information they make publicly available permits the construction of simple feedback mechanisms capable of inducing high average levels of cooperation that remain stable over time. More specifically, I study a family of mechanisms that:

- encourage buyers to rate transaction outcomes as either “positive” or “negative”
- hide detailed feedback histories from public view
- publish a statistic that roughly corresponds to the sum of negative ratings posted by buyers on a seller during the N most recent transactions.

These mechanisms are practically important because they resemble aspects of eBay’s well-known feedback mechanism. I will therefore refer to them as *eBay-like* mechanisms in the rest of the paper. In addition to their theoretical interest, the results of the paper, therefore, have important implications for eBay and other online communities that operate similar mechanisms.

Section 2 of the paper introduces the model. Section 3 analyzes the equilibrium outcomes induced by eBay-like mechanisms under the assumption that all sellers are rational and their payoffs are common knowledge. Previous attempts to develop game-theoretic models of online feedback mechanisms exogenized the behavior of either the buyers (Houser and Wooders, 2000) or the seller (Livingston, 2002). This is the first analysis I am aware of that endogenizes the behavior of both buyers and sellers. My analysis results in three major findings:

- If buyer valuations of high quality are sufficiently high (relative to the cost of high effort), eBay-like mechanisms induce high average levels of cooperation that remain stable over time. Furthermore, the buyer and seller strategies that maximize cooperation have a particularly simple stationary form
- eBay-like mechanisms incur small efficiency losses relative to the first-best case; however, these losses cannot be improved by *any* mechanism that publishes a seller’s past feedback history (or any truncation thereof)
- Surprisingly, the maximum efficiency attainable through the use of eBay-like mechanisms is independent of the number of ratings N summarized by the mechanism: mechanisms that only publish the *single* most recent rating posted for a seller are just as efficient as mechanisms that publish summaries of *arbitrarily large* numbers of a seller’s recent ratings

Section 4 examines the robustness of eBay-like mechanisms to the presence of incorrect or incomplete feedback, as well as to strategic manipulation of online identities. These contingencies are particularly important in large-scale, heterogeneous online environments and therefore essential to incorporate into the analysis of systems that are intended for use in such settings. The principal results can be summarized as follows:

- The efficiency induced through eBay-like mechanisms in the presence of incomplete feedback submission depends on the mechanism’s policy regarding missing feedback; the policy that maximizes efficiency is to treat missing ratings as positive ratings¹
- Under such a “no news is good news” policy, incomplete feedback submission *does not* lower the maximum seller payoff attainable through the mechanism (however, it *does* raise the minimum ratio of valuation to cost necessary in order for the most efficient equilibrium to obtain)
- eBay-like mechanisms are vulnerable to sellers that can costlessly disappear and re-enter an online community under new identities (following, say, a cheating incident); this vulnerability can be removed at some efficiency loss by setting the initial state of the feedback profile of newcomer sellers so that it corresponds to the “worst” possible reputation.
- The efficiency loss associated with preventing easy name changes is minimized when $N = 1$. Interestingly, in environments where players can costlessly change their identities, the simplest eBay-like mechanisms, i.e. ones that publish the single most recent rating, are *strictly more efficient* than mechanisms that summarize larger numbers of ratings.

Section 5 compares the outcomes predicted by our theoretical analysis to some of the results of empirical studies on eBay. I find a remarkable consistency between the predictions of theory and the findings of empirical work and offer new, theory-backed, explanations to hitherto poorly understood phenomena such as the remarkably low fraction of negative feedback on eBay. Finally, Section 6 summarizes the implications of the results of this paper for eBay and other online marketplaces and discusses possible extensions of this work.

2 eBay-like feedback mechanisms in settings with moral hazard

This section presents a model that will allow us to study the properties of an important class of feedback mechanisms (that resemble the one used by online auction house eBay) in settings with

¹eBay currently ignores missing feedback. My analysis suggests that eBay can increase the efficiency induced by its mechanism by adopting a policy where missing feedback is, by default, counted as positive feedback.

moral hazard and two possible outcomes.

2.1 The setting

The setting involves a single long-run player (“seller”) who in each period provides one unit of a product or a service (“good”) to one of multiple short-run players (“buyers”). The good is either “high quality” or “low quality”, but only high quality is acceptable to the buyers. Following receipt of payment, the seller can exert either “high effort” (“cooperate”) or “low effort” (“cheat”). The buyer *privately* observes the quality of the good delivered, but not the effort exerted by the seller. Moral hazard is introduced because high effort is costlier to the seller, who can reduce his costs by failing to exert high effort, providing the buyer with a good of lower expected quality.

More formally, I analyze a setting with a monopolist seller who each period offers for sale a single unit of a good to m buyers. Buyer i has valuation w_i for a high quality good and all buyers value a low quality good at zero. Buyer lifetime is exactly one period and in each period the m buyers are drawn from the same probability distribution, thus buyer valuations are independent and identically distributed within and across periods. There are an infinite number of periods and the seller has a period discount factor δ reflecting the time value of money, or the probability that the game will end after each period. Seller effort determines the probability that the good provided will be perceived by the buyer as being low quality: if the seller exerts low effort, the good will be of low quality with probability β , whereas if the seller exerts high effort he will incur an additional cost c and the good will be of low quality with a smaller probability α ($\alpha < \beta$). The seller’s objective is to maximize the present value of his payoffs over the entire span of the game, while the buyers’ objective is to maximize their short-term (stage game) payoff.

In each period a mechanism is used to allocate the good among the m buyers by determining the buyer that receives the good and the price she pays to the seller. Without loss of generality we assume that buyers are indexed according to their valuations ($w_1 \geq w_2 \geq \dots \geq w_m$). Furthermore, we assume that a second price Vickrey auction is used to award the good to the buyer with the highest valuation w_1 for a high quality good. The winning bidder pays a price equal to the -second-highest bid G ; the valuation of the second-highest bidder for a high quality good is w_2 .

While stylized, the above setting captures the essential properties of a large number of important real-life economic settings, ranging from the provision of professional services, to online purchasing and auctions like eBay. In professional services (medical consultations, auditing, construction projects, etc.) there are well defined standards of high quality service and the uncertainty is focused on

whether the provider will adhere to those standards or try to “cut corners”. In mail order or online purchasing the moral hazard is focused on whether, following receipt of payment, the seller will provide a good of the quality advertised.

My formulation assumes that all sellers have identical costs and conditional probabilities of outcomes given effort. This assumption is reasonable in online retail environments where seller actions are simple (ship/not ship), cost structures relatively standard, and conditional probabilities of transaction outcomes given a seller’s action primarily determined by factors (such as the probability of damage or delay during shipping, the probability that a buyer misunderstood a seller’s item description, etc.) that are exogenous to a seller’s innate abilities

2.2 eBay-like feedback mechanisms

In the above setting, I consider a family of feedback mechanisms that allow buyers to rate the seller based on the quality of the good received. Buyers report the outcome of a transaction as either “positive” or “negative”, with positive ratings indicating that a high quality good was received, and negative ratings indicating low quality. The mechanisms hide a seller’s detailed ratings history and provide buyers with a summary of that seller’s most recent ratings. Specifically, the only information available to buyers is a statistic x that is *approximately* equal to the total number of negative ratings posted on the seller during the most recent N transactions, where N is a parameter under the control of the mechanism designer. These mechanisms attempt to capture the essence of online feedback tracking systems, such as eBay’s “ID Card”, which summarizes ratings received during the most recent six-month period (Figure 1). I will, therefore, refer to them as *eBay-like feedback mechanisms* in the rest of the paper².

If N remains constant for the duration of the game, a seller’s feedback profile is completely characterized by the number of negative ratings x in the current “time window”. eBay-like feedback mechanisms initialize and update a seller’s feedback profile according to the following procedure:

- *Initialization*: At the beginning of the game, the mechanism creates an unordered set W of cardinality N . W is initialized to contain x_0 negative reports and $N - x_0$ positive reports, where x_0 is a parameter under the control of the mechanism designer. The seller’s feedback profile is initialized to x_0

²The reader should not view *eBay-like* mechanisms as an attempt to create a precise model of eBay’s *current* mechanism, but rather, as useful parameterized abstractions that capture the essential properties of that family of mechanisms for the purpose of suggesting ways in which eBay’s current mechanism can be improved.

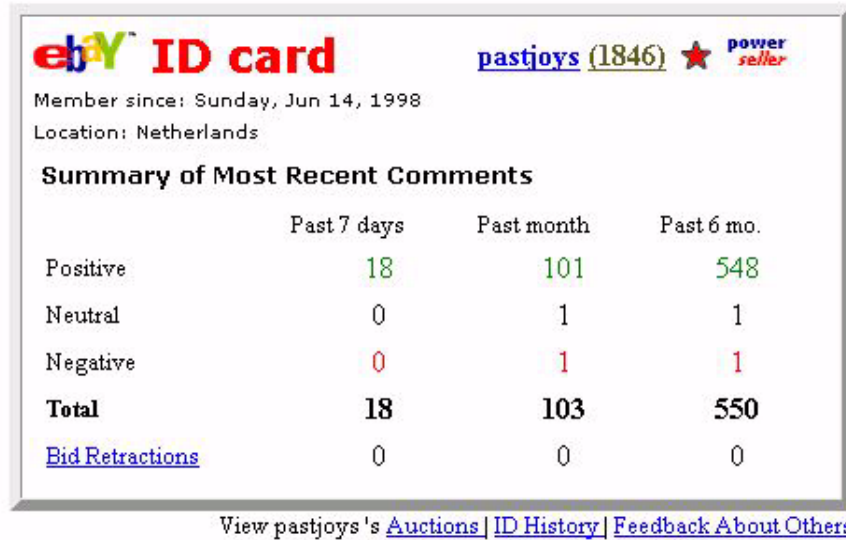


Figure 1: eBay's ID Card

- *Update:* At the end of every period, the mechanism selects a report $r' \in W$ at random and replaces it with the report (“+” or “-”) submitted by the buyer in the current period. It then sets the seller’s feedback profile equal to the new sum x of negative reports in the updated set W

It is easy to see that the above profile updating procedure can be equivalently expressed as a function τ of the current profile x and report (“+” or “-”) as follows:

$$\tau(x, +) = \begin{cases} x & \text{with probability } 1 - x/N \\ x - 1 & \text{with probability } x/N \end{cases}$$

$$\tau(x, -) = \begin{cases} x + 1 & \text{with probability } 1 - x/N \\ x & \text{with probability } x/N \end{cases}$$

The objective of the random replacement algorithm is to eliminate whatever advantages long-run sellers may have from knowing the exact sequence of past ratings (which buyers ignore since I assumed that the only information that is available to them is the statistic x). Proposition 3.2

1. Seller offers a single unit of a good, promising to deliver a high quality good (as there is no demand for a low quality good).
2. System publishes the seller's current feedback profile x , roughly corresponding to the number of negative ratings posted on the seller in the N most recent transactions.
3. Buyers bid their expected valuations for the good in a second price Vickrey auction; the winning bidder pays G , which is the second-highest bid; we denote by w_1 and w_2 the respective valuations for a high quality good of the winning bidder and the second-highest bidder.
4. Seller decides whether to exert high effort at cost c , or low effort at cost 0, with corresponding probabilities that the resulting good is of low quality being α and β ($\alpha < \beta$).
5. Buyer receives the good, experiences its quality, and realizes the corresponding valuation w_1 for a high quality good or 0 for a low quality good. Buyer reports on the quality of the good received to the system, and the feedback profile of the seller is updated accordingly.

Table 1: Stage game of repeated bilateral exchange game studied in this paper.

shows that this simple algorithm succeeds in inducing stationary equilibria in which seller payoffs are as high as in settings where the exact sequence of past ratings is known to *all* players.

Table 1 summarizes the stage game of the repeated bilateral exchange game that serves as the setting of this work. Our objective in the rest of this paper is to study the behaviors induced by eBay-like feedback mechanisms in the above setting and to explore the effects of mechanism parameters N and x_0 on efficiency and robustness.

3 Equilibrium play and payoffs under complete information

This section derives the equilibrium strategies and maximum expected auction revenue and seller payoffs induced by eBay-like feedback mechanisms under the assumption that all buyers and sellers are perfectly rational and their respective payoffs are common knowledge.

3.1 Derivation of equilibrium play and payoffs

Let $s(x, t, h(t)) \in [0, 1]$ denote the seller's strategy in period t , equal to the probability that the seller will cooperate (i.e., exert high effort) in period t if the past history of play is $h(t)$ and his

current feedback profile contains $x \in \{0, \dots, N\}$ negative ratings at the beginning of the period³. I will restrict the seller to stationary strategies, where $s(x, t, h(t))$ does not depend on past history or time, and thus $s(x, t, h(t)) \equiv s(x)$ ⁴. Let $\mathbf{s} = [s(0), \dots, s(N)]$ denote the seller's stationary strategy vector.

Given the seller's strategy, short-term buyers simply play the corresponding stage-game static best response. Since they compete with each other on a Vickrey auction, each buyer's optimal action in each period will be to bid an amount equal to her expected valuation

$$G_i(x, \mathbf{s}) = \{s(x)(1 - \alpha) + [1 - s(x)](1 - \beta)\} w_i = [s(x)(\beta - a) + (1 - \beta)] w_i \quad (1)$$

resulting in expected auction revenue for that period

$$G(x, \mathbf{s}) = [s(x)(\beta - a) + (1 - \beta)] w_2 \quad (2)$$

where w_2 is the second highest bidder's valuation of a high quality good. The seller's corresponding current period payoff is:

$$h_s(x, \mathbf{s}) = G(x, \mathbf{s}) - s(x)c = [s(x)(\beta - a) + (1 - \beta)] w_2 - s(x)c \quad (3)$$

The expected surplus for the winning bidder is

$$h_b(x, \mathbf{s}) = [s(x)(\beta - a) + (1 - \beta)] (w_1 - w_2) \quad (4)$$

where w_1 is the winning bidder's valuation of a high quality good.

The following analysis assumes that buyers always submit ratings that truthfully reflect their quality observations. From a theoretical perspective this can be weakly justified if we make the assumption that buyers only transact with a given seller once (an assumption that is quite reasonable in

³The past history of play includes the bid amounts, seller actions, and buyer ratings for all past periods and the bid amounts for the current period.

⁴In Proposition 3.1 I show that the equilibrium seller payoffs attainable through stationary strategies cannot be improved by any other strategy that makes use of his (private) knowledge of the past history of play. The seller, thus, does not need to consider more complex strategies.

If seller cooperates:	If seller cheats:
$x + 1$ with probability $\alpha(1 - x/N)$	$x + 1$ with probability $\beta(1 - x/N)$
x with probability $\alpha x/N + (1 - \alpha)(1 - x/N)$	x with probability $\beta x/N + (1 - \beta)(1 - x/N)$
$x - 1$ with probability $(1 - \alpha)x/N$	$x - 1$ with probability $(1 - \beta)x/N$

Table 2: Transition probabilities of seller’s feedback profile if his current state is x

large-scale electronic markets). Buyers are then indifferent between truthful reporting, untruthful reporting and no reporting. In reality, however, submission of online ratings incurs a small cost associated with the time required to log on to the feedback site and fill the necessary feedback forms. Fortunately, it is not difficult to devise a side payment mechanism that provides buyers with strict incentives to always rate truthfully (see Kandori and Matsushima, 1998; Miller, Resnick and Zeckhauser, 2002). Such a mechanism can be easily combined with the mechanism I present in this paper⁵.

Since a seller’s choice of effort level takes place *after* payment for the current period has been received, the seller’s objective is to select \mathbf{s} so as to maximize the present value of his payoff in the remaining game. Let $U(x, \mathbf{s})$ denote the seller’s expected future payoff immediately *after* he receives payment for the current period if his current feedback profile contains x negatives.

If the seller exerts high effort in the current period, he incurs an immediate cost c ; the resulting quality of the good will be perceived as high with probability $1 - \alpha$ and low with probability α . If, on the other hand, he exerts low effort he incurs no immediate cost but the probability of high (low) quality becomes $1 - \beta$ (β).

Under the assumption of complete and truthful reporting, it is easy to see that the profile updating function τ described in Section 2.2 gives rise to the feedback profile transition probabilities listed in Table 2.

Given the transition probabilities of Table 2, the expected future payoff is

$$U_{coop}(x, \mathbf{s}) = -c + \delta[\alpha(1 - \frac{x}{N})V(x + 1, \mathbf{s}) + [\alpha \frac{x}{N} + (1 - \alpha)(1 - \frac{x}{N})]V(x, \mathbf{s}) + (1 - \alpha)\frac{x}{N}V(x - 1, \mathbf{s})] \quad (5)$$

⁵Of course, even a side-payment mechanism that induces truth telling cannot eliminate the possibility of buyer mistakes or the existence of irrational buyers who “forget” to leave feedback. In Section 4 I study the implications of non-strategic feedback errors (trembles) and incomplete feedback submission.

if the seller cooperates

$$U_{cheat}(x, \mathbf{s}) = \delta[\beta(1 - \frac{x}{N})V(x+1, \mathbf{s}) + [\beta\frac{x}{N} + (1-\beta)(1 - \frac{x}{N})]V(x, \mathbf{s}) + (1-\beta)\frac{x}{N}V(x-1, \mathbf{s})] \quad (6)$$

if the seller cheats, and

$$U(x, \mathbf{s}) = s(x)U_{coop}(x, \mathbf{s}) + [1 - s(x)]U_{cheat}(x, \mathbf{s}) \quad (7)$$

if he follows a mixed strategy. In the above equations $V(x, \mathbf{s}) = G(x, \mathbf{s}) + U(x, \mathbf{s})$ denotes the seller's expected future payoff at the *beginning* of a period where the seller's feedback profile contains x negatives.

In this setting, a strategy \mathbf{s} is an equilibrium strategy if and only if it satisfies the incentive compatibility constraints:

$$\begin{aligned} s(x) = 0 &\Rightarrow U_{coop}(x, \mathbf{s}) \leq U_{cheat}(x, \mathbf{s}) \\ 0 < s(x) < 1 &\Rightarrow U_{coop}(x, \mathbf{s}) = U_{cheat}(x, \mathbf{s}) \text{ for all } x = 0, \dots, N \\ s(x) = 1 &\Rightarrow U_{coop}(x, \mathbf{s}) \geq U_{cheat}(x, \mathbf{s}) \end{aligned} \quad (8)$$

Not surprisingly, this game has multiple equilibrium strategies. In the rest of the paper, we will focus our attention on the "optimal" equilibrium strategy \mathbf{s}^* that maximizes the seller's expected discounted lifetime payoff $V(x_0, \mathbf{s})$, where x_0 is the initial state of the feedback profile for new sellers. For discount factors close to one this strategy also maximizes a buyer's average per-period surplus and, therefore, the social welfare⁶. This optimal strategy \mathbf{s}^* is the solution to the $(N+1)$ -dimensional constrained optimization problem:

$$V(x, \mathbf{s}^*) \geq V(x, \mathbf{s}) \text{ for all } \mathbf{s} \in [0, 1]^{N+1} \text{ and all } x = 0, \dots, N, \text{ subject to the constraints (8).}$$

It turns out that the solution of the above problem has a particularly simple closed form. Let $\rho = w_2/c$; ρ is the ratio of the valuation of a high quality good to the cost of high effort and is also a rough measure of the profit margin of a fully cooperating seller. The following proposition summarizes the seller's optimal strategy:

⁶To see this, from (3), if $(\beta - \alpha)w_2 > c$, a seller's stage game payoff is a linearly increasing function of his probability of cooperation $s(x)$. Therefore, for discount factors close to one the maximum seller discounted lifetime payoff corresponds to higher average levels of cooperation. From (4) the winning bidder's surplus is also a linearly increasing function of the seller's probability of cooperation. Therefore, as the discount factor tends to one, the strategy profile that maximizes the seller's payoff also maximizes the buyer's average per-period surplus.

Proposition 3.1:

1. If $\rho < \beta/(\beta - \alpha)^2$ then:

- (a) the seller's optimal stationary strategy is $\mathbf{s}^* = [s(x) = 0, x = 0, \dots, N]$: always exert low effort.
- (b) the expected single-period auction revenue is given by $G(\mathbf{s}^*) = (1 - \beta)w_2$
- (c) the seller's payoff is equal to $V(\mathbf{s}^*) = (1 - \beta)w_2/(1 - \delta)$ independently of the initial state of his feedback profile.

2. If $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$ then:

(a) the seller's optimal stationary strategy is:

$$\mathbf{s}^* = [s(x) = 1 - x(1 - \delta + \frac{\delta}{N})\frac{c/w_2}{\delta(\beta - \alpha)^2}, x = 0, \dots, N]$$

(the seller always exerts high effort if he has zero negative ratings, otherwise he follows a mixed strategy in which the probability of cooperation is a linearly decreasing function of the current number of negative ratings in his profile.)

(b) the expected single-period auction revenue is a linearly decreasing function of the current number of negative ratings in the seller's profile and given by:

$$G(x, \mathbf{s}^*) = (1 - \alpha)w_2 - x(1 - \delta + \frac{\delta}{N})\frac{c}{\delta(\beta - \alpha)}$$

(c) the seller's payoff is a function of his initial profile state and equal to:

$$V(x_0, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right] - x_0 \frac{c}{\delta(\beta - \alpha)}$$

3. The above payoffs constitute an upper bound on the payoffs attainable in any sequential equilibrium of the game; the seller, therefore, cannot improve his payoffs by considering non-stationary strategies.

Proof: See appendix.

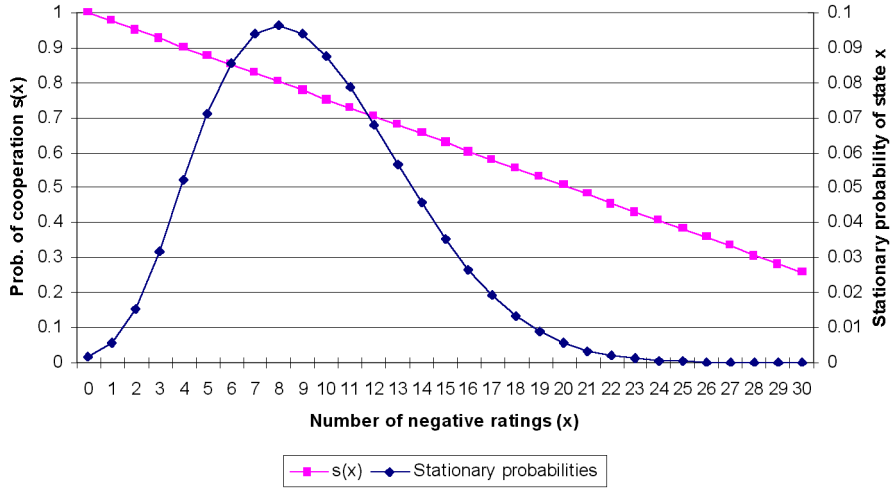


Figure 2: Equilibrium probabilities of seller cooperation and stationary probabilities of having x negatives that maximize seller payoff for $\delta = 0.999, \alpha = 0.1, \beta = 1, \rho = 1.5$ and $N = 30$

3.2 Discussion

Proposition 3.1 states that, if $\rho = w_2/c$ is too small, then the feedback mechanism fails to sustain cooperation, whereas if ρ is large enough, the most efficient equilibrium induces the seller to cooperate fully as long as he has no negative ratings in his profile and otherwise to follow a mixed strategy in which the probability of cooperation is a linearly decreasing function of the number of negative ratings. Knowing this, buyers also bid amounts that decrease linearly with the number of negative ratings in the seller's profile. Figure 2 depicts some aspects of the most efficient equilibrium for an illustrative set of model parameters.

The intuition behind Proposition 3.1 is the following: From (3) it is easy to see that, if $(\beta - \alpha)w_2 > c$ (a condition that always holds if $\rho > [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$) a seller's profit from a single transaction is an increasing function of $s(x)$, where $s(x)$ is the probability that the seller will cooperate during periods when his feedback profile has x negatives. From (4) we see that buyer surplus also increases with $s(x)$. It is thus to everyone's benefit to cooperate as much as possible. Unfortunately, sellers decide whether to cooperate *after* they receive payment and then they always have a short-term gain equal to c if they cheat. Therefore, the only way that a seller will credibly cooperate following receipt of payment is if there is a longer-term loss for him associated with cheating. The only consequence of cheating in this game is a higher probability of transitioning to a state with more negative ratings and the only way that a seller can have a lower payoff in such a state is by cooperating with lower

probability (because, expecting this, buyers will then place lower bids). Therefore, a seller can give himself incentives to cooperate during periods when his feedback profile has few negatives by “promising” to cooperate less (effectively “punishing himself” by doing so) if he accumulates more negatives. Proposition 3.1 shows that, if $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$, such an approach indeed results in optimal seller payoffs for $s(x) = 1 - x(1 - \delta + \frac{\delta}{N})\frac{c/w_2}{\delta(\beta - \alpha)^2}$: full cooperation when $x = 0$ and progressively lower probability of cooperation as the number of negative ratings grows.

The condition $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$ expresses the fact that, for the feedback mechanism to succeed in inducing *any* amount of cooperation, the buyers’ valuation of high quality must be high enough (relative to cost of exerting high effort) so that discounted future payoffs from sustained cooperation are greater than short-term wealth increases obtained from cheating. This seems to be a general property of reputation mechanisms, first pointed out in (Klein and Leffler, 1981) and explored more formally in (Shapiro, 1983). The condition on ρ limits the applicability of such mechanisms to environments where the profit margin of cooperating sellers is sufficiently high so that the promise of future gains offsets the short-term temptation to cheat.

An attractive property of eBay-like feedback mechanisms is that they allow players to maximize their payoffs by relying on simple stationary strategies; players need not even consider more complex strategies (or the possibility that other players may use such strategies). This property makes these mechanisms especially well suited to large-scale online trading environments, as it makes it easy even for relatively inexperienced and naïve traders to perform well without having to perform an excessively complicated analysis.

One final observation concerns the applicability of the above analysis to competitive environments. Although the setting of Section 2 involves a monopolist seller, I conjecture that the results of this paper will not substantially change in settings where, in each period, several sellers offer competing auctions for the same (or for substitute) goods and each buyer elects to bid in exactly one of these auctions. This is easy to see if we assume that buyers have identical valuations or that the number of buyers is very large. In both cases I can assume that $w_1 \approx w_2$. From equation (4) the expected surplus for the winning bidder is then equal to zero and independent of the seller’s current feedback profile state x . If buyers have a weak preference for winning an auction with expected surplus zero, when faced with a number of competing auctions by sellers with different reputations, they will attempt to maximize their probability of winning by selecting one of these auctions at random⁷. A seller’s feedback profile thus does not affect the average number of bidders or the distribution of the bidder valuations. Each seller can, therefore, perform the preceding analysis as if he were a monopolist.

⁷Randomization is a symmetric equilibrium strategy for all buyers in this case.

3.3 Efficiency considerations

In the presence of noise, even fully cooperating sellers will eventually accumulate negative ratings and will, therefore, transition to periods of partial cooperation. This unfortunate property results in efficiency losses relative to the first-best case. In our setting, the first-best case corresponds to environments where sellers can credibly commit to full cooperation. Their payoff would then be equal to $V_{first-best} = [(1 - \alpha)w_2 - c] / (1 - \delta)$. According to Proposition 3.1, if $\rho > [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$ the maximum seller payoff, attainable when the seller starts the game with $x_0 = 0$ negative ratings (i.e. a “clean record”), is equal to

$$V(0, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right]$$

We see that eBay-like feedback mechanisms incur an efficiency loss equal to $\alpha c / (1 - \delta)(\beta - \alpha)$ relative to the first-best case.

A remarkable result is that, under complete information, this efficiency loss cannot be improved by any mechanism that publishes a seller’s entire feedback history (or any truncation thereof). More specifically, the maximum seller payoff achievable in stationary strategies through the use of eBay-like feedback mechanisms is equal to the maximum seller payoff achievable in *any sequential equilibrium* of a perturbation of the above repeated game in which there is no state variable x but instead, (a) buyers submit public reports (b) the entire public history of buyer reports is visible to everybody and (c) sellers are allowed to condition their strategies on both public and private histories.

Proposition 3.2: If $\rho > [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$ then the set of sequential equilibrium seller payoffs of a repeated game with the stage game structure described in Table 1 and where the entire public history of buyer reports is available to short-run players is bounded above by

$$V^* = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right]$$

Proof: The proof makes use of the maximal score method, introduced by Fudenberg and Levine (1994) for computing the limiting set of payoffs of games with long-run and short-run players. The details of the proof are given in the appendix.

Another interesting, and rather unexpected, corollary of Proposition 3.1 is that, if ρ is large enough so that the most efficient equilibrium (Case 2) obtains, a seller’s maximum payoff is independent of the window width N .

Corollary 3.3: If $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$ then the maximum seller payoff attainable by an eBay-like feedback mechanism is independent of the size N of the time window.

Proof: Obvious from the fact that, for $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$ the expression for $V(0, \mathbf{s}^*)$ is independent of N ⁸.

Contrary to intuition, Proposition 3.2 and Corollary 3.3 show that publishing or summarizing larger amounts of feedback information does *not* improve the ability of eBay-like mechanisms to induce cooperation in two-outcome moral hazard settings. The simplest eBay-like feedback mechanism is one where $N = 1$: in this case, a seller’s profile $x \in \{0, 1\}$ essentially shows the seller’s *single* most recent rating ($x = 0$ means that the last rating was positive, while $x = 1$ that it was negative). The above results show that a mechanism that publishes the single most recent rating of a seller is just as efficient in dealing with two-outcome moral hazard as a mechanism that summarizes larger numbers of recent ratings or that publishes the seller’s entire feedback history. In Section 4.3 I further show that, if sellers can costlessly disappear and reappear with new identities, setting $N = 1$ results in a mechanism that is *strictly* more efficient than eBay-like mechanisms that summarize ratings over larger time windows. Although specific to two-outcome moral hazard settings, this surprising result is an encouraging indication of the power of properly designed online feedback mediators to simplify decision-making without sacrificing efficiency.

4 Robustness Analysis

Online environments are characterized by large numbers of heterogeneous players whose behavior does not always conform to the traditional game theoretic assumptions of full rationality. Furthermore, in such environments players can easily disappear and re-appear under new online identities if it is profitable for them to do so. It is therefore important that mechanisms intended for practical use in such settings be robust to a number of different contingencies arising from these properties. This section examines the robustness of eBay-like mechanisms to the presence of incorrect or incomplete feedback as well as to strategic changes of online identities.

⁸The threshold $\rho = [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$, above which the result holds, grows with N . However, for sellers with frequent transactions ($1 - \delta \approx 0$) this dependence is very weak

4.1 Quality misreporting (buyer trembles)

The analysis of Section 3 was based on the assumption that buyers truthfully report their quality observations. Nevertheless, it is plausible that some buyers might make reporting mistakes or that there might exist some irrational buyers who always submit untruthful reports. Since we have argued that no short-run player can benefit from strategically misreporting her observed quality, such mistakes can be modeled as noise. More specifically, let ε be the (exogenous) probability that a buyer misreports her observed quality. Then, the conditional probabilities of a negative rating become:

$$\begin{aligned}\tilde{\alpha} &= Pr[-|\text{high eff}] \\ &= Pr[-|\text{high qual}]Pr[\text{high qual}|\text{high eff}] + Pr[-|\text{low qual}]Pr[\text{low qual}|\text{high eff}] \\ &= \varepsilon(1 - \alpha) + (1 - \varepsilon)\alpha\end{aligned}\tag{9}$$

$$\begin{aligned}\tilde{\beta} &= Pr[-|\text{low eff}] \\ &= Pr[-|\text{high qual}]Pr[\text{high qual}|\text{low eff}] + Pr[-|\text{low qual}]Pr[\text{low qual}|\text{low eff}] \\ &= \varepsilon(1 - \beta) + (1 - \varepsilon)\beta\end{aligned}\tag{10}$$

The new conditional probabilities of negative ratings affect the expected future payoffs (equations (5) and (6)) but not the stage-game payoffs (equations (1)-(4)), since the latter depend only on the probabilities of privately observed qualities. By substituting $\tilde{\alpha}, \tilde{\beta}$ for α, β in (5) and (6) we get the following result, analogous to Proposition 3.1:

Proposition 4.1:

1. If $\rho < \tilde{\beta}/(\beta - \alpha)(\tilde{\beta} - \tilde{\alpha})$ then:
 - (a) the seller's optimal strategy is $\mathbf{s}^* = [s(x) = 0, x = 0, \dots, N]$: always exert low effort.
 - (b) the expected single-period auction revenue is given by $G(\mathbf{s}^*) = (1 - \beta)w_2$
 - (c) the seller's payoff is equal to $V(\mathbf{s}^*) = (1 - \beta)w_2/(1 - \delta)$ independently of the initial state of his feedback profile.
2. If $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)(\tilde{\beta} - \tilde{\alpha})$ then:

(a) the seller's optimal strategy is

$$\mathbf{s}^* = [s(x) = 1 - x(1 - \delta + \frac{\delta}{N}) \frac{c/w_2}{\delta(\beta - \alpha)(\tilde{\beta} - \tilde{a})}, x = 0, \dots, N]$$

(the seller always exerts high effort if he has zero negative ratings, otherwise he follows a mixed strategy in which the probability of cooperation is a linearly decreasing function of the number of negative ratings in his profile.)

(b) the expected auction revenue is a linearly decreasing function of the current number of negative ratings in the seller's profile and given by

$$G(x) = (1 - \alpha)w_2 - x(1 - \delta + \frac{\delta}{N}) \frac{c}{\delta(\tilde{\beta} - \tilde{a})}$$

(c) the seller's payoff is a function of seller's initial profile state and equal to

$$V(x_0, \mathbf{s}) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\tilde{a}c}{\tilde{\beta} - \tilde{a}} \right] - x_0 \frac{c}{\delta(\tilde{\beta} - \tilde{a})}$$

3. The above payoffs constitute an upper bound on the payoffs attainable in any sequential equilibrium of the game; the seller, therefore, cannot improve his payoffs by considering non-stationary strategies.

Proof: See appendix.

For $\alpha < 0.5 < \beta$, from (9) and (10) it is easy to see that, for $\varepsilon > 0$, $\tilde{\alpha} > \alpha$, $\tilde{\beta} < \beta$ and $\tilde{\beta} - \tilde{\alpha} < \beta - \alpha$. Proposition 4.1 then states that, consistent with intuition, quality misreporting reduces the resulting efficiency of the system and raises the threshold ρ , above which the most efficient equilibrium obtains. It is also fairly easy to see that the equivalent of Proposition 3.2 holds in this case as well. Therefore, even in the presence of occasional misreporting, eBay-like feedback mechanisms are as efficient as mechanisms that publish the entire history of past feedback.

4.2 Incomplete reporting

Reputation mechanisms rely on voluntary feedback submission. Since this incurs a (small) cost associated with connecting to a website and filling the necessary feedback forms, in the absence of

concrete incentives a fraction of buyers might submit no report to the system⁹. Although, as I mentioned in Section 3.1, such incentive schemes are not difficult to construct, in this section I examine the impact of incomplete reporting on the equilibria induced by eBay-like feedback mechanisms. This analysis applies to environments where incentive schemes are currently absent (such as eBay) or where, despite the existence of incentive schemes, there might be irrational buyers who forget or dislike to leave feedback.

The possibility of incomplete feedback introduces the need for a policy regarding the treatment of missing feedback. There are three possible options:

- Policy 1: Treat missing feedback as positive feedback
- Policy 2: Treat missing feedback as negative feedback
- Policy 3: Ignore missing feedback (do not update feedback profile)

The following proposition compares the above policies in terms of the maximum seller payoffs achievable under each of them.

Proposition 4.2: Let η_+, η_- denote the fractions of buyers who submit (truthful) ratings to the system when they observe good or bad quality respectively. These fractions are determined exogenously. Assume that the remaining buyers submit no rating. Furthermore, let V_1^*, V_2^* and V_3^* be the maximum seller payoffs attainable through an eBay-like feedback mechanism under Policies 1, 2 and 3 respectively. The following statements are true:

1. For fixed $0 \leq \eta_+, \eta_- \leq 1$, $V_1^* \geq V_2^*$. Equality holds if and only if $\eta_+ = 1$.
2. For fixed $0 \leq \eta_+, \eta_- \leq 1$, $V_1^* \geq V_3^*$. Equality holds only if $\eta_+ = \eta_-$.
3. Under Policy 1, if $\rho > \left(\frac{1}{\eta_-}\right) [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$ the seller's maximum discounted lifetime payoff is equal to

$$V(x_0, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - a} \right] - x_0 \left(\frac{1}{\eta_-} \right) \frac{c}{\delta(\beta - a)}$$

⁹eBay currently does not provide concrete incentives for feedback submission. According to recent empirical evidence, about 50% of all eBay transactions receive no ratings (Resnick and Zeckhauser, 2002).

Proof: See Appendix.

Proposition 4.2 shows that the most efficient policy regarding missing ratings is Policy 1, that is, a policy of treating missing feedback as positive feedback (“no news is good news”). Policy 2 is strictly less efficient than Policy 1 if there is incomplete reporting of positive outcomes ($\eta_+ < 1$). Policy 3 is equivalent to Policy 1 in the special case where $\eta_+ = \eta_-$ but less efficient otherwise.

Under a “no news is good news” policy, incomplete reporting impacts the minimum ρ required in order for eBay-like mechanisms to induce the most efficient equilibrium (it gets multiplied by the factor $1/\eta_- > 1$) and therefore restricts their applicability to environments with high profit margins. Perhaps surprisingly, however, if ρ is high enough, incomplete reporting does *not* lower the maximum seller payoff (attainable for $x_0 = 0$).

The results of this section underline the importance of careful design in optimizing the efficiency of online feedback mechanisms. It is interesting to note that eBay’s feedback forum currently follows Policy 3: eBay feedback profiles only show information about transactions for which feedback was provided; the remaining transactions are ignored. Proposition 4.2 shows that eBay can increase the efficiency of its mechanism by replacing that policy with a “no news is good news policy”: if no feedback is submitted for a transaction within a specified period, by default award the seller a positive rating for that transaction.

4.3 Easy name changes

In most online communities it is relatively easy for members to disappear from the community and re-register under a different online identity. Friedman and Resnick (2001) refer to this property as “cheap pseudonyms”. This property hinders the effectiveness of feedback mechanisms: online community members can build a reputation, milk it by cheating other members and then disappear and re-enter the community with a new identity and a clean record.

In our model, the ability to costlessly change one’s identity has a simple but disruptive effect: rational sellers will disappear and re-enter the community with a new identity whenever their feedback profile transitions to a state whose (lifetime discounted) payoff is lower than that corresponding to the initial state of a new seller. This has particularly severe consequences if the initial state is $x_0 = 0$ (i.e. newcomers start with a “clean record”): sellers will then be tempted to disappear as soon as they receive a single negative rating. This gives them no incentive to avoid negative ratings. Therefore, sellers will always cheat, buyers will expect them to do so, and the equilibrium outcome would correspond to Case 1 of Proposition 3.1 (no cooperation) independently of the value of ρ .

The fact that, in online settings, the mechanism designer can control the initial state of the feedback profile of newcomers suggests a solution to the problem of easy name changes: In order for the mechanism to prevent sellers from changing their identity, it suffices to set the initial state of a new seller’s feedback profile so that it corresponds to the state with the lowest possible payoff. The seller would then never find it optimal to disappear and re-enter under a new identity. In the setting of this paper, by Proposition 3.1 this means that new sellers must start the game with profiles corresponding to the “worst” possible reputation ($x_0 = N$) and must then slowly “build” their reputation by transitioning to states with *fewer* negative ratings.

Although effective in addressing the adverse consequences of easy name changes, this solution lowers the seller payoffs induced by the mechanism. In settings where players cannot change their identities the socially optimal policy is to start new sellers with a “clean” feedback profile ($x_0 = 0$), resulting in payoff

$$V(0, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right]$$

If easy name changes are a concern, new sellers start at $x_0 = N$. This lowers their payoff to

$$V(N, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right] - N \frac{c}{\delta(\beta - \alpha)}$$

The possibility of easy name changes thus forces a community to shift from an optimistic (and more efficient) policy to a more cautious (and less efficient) policy where newcomers have to “pay their dues”.

It is interesting to note that the loss associated with preventing easy name changes is minimized when $N = 1, x_0 = 1$. This result is remarkable. In Section 3.3 I showed that the maximum efficiency attainable through eBay-like mechanisms is independent of the size of the parameter N and, therefore, a mechanism that simply publishes a seller’s single most recent rating is as efficient as mechanisms that summarize larger numbers of recent ratings. In environments where players can costlessly change their identities, the result becomes even stronger: eBay-like mechanisms that publish the single most recent rating are *strictly more efficient* than mechanisms that summarize larger numbers of ratings. This is yet another unexpected result that supports the “simpler is better” principle of online feedback mechanisms.

I conclude this section by showing that, in the presence of easy name changes, no mechanism that publishes the seller’s entire feedback history (or any truncation thereof) can perform better than an eBay-like mechanism with parameters . Efficiency losses associated with easy name changes are thus inevitable and an eBay-like mechanism with $N = 1, x_0 = 1$ constitutes the best possible solution.

Proposition 4.3: If $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2$ and players can costlessly change identities, the set of sequential equilibrium seller payoffs of a repeated game with the stage game structure described in Table 1 and where the entire public history of buyer reports is available to short-run players is bounded above by

$$V(1, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right] - \frac{c}{\delta(\beta - \alpha)}.$$

Proof: See appendix.

5 Relationship to empirical findings

This section illustrates how the predictions of my theoretical model apply to online trading environments, such as eBay. I find remarkable consistency between theory and empirically observed outcomes and offer new, theory-backed explanations to hitherto poorly understood phenomena, such as the surprisingly low fraction of negative feedback on eBay.

A number of recent empirical studies have looked at various aspects of eBay’s feedback mechanism (see Dellarocas, 2002; Resnick, Zeckhauser, Swanson and Lockwood, 2002 for surveys). Although many of the results of these studies are mutually inconsistent, one particularly striking fact seems to stand out from the majority of them: the amount of negative feedback submitted to eBay’s feedback mechanism is surprisingly low.

For example, Resnick and Zeckhauser (2002) report that, based on a data set of 36,233 randomly chosen completed eBay transactions, less than 1% of ratings submitted by buyers are negative or neutral. They speculate that part of the reason for this phenomenon is that, fearing retaliation from sellers, many dissatisfied buyers are reluctant to provide negative feedback. Miller, Resnick and Zeckhauser (2002) treat this overwhelmingly positive feedback as evidence of poor functioning of eBay’s “primitive” mechanism and comment that “that such a high proportion of assessments are positive suggests that little information is conveyed”. Our theory-driven analysis suggests another




Figure 3: Maximum efficiency (relative to first-best case) induced by eBay-like feedback mechanisms in environments where $\alpha = 1 - \beta = 0.01$

explanation for the very low fraction of negative feedback on eBay: quite simply, in the two-outcome settings studied in this paper, this “primitive” mechanism happens to be so effective that, at equilibrium, it induces very high levels of cooperation. Therefore, the amount of negative feedback is expected to be very low simply because, thanks to the mechanism, the fraction of dissatisfied buyers is expected to be low as well¹⁰.

On environments like eBay the probability of a “low quality” outcome if the seller cooperates (i.e. promptly ships back goods of promised quality after receiving payment from buyer) is arguably quite low and corresponds to relatively improbable circumstances where goods are damaged or lost in the mail, delayed because of post office strikes, the buyer misunderstood the seller’s description or made a mistake when rating, etc. Therefore, I argue that on online trading environments a reasonable value for α would be 0.01 or less. The probability of a “high quality” outcome if a seller cheats is similarly very low. Assume that $\alpha = 1 - \beta = 0.01$. Then, from Proposition 3.1, for sellers with frequent trades ($\delta \approx 1$) the threshold required in order for the feedback mechanism to induce the most efficient equilibrium is a reasonable $\rho \geq 1.04$. This value corresponds to the requirement that the second highest bidder’s valuation of high quality is more than 4% higher than the cost of high effort to the seller, a condition that is arguably satisfied by the vast majority of items sold through

¹⁰The “fear of retaliation” hypothesis should not be ruled out, of course. An interesting extension of this paper would be to allow both players to rate one another and explore what strategic side effects are introduced by this new mechanism dimension.

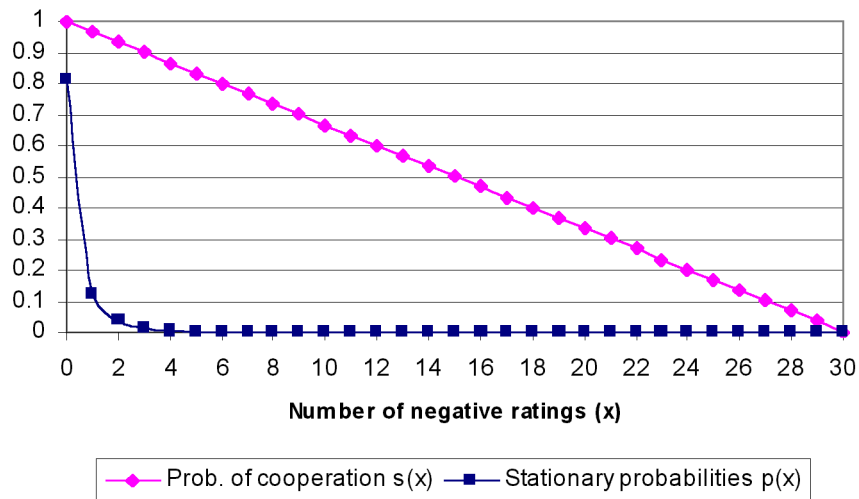


Figure 4: Equilibrium outcomes induced by eBay-like feedback mechanisms with $N = 30$ in environments where $\alpha = 1 - \beta = 0.01$, $\eta_+ = \eta_- = 0.5$, $\delta = 0.999$ and $\rho = 2.15$

such environments.

Figure 3 plots the maximum relative efficiency $V(0, \mathbf{s}^*)/V_{first-best}$ induced by eBay-like mechanisms for the above values of α and β as a function of ρ . We see that, for $\rho > 1.22$, relative efficiency is higher than 95% and for $\rho > 2.04$ it becomes higher than 99%.

Resnick and Zeckhauser report that, on their data set, roughly 50% of transactions received no rating from buyers. Although their data provides no basis for distinguishing what fractions of the unrated transactions would have received what type of rating, let us make the conservative assumption $\eta_+ = \eta_- = 0.5$. Figure 4 depicts the probabilities of cooperation $s(x)$ and stationary probabilities $p(x)$ of having x negative ratings for sellers that behave according to the optimal strategy described by Proposition 4.1 in an environment where $\alpha = 1 - \beta = 0.01$, $\eta_+ = \eta_- = 0.5$, $\rho = 2.15$ and $N = 30$ ¹¹. We see that the model predicts high levels of cooperation and low probabilities of being in states with many negative ratings (the probability of having zero negative ratings is 81%, the probability of having more than 2 negative ratings is less than 1%).

The probability of receipt of a negative rating during a period where the seller's profile has x negatives is equal to $s(x)\tilde{\alpha} + [1 - s(x)]\tilde{\beta}$. The average probability of a negative rating (or equivalently,

¹¹From Proposition 4.2, for $\eta_+ = \eta_-$, the most efficient equilibrium induced by eBay's current policy regarding missing feedback (Policy 3) is identical to the most efficient equilibrium induced by a "no news is good news policy" (Policy 1). From the proof of Proposition 4.2, under such a policy, $\tilde{\alpha} = \eta_- \alpha = 0.005$, $\tilde{\beta} = \eta_- \beta = 0.495$.




Figure 5: Expected fraction of negative feedback induced by eBay-like feedback mechanisms under incomplete reporting ($\eta_+ = \eta_- = 0.5$) as a function of the noise parameter α ($\beta = 1 - \alpha$)

the expected fraction of negative ratings in a large sample of ratings) is therefore equal to $p = \sum_{x=0}^N p(x) [s(x)\tilde{\alpha} + [1 - s(x)]\tilde{\beta}]$, where $p(x)$ are the stationary probabilities that a seller will have x negative ratings in his profile at the beginning of a period. Figure 5 plots the expected fraction of negative ratings as a function of α (under the assumption that $\beta = 1 - \alpha$ and $\eta_+ = \eta_- = 0.5$) for the minimum ρ required by Case 2 of Proposition 4.1 in order for the most efficient equilibrium to obtain (higher values of ρ result in even lower fractions of negative feedback). We see that, for $\alpha = 0.01$ only 0.98% of feedback is expected to be negative. These results are consistent with the empirical observations of eBay’s feedback mechanism and provide arguments, not only for the validity of the model presented in this paper, but also for the effectiveness of eBay’s “primitive” (according to some authors) feedback mechanism.

6 Concluding Remarks

This paper presented a theoretical analysis of the impact of feedback mechanisms in trading environments with imperfect monitoring of a seller’s effort level and two possible outcomes (corresponding to “high” and “low” quality respectively), one of which has no value to buyers. On a more abstract level, the paper provides a stylized exploration of online feedback mechanism design issues in online

purchasing environments. The analysis highlights the new design possibilities offered by the information processing capabilities of automated feedback mediators and the importance of carefully considering and fine-tuning the parameters that these new systems make available to community operators.

On a practical level, the results of this paper have important implications for eBay and other online communities that use similar feedback mechanisms.

- First, they establish that, if buyer valuations of high quality are sufficiently high (relative to the cost of high effort), eBay’s relatively “primitive” mechanism is capable of inducing high average levels of cooperation that remain stable over time. Furthermore, the buyer and seller strategies that maximize cooperation have a particularly simple stationary form. Surprisingly, efficiency cannot be improved by summarizing more ratings or by publishing detailed seller feedback histories. Far from being “primitive”, therefore, eBay’s current mechanism might, in a way, even be thought of as an “overkill”, because it summarizes large numbers of past ratings when it could induce equally efficient outcomes simply by publishing a seller’s *single* most recent rating.
- Second, they provide theoretical arguments that show that the remarkably low fraction of negative feedback on eBay arises naturally from equilibrium behavior and is not, as some authors previously speculated, an indication of the mechanism’s poor functioning.
- Third, they suggest ways in which eBay’s current mechanism can be improved. More specifically, to prevent members from changing their identities following bad ratings, eBay might want to consider starting all new members with a “bad” profile (a profile that already contains a number of negative ratings). Furthermore, in the presence of incomplete feedback submission, eBay can increase the efficiency induced by its mechanism by replacing its current policy regarding unrated transactions with the more efficient “no news is good news” policy.

As stated in the Introduction, this work is part of a larger research program whose objective is to identify “good” online feedback mechanism architectures for a variety of practically important settings. The results of this paper can be extended in a number of important directions:

- First, this paper studies a setting where there are only two seller effort levels and two stage game outcomes (quality levels). An interesting open question is how the results of this paper extend to environments where multiple effort levels and/or outcomes are possible.

- Second, the analysis of this paper assumes that all sellers are identical (in terms of their costs and conditional probabilities of outcomes given effort). In some online environments it is plausible that there might be different seller types with different cost structures and/or conditional probabilities of outcomes, unknown to buyers. For example, in online marketplaces for professional services, such as eLance.com, professionals of varying (and privately known) abilities advertise their services. In such settings, an additional objective of feedback mechanisms is to help buyers assess the type of the seller they are facing. It is therefore interesting to explore how the models of this paper extend to environments with incomplete information.
- Third, although reasonably efficient, in environments with noisy monitoring of quality eBay-like mechanisms incur efficiency losses relative to the first-best case. This paper has shown that these efficiency losses cannot be improved by following the “obvious” paths of summarizing more ratings or by publishing the entire feedback history. Nevertheless, it has not ruled out the existence of more efficient mechanisms that are based on more novel architectures. Inventing such mechanisms (or proving that they do not exist) is an endeavor of both theoretical and practical importance.

References

- Dellarocas, C. (2002) The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms. Working Paper.
- Friedman, E. and Resnick, P. (2001). The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy* 10 (1), pp. 173-199.
- Fudenberg, D. and Levine, D. K. (1994) Efficiency and Observability with Long-Run and Short-Run Players. *Journal of Economic Theory*, 62, pp. 103-135.
- Greif, A. (1993). Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition. *American Economic Review* 83, pp. 525-548.
- Houser, D. and Wooders, J. (2000) Reputation in Auctions: Theory and Evidence from eBay. Working Paper, University of Arizona.
- Kandori, M. and Matsushima, H. (1998) Private Observation, Communication and Collusion. *Econometrica*, 66 (3), May 1998, pp. 627-652.

- Klein, B. and Leffler, K. (1981) The Role of Market Forces in Assuring Contractual Performance. *Journal of Political Economy*, 89 (4), pp. 615-641.
- Klein, D. ed. (1997) *Reputation: Studies in the Voluntary Elicitation of Good Conduct*. University of Michigan Press, Ann Arbor, MI.
- Livingston, J. (2002) Can a Contract Be Self-Enforced Using Reputation? A Sample Selection Model of Internet Auctions. Working Paper, University of Maryland.
- Milgrom, P. R., North, D., and B. R. Weingast. (1990). The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs. *Economics and Politics* 2, pp. 1-23.
- Miller, N., Resnick, P., and Zeckhauser, R. (2002) Eliciting Honest Feedback in Electronic Markets, Working Paper.
- Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K. (2000) Reputation Systems. *Communications of the ACM*, 43 (12), pp. 45-48.
- Resnick, P., and Zeckhauser, R. (2002) Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In *The Economics of the Internet and E-Commerce*. Michael R. Baye, editor. Volume 11 of Advances in Applied Microeconomics. Amsterdam, Elsevier Science (forthcoming).
- Resnick, P., Zeckhauser, R., Swanson, J. and Lockwood, K. (2002) The Value of Reputation on eBay: A Controlled Experiment. Working Paper.
- Shapiro, C. (1983). Premiums for High Quality Products as Returns to Reputations. *The Quarterly Journal of Economics*, 98 (4), pp. 659-680.

Appendix

Proof of Proposition 3.1

In order for a stationary strategy \mathbf{s}^* to be an equilibrium, it must satisfy the incentive compatibility constraints (8). From (2), (7) and (8) it follows that:

$$\begin{aligned}
s^*(x) = 0 &\Rightarrow U_{coop}(x, \mathbf{s}^*) \leq U_{cheat}(x, \mathbf{s}^*) \\
0 < s^*(x) \leq 1 &\Rightarrow U_{coop}(x, \mathbf{s}^*) \geq U_{cheat}(x, \mathbf{s}^*)
\end{aligned} \tag{11}$$

and

$$V(x, \mathbf{s}^*) = G(x, \mathbf{s}^*) + U(x, \mathbf{s}^*) = \begin{cases} (1 - \beta)w_2 + U_{cheat}(x, \mathbf{s}^*) & \text{if } s^*(x) = 0 \\ G(x, \mathbf{s}^*) + U_{coop}(x, \mathbf{s}^*) & \text{if } 0 < s^*(x) \leq 1 \end{cases} \tag{12}$$

Substituting (5) and (6) into (11) and (12), the above set of constraints can be rewritten as:

$$\begin{aligned}
s^*(x) = 0 &\Rightarrow (1 - \frac{x}{N})[V(x, \mathbf{s}^*) - V(x + 1, \mathbf{s}^*)] + \frac{x}{N}[V(x - 1, \mathbf{s}^*) - V(x, \mathbf{s}^*)] \leq \frac{c}{\delta(\beta - \alpha)} \\
0 < s^*(x) \leq 1 &\Rightarrow (1 - \frac{x}{N})[V(x, \mathbf{s}^*) - V(x + 1, \mathbf{s}^*)] + \frac{x}{N}[V(x - 1, \mathbf{s}^*) - V(x, \mathbf{s}^*)] \geq \frac{c}{\delta(\beta - \alpha)}
\end{aligned} \tag{13}$$

and

$$V(x, \mathbf{s}^*) = \begin{cases} (1 - \beta)w_2 + \delta[\beta(1 - \frac{x}{N})V(x + 1, \mathbf{s}^*) + [\beta\frac{x}{N} + (1 - \beta)(1 - \frac{x}{N})]V(x, \mathbf{s}^*) \\ \quad + (1 - \beta)\frac{x}{N}V(x - 1, \mathbf{s}^*)] & \text{if } s^*(x) = 0 \\ G(x, \mathbf{s}^*) - c + \delta[\alpha(1 - \frac{x}{N})V(x + 1, \mathbf{s}^*) + [\alpha\frac{x}{N} + (1 - \alpha)(1 - \frac{x}{N})]V(x, \mathbf{s}^*) \\ \quad + (1 - \alpha)\frac{x}{N}V(x - 1, \mathbf{s}^*)] & \text{if } 0 < s^*(x) \leq 1 \end{cases} \tag{14}$$

I will show that the equilibrium strategy \mathbf{s}^* that maximizes seller payoffs depends on the ratio $\rho = w_2/c$; ρ provides a measure for the ratio of the valuation of a high quality good to the cost of high effort.

The proof is organized as follows: First I establish an upper bound on seller payoffs attainable in any sequential equilibrium of a modified game that includes all stationary equilibria of the original game. Then, I show that there exists a simple stationary equilibrium of the original game that attains payoffs equal to this upper bound. Together, the two results prove Proposition 3.1.

Consider a repeated game where the stage game structure is identical to that of Table 1 but where, instead of the statistic x , the mechanism publishes the entire history of past ratings plus the sequence of random numbers based on which it would have updated the statistic x at the end of each period

(see Section 2.2). It is easy to see that any stationary equilibrium of the original game is also an equilibrium of the modified game (because the information structure of the modified game allows all players to unambiguously compute x at each period and condition their play on that x only). The proof of this and several subsequent propositions of this paper are based on the following important lemma.

Lemma 1. The set of sequential equilibrium seller payoffs of a repeated game with the information structure described above and with the stage game structure described in Table 1 is bounded above by

$$V^* = \begin{cases} \frac{(1-\beta)w_2}{1-\delta} & \text{if } \rho < \beta/(\beta - \alpha)^2 \\ \frac{1}{1-\delta} \left[(1 - \alpha)w_2 - c - \frac{\alpha c}{\beta - \alpha} \right] & \text{if } \rho \geq \beta/(\beta - \alpha)^2 \end{cases}$$

Proof: In their 1994 paper, Fudenberg and Levine introduced an algorithm (known as the *maximal score method*) for computing the limiting set of payoffs of perfect public equilibria¹² of games with long-run and short-run players. One class of games for which the maximal score method acquires a particularly simple form are games with a *product structure*. Such games have the property that there is a separate public signal for each long-run player, which is independent of the signal for other long-run players and depends only on his own play and that of the short-run players. According to this definition, all games with a single long-run player have a product structure.

Consider a game with a single long-run player and n short-run players. Denote by A the long-run player's pure action set, ΔA the corresponding space of mixed actions, S the set of (publicly observable) stage-game outcomes and $\mathbf{B} : \Delta A \rightarrow \Delta A_1 \times \dots \times \Delta A_n$ the correspondence that maps any mixed action profile for the long run player to the corresponding static equilibria for the short-run players. Furthermore, let $h(a, \alpha_{SR})$ denote the long-run player's stage-game payoff and $p(s|a, \alpha_{SR})$ denote the probability that the stage-game outcome will be $s \in S$ if the long-run player plays $a \in A$ and the short run players play a mixed action $\alpha_{SR} \in \mathbf{B}(a)$ for some $a \in \Delta A$. Let $\Pi(\alpha_{SR})$ be the matrix with rows corresponding to actions a , columns to outcomes s and with the (a, s) component equal to $p(s|a, \alpha_{SR})$. If a game has a product structure and, in addition, has the property that $\Pi(\alpha_{SR})$ has full row rank (i.e. rank equal to the number of the long-run player's actions), then Fudenberg and Levine show that the maximum long-run player payoff v is the solution of the following linear programming problem:

¹²According to Fudenberg and Levine (1994) a strategy for long-run player i is public if at each time t , it depends only on the public information and not on the private information of that player. A perfect public equilibrium is a profile of public strategies such that at every date t and for every public history the strategies are a Nash equilibrium from that date on.

$$\begin{aligned}
& \max_{\alpha \in \Delta A, w(s) \in \mathbf{R}} v && \text{subject to} \\
& v = h(a, \alpha_{SR}) + \delta \sum_{s \in S} p(s|a, \alpha_{SR}) w(s) && \text{for } a \in A \text{ such that } \alpha(a) > 0 \\
& v \geq h(a, \alpha_{SR}) + \delta \sum_{s \in S} p(s|a, \alpha_{SR}) w(s) && \text{for } a \in A \text{ such that } \alpha(a) = 0 \\
& v \geq w(s) && \text{for } s \in S
\end{aligned}$$

If, in addition, short-run player actions are observable, Theorem 5.2 of the same paper asserts that the set of sequential equilibrium payoffs is the same as the set of perfect public equilibrium payoffs.

Under the assumption that, at the end of each period, the mechanism publishes both the rating posted by the buyer for the seller as well as a random number between 1 and N , the moral hazard game of Section 2 has $A = \{H, L\}$ (high effort, low effort), $S = \{(+, r), (-, r)\}, r \in \{1, 2, \dots, N\}$ and mixed seller actions characterized by a probability $s \in [0, 1]$ of playing H. The corresponding static best response of the short-run buyers is to bid amounts equal to $G_i = (s(\beta - \alpha) + 1 - \beta)w_i$. This results in stage-game seller payoffs $h(H, \alpha_{SR}) = (s(\beta - \alpha) + 1 - \beta)w_2 - c$ and $h(L, \alpha_{SR}) = (s(\beta - \alpha) + 1 - \beta)w_2$. Finally,

$$\begin{aligned}
\Pi(\alpha_{SR}) &= \begin{bmatrix} p[(+, 1)|H, \alpha_{SR}] & \dots & p[(+, N)|H, \alpha_{SR}] & p[(-, 1)|H, \alpha_{SR}] & \dots & p[(-, N)|H, \alpha_{SR}] \\ p[(+, 1)|L, \alpha_{SR}] & \dots & p[(+, N)|L, \alpha_{SR}] & p[(-, 1)|L, \alpha_{SR}] & \dots & p[(-, N)|L, \alpha_{SR}] \end{bmatrix} \\
&= \begin{bmatrix} (1 - \alpha)/N & \dots & (1 - \alpha)/N & \alpha/N & \dots & \alpha/N \\ (1 - \beta)/N & \dots & (1 - \beta)/N & \beta/N & \dots & \beta/N \end{bmatrix}
\end{aligned}$$

The above stage game satisfies the full row rank condition (since $\text{Rank}\Pi(\alpha_{SR}) = 2$) and has observable short-run player actions. According to the above, the maximum long-run player sequential equilibrium payoff v is the solution of the following linear programming problem:

$$\begin{aligned}
& \max_{s \in [0, 1], w(+, r), w(-, r)} v && \text{subject to} \\
& v = (s(\beta - \alpha) + 1 - \beta)w_2 - c + \delta[(1 - \alpha) \sum_{r=1}^N w(+, r) + \alpha \sum_{r=1}^N w(-, r)]/N && \text{for } a = H \text{ and } s > 0 \\
& v \geq (s(\beta - \alpha) + 1 - \beta)w_2 - c + \delta[(1 - \alpha) \sum_{r=1}^N w(+, r) + \alpha \sum_{r=1}^N w(-, r)]/N && \text{for } a = H \text{ and } s = 0 \\
& v = (s(\beta - \alpha) + 1 - \beta)w_2 + \delta[(1 - \beta) \sum_{r=1}^N w(+, r) + \beta \sum_{r=1}^N w(-, r)]/N && \text{for } a = L \text{ and } s > 0 \\
& v \geq (s(\beta - \alpha) + 1 - \beta)w_2 + \delta[(1 - \beta) \sum_{r=1}^N w(+, r) + \beta \sum_{r=1}^N w(-, r)]/N && \text{for } a = L \text{ and } s = 1 \\
& v \geq w(+, r), v \geq w(-, r) && \text{for all } r = 1, \dots, N
\end{aligned}$$

For $\rho < \beta/(\beta - \alpha)^2$ the above problem has solution:

$$\begin{aligned} v &= \frac{(1-\beta)w_2}{1-\delta} \\ s &= 0 \\ w(+, r) &= v \quad \text{for all } r = 1, \dots, N \\ w(-, r) &= v \quad \text{for all } r = 1, \dots, N \end{aligned}$$

whereas for $\rho \geq \beta/(\beta - \alpha)^2$ the solution becomes:

$$\begin{aligned} v &= \frac{1}{1-\delta} \left[(1-\alpha)w_2 - c - \frac{\alpha c}{\beta-\alpha} \right] \\ s &= 1 \\ w(+, r) &= v \quad \text{for all } r = 1, \dots, N \\ w(-, r) &= v - \frac{c}{\delta(\beta-\alpha)} \quad \text{for all } r = 1, \dots, N \end{aligned}$$

This completes the proof of Lemma 1.

Case 1: $\rho < \beta/(\beta - \alpha)^2$

From (2), (3) it is easy to see that the strategy $\mathbf{s}^* = [s(x) = 0, x = 0, \dots, N]$ induces expected auction revenue $G(x, \mathbf{s}^*) = (1 - \beta)w_2$, seller per-period payoff $h_s(x, \mathbf{s}^*) = G(x, \mathbf{s}^*) = (1 - \beta)w_2$ and expected lifetime payoff $V(x, \mathbf{s}^*) = (1 - \beta)w_2/(1 - \delta)$.

Since $U_{coop}(x, \mathbf{s}^*) = -c + \delta(1 - \beta)w_2/(1 - \delta)$, $U_{cheat}(x, \mathbf{s}^*) = \delta(1 - \beta)w_2/(1 - \delta)$ and $U_{coop}(x, \mathbf{s}^*) < U_{cheat}(x, \mathbf{s}^*)$ for all $x = 0, \dots, N$, strategy \mathbf{s}^* satisfies the incentive compatibility constraints (11) and is, therefore, an equilibrium strategy. Furthermore, according to Lemma 1, no other equilibrium strategy (stationary or non-stationary) can attain higher payoffs.

Case 2: $\rho > [\delta + N(1 - \delta)]/\delta(\beta - \alpha)^2 > \beta/(\beta - \alpha)^2$

I will derive the optimal strategy \mathbf{s}^* from among the subset of stationary equilibrium strategies where $0 < s^*(x) \leq 1$ for all $x = 0, \dots, N$ and will show that it induces payoffs equal to the upper bound established by Lemma 1. From Lemma 1, it follows that no other equilibrium strategy, stationary or otherwise, can induce higher payoffs.

Let $\mathbf{s}^* = [0 < s^*(x) \leq 1, x = 0, \dots, N]$. Since $0 < s^*(0) \leq 1$, from (14) it must be:

$$V(0, \mathbf{s}^*) = G(0, \mathbf{s}^*) - c + \delta[\alpha V(1, \mathbf{s}^*) + (1 - \alpha)V(0, \mathbf{s}^*)] \quad (15)$$

where $G(0, \mathbf{s}^*) = [s^*(0)(\beta - \alpha) + (1 - \beta)]w_2$. Furthermore, from (13) it must be:

$$V(0, \mathbf{s}^*) - V(1, \mathbf{s}^*) \geq \frac{c}{\delta(\beta - \alpha)} \quad (16)$$

It is easy to see that the maximum values of $V(0, \mathbf{s}^*)$ and $V(1, \mathbf{s}^*)$ that satisfy (15) and (16) are attained when $s^*(0) = 1$ and are independent of $s^*(x)$ for $x > 0$. More specifically:

$$\begin{aligned} s^*(0) &= 1 \\ G(0, \mathbf{s}^*) &= (1 - \alpha)w_2 \\ V(0, \mathbf{s}^*) &= \frac{1}{1 - \delta} \left([(1 - \alpha)w_2 - c] - \frac{\alpha c}{\beta - \alpha} \right) \quad x = 0, \dots, N \\ V(1, \mathbf{s}^*) &= V(0, \mathbf{s}^*) - \frac{c}{\delta(\beta - \alpha)} \end{aligned} \quad (17)$$

Having calculated $V(0, \mathbf{s}^*)$, $V(1, \mathbf{s}^*)$, it is fairly straightforward to obtain $s^*(x)$, $G(x, \mathbf{s}^*)$, $V(x, \mathbf{s}^*)$ for higher x by induction. The key hypothesis is that, for all $x = 1, \dots, N$

$$V(x, \mathbf{s}^*) = V(x - 1, \mathbf{s}^*) - \frac{c}{\delta(\beta - \alpha)} \quad (18)$$

which, by (17), is satisfied for $x = 1$. For $1 \leq x < N$ and $0 < s^*(x) \leq 1$, (18) and (13) then imply that the maximum value of $V(x + 1, \mathbf{s}^*)$ is equal to:

$$V(x + 1, \mathbf{s}^*) = V(x, \mathbf{s}^*) - \frac{c}{\delta(\beta - \alpha)} \quad (19)$$

This verifies the induction hypothesis. Combining (17) and (19) gives:

$$V(x, \mathbf{s}^*) = V(0, \mathbf{s}^*) - x \frac{c}{\delta(\beta - \alpha)} = \frac{1}{1 - \delta} \left([(1 - \alpha)w_2 - c] - \frac{\alpha c}{\beta - \alpha} \right) - x \frac{c}{\delta(\beta - \alpha)} \quad (20)$$

Substituting (20) into (14) and solving for $s^*(x)$, $G(x, \mathbf{s}^*)$ gives:

$$\begin{aligned}
s^*(x) &= 1 - x(1 - \delta + \frac{\delta}{N})\frac{c/w_2}{\delta(\beta - \alpha)^2} \\
G(x, \mathbf{s}^*) &= (1 - \alpha)w_2 - x(1 - \delta + \frac{\delta}{N})\frac{c}{\delta(\beta - \alpha)} \quad \text{for } 0 \leq x \leq N
\end{aligned} \tag{21}$$

Since we have assumed that $s^*(x) > 0$ for all $x = 0, \dots, N$, strategy \mathbf{s}^* requires that $1 - x(1 - \delta + \frac{\delta}{N})\frac{c/w_2}{\delta(\beta - \alpha)^2} > 0$ for all $x = 0, \dots, N$. This implies $\rho > [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$, which is consistent with our hypothesis.

Proof of Proposition 3.2

This proposition is a simple corollary of Lemma 1. All sequential equilibria of a game where the entire history of past ratings is public information are also equilibria of the modified game analyzed in Lemma 1 (the information structure of the game of Lemma 1 includes all past ratings plus the random numbers used by the mechanism to update the statistic x ; all equilibria of the game of Proposition 3.2 are, thus, equilibria of the game of Lemma 1 where players simply ignore the random numbers).

Proof of Proposition 4.1

The proof is identical to that of Proposition 3.1 with the only difference that the seller's expected future payoff now becomes:

$$U_{coop}(x, \mathbf{s}) = -c + \delta[\tilde{\alpha}(1 - \frac{x}{N})V(x + 1, \mathbf{s}) + [\tilde{\alpha}\frac{x}{N} + (1 - \tilde{\alpha})(1 - \frac{x}{N})]V(x, \mathbf{s}) + (1 - \tilde{\alpha})\frac{x}{N}V(x - 1, \mathbf{s})]$$

if the seller cooperates and

$$U_{cheat}(x, \mathbf{s}) = \delta[\tilde{\beta}(1 - \frac{x}{N})V(x + 1, \mathbf{s}) + [\tilde{\beta}\frac{x}{N} + (1 - \tilde{\beta})(1 - \frac{x}{N})]V(x, \mathbf{s}) + (1 - \tilde{\beta})\frac{x}{N}V(x - 1, \mathbf{s})]$$

if the seller cheats. In the above equations $V(x, \mathbf{s}) = G(x, \mathbf{s}) + U(x, \mathbf{s})$, where

$$G(x, \mathbf{s}) = [s(x)(\beta - \alpha) + (1 - \beta)]w_2$$

Proof of Proposition 4.2

Under the assumption that submitted feedback is truthful feedback, under Policy 1 the value of η_+ becomes irrelevant and the conditional probabilities of a negative rating become $\tilde{\alpha} = \eta_- \alpha$, $\tilde{\beta} = \eta_- \beta$.

The properties of the optimal seller strategy can be obtained by simple substitution of the above values of $\tilde{\alpha}$, $\tilde{\beta}$ in Proposition 4.1. More specifically, the minimum ρ required in order for the most efficient equilibrium (Case 2) to obtain is given by $\rho > \left(\frac{1}{\eta_-}\right) [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$. The corresponding seller payoffs are then equal to

$$V_1(x_0, \mathbf{s}^*) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{ac}{\beta - a} \right] - x_0 \left(\frac{1}{\eta_-} \right) \frac{c}{\delta(\beta - a)}$$

These payoffs are maximized for $x_0 = 0$. Therefore, $V_1^* = V_1(0, \mathbf{s}^*)$.

Under Policy 2, the corresponding conditional probabilities of a negative rating are $\tilde{\alpha} = 1 - \eta_+(1 - \alpha)$, $\tilde{\beta} = 1 - \eta_+(1 - \beta)$.

By substitution of the above values of $\tilde{\alpha}$, $\tilde{\beta}$ in Proposition 4.1, the minimum ρ required in order for the most efficient equilibrium to obtain is given by $\rho > \left(\frac{1}{\eta_+}\right) [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$. The corresponding seller payoffs are then equal to

$$V_2(x_0, \mathbf{s}) = \frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \left(\frac{1}{\eta_+} - 1 + \alpha \right) \frac{c}{\beta - a} \right] - x_0 \left(\frac{1}{\eta_+} \right) \frac{c}{\delta(\beta - a)}$$

These payoffs are maximized for $x_0 = 0$. Therefore, $V_2^* = V_2(0, \mathbf{s})$.

It is easy to see that $V_1^* \geq V_2^*$ with equality if and only if $\eta_+ = 1$.

The modeling of Policy 3 is more complicated. The basic observation underlying this policy is that, if no rating is posted for a seller in the current period then his feedback profile will remain unchanged in the next period. The conditional probabilities of a rating are

$$\begin{aligned} p(+|H) &= \eta_+(1 - \alpha) & p(-|H) &= \eta_- \alpha \\ p(+|L) &= \eta_+(1 - \beta) & p(-|L) &= \eta_+ \beta \end{aligned}$$

This gives:

$$\begin{aligned}
U_{coop}(x, \mathbf{s}) &= -c + \delta[p(-|H)(1 - \frac{x}{N})V(x+1, \mathbf{s}) + [1 - p(-|H)(1 - \frac{x}{N}) - p(+|H)\frac{x}{N}]V(x, \mathbf{s}) \\
&\quad + p(+|H)\frac{x}{N}V(x-1, \mathbf{s})] \\
U_{cheat}(x, \mathbf{s}) &= \delta[p(-|L)(1 - \frac{x}{N})V(x+1, \mathbf{s}) + [1 - p(-|L)(1 - \frac{x}{N}) - p(+|L)\frac{x}{N}]V(x, \mathbf{s}) \\
&\quad + p(+|L)\frac{x}{N}V(x-1, \mathbf{s})]
\end{aligned}$$

In the special case where $\eta_+ = \eta_- = \eta$, by following a procedure identical to that of the proof of Proposition 3.1 we find that for $\rho > \left(\frac{1}{\eta_-}\right) [\delta + N(1 - \delta)] / \delta(\beta - \alpha)^2$ the maximum seller payoffs are identical to those of Policy 1.

The exact expression for the payoffs in the general case where $\eta_+ \neq \eta_-$ is difficult to derive. Instead, we will show that $V_3^* < V_1^*$ using a two step procedure:

- Step 1: Show that $V_3^* \leq V_1^*$
- Step 2: Show that, under Policy 3, if $\eta_+ \neq \eta_-$ there cannot be an equilibrium with payoffs equal to V_1^* .

Step 1: Consider the moral hazard game of Section 2 augmented with a mechanism that, following each stage game, publishes a public signal taking values from a set $R = \{+, -, 0\}$ where “+”, “-“ indicate the submission of a positive or negative rating and 0 indicates the absence of any rating. In addition, the mechanism publishes the sequence of random numbers based on which it would have updated the statistic x at the end of each period (see Section 2.2). It is easy to see that any stationary equilibrium of the original game under any policy regarding missing feedback is also an equilibrium of the modified game (because, given the information structure of the modified game plus a policy regarding missing ratings, all players can unambiguously compute x at each period and therefore condition their play on that x only). Through the use of the maximal score method on this setup, it follows that, for $\rho \geq \beta / (\beta - \alpha)^2$, the payoffs attainable in any sequential equilibrium of the modified game are bounded above by

$$\frac{1}{1 - \delta} \left[(1 - \alpha)w_2 - c - \frac{ac}{\beta - a} \right] = V_1^*$$

Therefore, $V_3^* \leq V_1^*$

Step 2: Let $\Delta(x) = V(x - 1) - V(x)$. A stationary equilibrium strategy of the repeated game studied in this paper defines (a) a set of $N + 1$ parameters $s(0), \dots, S(N)$ and (b) an additional set of $N + 1$ parameters $V(0), \Delta(1), \dots, \Delta(N)$. These $2N + 2$ parameters are a solution of a system of $2N + 2$ constraints that includes (i) $N + 1$ incentive compatibility (IC) constraints (11) for $x = 0, \dots, N$ and (ii) $N + 1$ Bellman equations (12), again for $x = 0, \dots, N$.

I define the rank of a set of constraints as the number of linearly independent constraints. It is easy to see that, if $\eta_+ = \eta_-$ the $N + 1$ IC constraints have rank N (i.e. any one of them can be expressed as a linear combination of the other N) whereas if $\eta_+ \neq \eta_-$ they are linearly independent. The Bellman equations are also linearly independent. Therefore, for $\eta_+ \neq \eta_-$ the total system of constraints has rank $2N + 2$.

Assume that, under Policy 3 and for $\eta_+ \neq \eta_-$ there exists an equilibrium strategy \mathbf{s}^* with payoff $V_3(0, \mathbf{s}^*) = V_1^*$. From the Proof of Proposition 3.1 (Case 2) we know that such an equilibrium must also have $s(0) = 1$ and $\Delta(1) = c/\eta_- \delta(\beta - \alpha)$. These requirements satisfy one IC constraint and one Bellman equation (both for $x = 0$). We are, therefore, left with a system of $2N$ linearly independent constraints and $2N - 1$ unspecified parameters. This is obviously an over-specified system that has no solution. Our assumption that there exists a strategy with payoff $V_3(0, \mathbf{s}^*) = V_1^*$ has led to a contradiction.

Since, from Step 1 it must be $V_3(0, \mathbf{s}^*) \leq V_1^*$ whereas from Step 2 it follows that there is no equilibrium strategy with payoffs $V_3(0, \mathbf{s}^*) = V_1^*$, I have proven that, under Policy 3 and for $\eta_+ \neq \eta_-$, any stationary equilibrium strategy induces payoffs $V_3(0, \mathbf{s}^*) < V_1^*$.

Proof of Proposition 4.3

The proof is similar to the proof of Lemma 1 and makes use of Fudenberg and Levine's maximal score method. The basic difference is that the assumption of costless identity changes requires that the payoffs of intermediate states of the game are at least as high as the payoffs of the initial state (otherwise, whenever players are faced with the prospect of entering such states, they will simply disappear and restart the game under a different identity). The constraint $v \geq w(s)$ for $s \in S$ of the linear programming problem of Lemma 1 must thus be replaced with the constraint $v \leq w(s)$ for $s \in S$.