

Sanctioning Reputation Mechanisms in Online Trading Environments with Moral Hazard

Chrysanthos Dellarocas *

July 2004

Abstract

This paper offers a systematic exploration of reputation mechanism design in trading environments with opportunistic sellers, imperfect monitoring of a seller's actions and two possible seller effort levels, one of which has no value to buyers. The objective of reputation mechanisms in such settings is to induce sellers to exert high effort as often as possible. I study the impact of various mechanism parameters (such as the granularity of solicited feedback, the format of the public reputation profile, the policy regarding missing feedback, and the rules for admitting new sellers) on the resulting market efficiency. I find that the maximum efficiency that is attainable through reputation mechanisms is bounded away from the hypothetical first-best case where sellers could credibly pre-commit to full cooperation by a factor that is related to the probability that cooperating sellers may receive "unfair" bad ratings. Furthermore, the maximum efficiency is independent of the length of past history summarized in a seller's public reputation profile. I apply my framework to a model of eBay's feedback mechanism and conclude that eBay's simple mechanism is capable of inducing the maximum theoretical efficiency independently of the number of recent ratings that are being summarized in a seller's profile. I also derive optimal policies for dealing with missing feedback and easy online identity changes.

*Address: Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139. This paper supersedes the manuscript "Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard", MIT Sloan Working Paper No. 4297-03. This material is based upon work supported by the National Science Foundation under CAREER Grant No. 9984147. I am grateful to Yannis Bakos, Erik Brynjolfsson, Jacques Crémer, Drew Fudenberg, Roy Radner, Jean Tirole, Dimitri Vayanos, and the participants of seminars at Toulouse, London Business School, MIT, Stanford, U. Minnesota, and U. Pompeu Fabra for helpful comments.

1 Introduction

Reputation mechanisms are emerging as an increasingly important component of electronic markets, helping build trust and elicit cooperation among loosely connected and geographically dispersed economic agents (Resnick et al. 2000; Dellarocas 2003). For example, eBay’s feedback mechanism is the primary means through which eBay elicits honest behavior and, thus, facilitates transactions among strangers over the Internet (Resnick and Zeckhauser 2002). Several other online communities also rely on reputation mechanisms to promote trust and cooperation. Examples include *eLance* (online community of freelance professionals), *Slashdot* (online discussion forum where reputation scores help prioritize and filter postings), and *Epinions* (online consumer report forum where user feedback helps evaluate the quality of product reviews).

The use of reputation 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 has added some significant new dimensions to this age-old concept. Most important among them is the ability to systematically control the type and volume of reputational information exchanged within online communities through the use of properly architected information systems (*online reputation mediators*). Reputation mediators control who can participate in reputation communities, what type of information is solicited from participants, how it is aggregated, and how it is made publicly available to other community members. They enable community operators to exercise precise control over a number of parameters that are difficult or impossible to influence in traditional settings. For example, reputation mediators can replace detailed reputation histories with a wide variety of summary statistics, apply filtering algorithms to eliminate outlier or suspect postings, control the initial state of a new member’s reputation profile, etc. These design choices, in turn, impact the perceptions and, ultimately, the actions of community members towards one another.

The potential to engineer desirable social outcomes through the introduction of carefully crafted information systems is opening a new chapter on the frontiers of information systems research. Further progress in this area requires a deeper understanding of the role of reputation mechanisms in various types of communities, a careful scoping of the design space of online reputation mediators, and theory-driven guidelines for selecting the most appropriate mechanism architecture for a given class of settings.

This paper offers a systematic exploration of reputation mechanism design issues in trading environments where the primary trust issues are related to *moral hazard* (for example, bilateral exchange settings where the buyer pays first and the seller is then tempted to cheat). The objective of reputation mechanisms in such settings is to promote cooperative and honest behavior among self-interested economic agents. In contrast, in settings where the primary trust issues are related to *adverse selection*, the role of reputation mechanisms is to help the community learn the (initially unknown) attributes of community members (such as their ability, honesty, etc.).

A fundamental distinction between the two kinds of settings is that, in the former, all community members are assumed to be equally capable of behaving in “good” and “bad” ways and will act in ways that maximize their individual payoffs, whereas in the latter, community members differ in “type” (e.g. some are intrinsically more capable of providing good service than others) and behave in ways that are constrained by their type. In settings of the first kind, reputation mechanisms use the threat of future information disclosure to promote *discipline*. In contrast, in settings of the second kind reputation mechanisms publish information accumulated from past transactions to promote *social learning*. In common with other researchers (Kuwabara 2003) I will use the terms *sanctioning* and *signaling* to refer to these two roles of reputation mechanisms¹.

Epinions and Amazon reviews are examples of reputation mechanisms that primarily serve a signaling role: they spread information about the (initially privately known, but relatively “fixed”) qualities of the products being reviewed. eBay, on the other hand, is an example of a reputation mechanism whose primary objective is sanctioning. eBay users do not rate sellers on the quality of their products (in fact, it is not uncommon for the same eBay seller to offer products of various qualities) but rather on how well sellers were able to deliver what was promised on the item description. The role of eBay’s reputation mechanism is to promote honest trade rather than high quality products.

The objective of this paper is to explore what constitutes “good” sanctioning reputation mechanism design in the context of an abstract trading environment that captures the essential properties of online auction marketplaces. I consider a fairly general class of reputation mechanisms and study the

¹In many practical settings both moral hazard and adverse selection considerations are simultaneously present: Players *both* behave strategically *and* differ in their intrinsic ability levels, cost structures, etc. (This makes their relative preferences for good and bad behavior a function of their type.) In such settings, reputation mechanisms play both a sanctioning and a signaling role. In this paper we restrict our attention to the extreme case where only moral hazard considerations are present. The study of hybrid (sanctioning and signaling) reputation mechanisms is a natural next step of this line of research.

impact of various mechanism parameters (such as the granularity of solicited feedback, the amount of information published by the mechanism, the policy regarding missing feedback, and the rules for admitting new sellers) on the induced levels of cooperation and payoffs of buyers and sellers.

Section 2 of the paper introduces the model. Section 3 develops the basic intuitions by considering the simple class of binary reputation mechanisms (mechanisms where buyers can only report the outcome of a transaction as “positive” or “negative”). The analysis derives two important results: First, in environments with imperfect monitoring of the seller’s actions, the maximum efficiency that is attainable through the use of reputation mechanisms is bounded away from the hypothetical first-best case where sellers could credibly pre-commit to full cooperation by a factor that is related to the probability that cooperating sellers may receive “unfair” negative ratings. Second, in direct contrast to signaling reputation mechanisms, the maximum efficiency that can be induced by a sanctioning reputation mechanism is independent of the length of past history summarized in a seller’s public reputation profile. The length of history captured in a seller’s profile affects the duration of punishment associated with a bad rating but not the total amount of punishment, or the mechanism’s ability to distinguish a seller’s current action. I show that swifter punishment that lasts for fewer periods is preferable to more lenient punishment that lasts longer, because it reduces the minimum stage-game profit margin required for the mechanism to be effective. Therefore, a mechanism that, each period, only publishes a player’s single most recent rating minimizes the profit margin requirement (and thus maximizes the breadth of the mechanism’s applicability) without sacrificing efficiency.

I apply the above results to a theoretical model of eBay’s feedback mechanism. The principal conclusion is that eBay’s simple mechanism is capable of inducing the maximum efficiency attainable by any binary reputation mechanism independently of the amount of recent activity that is being summarized in the seller’s eBay feedback profile. Moreover, it allows traders to maximize their payoffs by relying on a simple stationary strategy according to which a seller’s expected probability of cooperation is a decreasing linear function of the sum of negative ratings in his current eBay “ID card”.

Section 4 extends the analysis to mechanisms that allow multiple buyer report types (for example, integer ratings between 1 and 5). I show that maximally efficient sanctioning reputation mechanisms can be constructed by designating a subset of those report types as “bad” reports and punishing the

seller by a fixed amount whenever he receives *any* report whose type belongs to the “bad” set. Given a set of “bad” report types, the mechanism’s maximum efficiency is inversely proportional to the likelihood ratio of punishment if the seller cooperates vs. if he cheats. The “bad” report set should, thus, be the subset of report types that minimizes the above likelihood ratio.

Finally, I consider the implications of easy online identity changes and incomplete feedback submission to the effectiveness of reputation mechanisms. I find that, in settings where feedback submission is voluntary, the optimal policy for dealing with missing feedback is to treat it as positive feedback. Moreover, in online settings where sellers can freely change identities, charging an entrance fee equal to the difference between the payoff of a seller with a clean record and the payoff of a seller with the worst possible reputation profile discourages sellers from the temptation to change their identity following the receipt of a bad rating and, thus, maintains their incentives to cooperate.

2 The Setting

The setting involves a marketplace where, in each period, a monopolist long-run seller provides one unit of a product or a service (“good”) to one of multiple short-run buyers. Following receipt of payment, the seller can exert either high effort (“cooperate”) or low effort (“cheat”). The seller’s action affects the probability distribution of the good’s quality and, thus, its expected value to a buyer. 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 value.

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 expected valuation w_i for a good produced by a cooperating seller. The expected valuation of a good produced by a cheating seller is zero for all buyers. The seller will, thus, not attract any buyers unless he can credibly promise to cooperate with positive probability. 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 costs c if the seller cooperates. The cost of cheating is normalized to zero. 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 highest bidder. The winning bidder pays a price equal to the second-highest bid.

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.

The only Nash equilibrium of the above setting is one where the seller will always cheat. Knowing this, no buyer will post a positive bid and thus no trade will ever take place. The objective of reputation mechanisms is to induce (at least partial) seller cooperation and thus to facilitate profitable trade despite the seller's short-term temptation to cheat.

A *reputation mechanism* allows each period's buyer to report the quality of the good she received to the center. I assume that the mechanism allows the current buyer to submit one out of a finite set of possible reports $r \in \{R_i | i = 1, \dots, m\}$, where reports indexed by a higher i indicate higher levels of buyer satisfaction². Given that, in most current systems, feedback submission is voluntary I also allow for the possibility that the current buyer does not submit any feedback by introducing a special report type R_0 that stands for "no report". Buyer reports are connected to seller actions through a conditional probability distribution $p(\text{report}|\text{action})$ that specifies the probability of a given type of report given the seller's (hidden) action. I assume that the conditional distribution $p(\cdot|\cdot)$ is

²For example, Amazon Auctions allows buyers to rate a transaction by posting an integer rating from 1 to 5 (higher ratings indicate higher satisfaction), whereas eBay allows three types of ratings ("negative", "neutral", or "positive"), again indicating increasing levels of buyer satisfaction.

1. The seller offers a single unit of a good, promising to exert high effort (as there is no demand for low effort).
2. The center publishes the seller's current reputation profile x .
3. Buyers bid their expected valuations for the good in a second price Vickrey auction; the winning bidder pays the second-highest bid.
4. The seller decides whether to exert high effort at cost c , or low effort at cost 0.
5. The buyer receives the good, experiences its quality, and reports on the quality of the good received to the center. The center updates the reputation profile of the seller accordingly.

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

exogenously given and depends on the set of reports supported by the mechanism and the properties of the trader population. The center aggregates all past reports and publishes a *reputation profile* for the seller at the beginning of each period. The reputation profile can be the unabridged history of past reports, any truncation of that history, or any statistic derived from that history. Table 1 summarizes the resulting stage-game.

The above specification implicitly assumes that, although buyers can make occasional mistakes or fail to submit feedback altogether, they are not acting strategically: the conditional distribution of ratings given a seller's action is, thus, independent of the seller's current profile or the history of play. 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 large-scale electronic markets). Buyers are then indifferent between truthful reporting and untruthful reporting³.

The goal of the mechanism designer is to select the parameters of the reputation mechanism to maximize the efficiency (total surplus) of the market. Specifically, the designer can control: (i) the set of allowed buyer reports, (ii) the format of the published reputation profile, (iii) the initial state of the reputation profile and other rules of entry for a new seller, and (iv) the policy for treating missing feedback. In addition, the designer must take into account the relative ease with which players can disappear and re-appear under different identities in online settings and must consider the robustness of their mechanism design to such contingencies.

The objective of the following sections is to derive a theoretical framework that connects the above

³Moreover, it is possible to devise a side payment mechanism that provides buyers with strict incentives to participate in the reputation mechanism as well as to 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 paper.

design dimensions to the maximum efficiency induced by the mechanism in the presence of incomplete feedback submission and identity changes. Section 3 derives the basic framework by considering the simple class of binary reputation mechanisms, that is mechanisms where buyers can only report the outcome of a transaction as “positive” or “negative”. Section 4 extends the results in settings that allow multiple report types.

3 Binary Reputation Mechanisms

The basic intuitions of this paper can be derived by considering the simple class of *binary* reputation mechanisms, that is, mechanisms where buyers can only report the outcome of a transaction as either “positive” or “negative”, with positive reports indicating that a good of satisfactory quality was received, and negative reports indicating a good of unsatisfactory quality. A binary reputation mechanism is characterized by its associated *information structure* (conditional probability distribution) $p(\text{report}|\text{action})$. I assume that both types of reports can occur with positive probability following both seller actions⁴. Specifically, I assume the following information structure:

$$\begin{aligned} p(+|coop) &= 1 - \alpha & p(-|coop) &= \alpha & 0 < \alpha < \beta < 1 \\ p(+|cheat) &= 1 - \beta & p(-|cheat) &= \beta \end{aligned} \tag{1}$$

where $+$, $-$ indicate a positive and negative report respectively.

3.1 Upper bounds on mechanism efficiency

I begin the analysis by assuming that buyers always submit feedback (i.e. $p(R_0|\cdot) = 0$) and that sellers cannot change identities. Given the above information structure, the reputation mechanism designer’s problem then focuses on deciding how to aggregate and publish buyer reports to maximize social welfare in the marketplace. The theory of repeated games provides powerful frameworks for answering such questions.

⁴This assumption intends to capture the idea that, even when a seller cooperates, it is still possible that he may get an occasional negative rating due to factors beyond his control (e.g. a shipped item was lost or damaged in the mail, the buyer misunderstood what the seller actually promised to deliver, etc.). Likewise, even when a seller cheats, it is still possible that he may get an occasional positive rating (e.g. the buyer mistakenly posted a rating destined for another transaction).

Let $s(h, x) \in [0, 1]$ denote the seller's strategy, equal to the probability that the seller will cooperate (i.e., exert high effort) following receipt of payment if the private history of play⁵ is h and his publicly available reputation profile at the beginning of the current period is equal to x . To simplify the analysis I initially restrict seller strategies to public strategies, i.e. strategies $s(x) \in [0, 1]$ that only depend on public information. I subsequently show that the seller can do no better by considering more complex strategies that are also conditioned on his private information.

Given their beliefs about the seller's strategy $s(x)$, 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 is to bid an amount equal to her expected valuation $G_i(x) = s(x)w_i$ resulting in expected auction revenue for that period $G(x) = s(x)\hat{w}_2$ where \hat{w}_2 is the expected value of the second highest bidder's valuation of a cooperating seller's output. The seller's corresponding expected current period payoff is $h_s(x) = G(x) - s(x)c = s(x)(\hat{w}_2 - c)$. The ex-ante expected surplus for the winning bidder is $h_b(x) = s(x)(\hat{w}_1 - \hat{w}_2)$ where \hat{w}_1 is the expected value of the highest bidder's valuation of a cooperating seller's output. Finally, the expected total surplus generated in the current period is equal to $h(s) = h_b(x) + h_s(x) = s(x)(\hat{w}_1 - c)$.

If $\hat{w}_2 > c$, both seller profits and buyer surplus are proportional to $s(x)$: higher levels of seller cooperation benefit *all* players. Therefore, if buyers have correct beliefs about the seller's strategy (a condition that, by definition, holds at equilibrium), a strategy that maximizes the seller's payoff also maximizes buyer surplus and thus social surplus (market efficiency). For that reason in the rest of the paper we can focus without loss of generality on finding reputation mechanism designs and associated seller strategies that maximize the seller's lifetime discounted payoff. For the same reason, I will be using the terms "seller payoffs" and "efficiency" interchangeably.

Each period the seller decides his level of effort after he receives payment. His decision problem is to select a strategy $s(x)$ that maximizes the present value of his remaining discounted payoff. The seller's remaining payoff at the time of effort selection is equal to:

$$U_{coop}(x) = -c + \delta[(1 - \alpha)V(x^+(x)) + \alpha V(x^-(x))] \quad (2)$$

⁵The private history of play includes the bid amounts, seller actions, and buyer ratings for all past periods and the bid amounts for the current period.

if the seller cooperates,

$$U_{cheat}(x) = \delta[(1 - \beta)V(x^+(x)) + \beta V(x^-(x))] \quad (3)$$

if the seller cheats, and

$$U(x) = s(x)U_{coop}(x) + [1 - s(x)]U_{cheat}(x) \quad (4)$$

if he follows a mixed strategy. In the above equations $x^+(x)$, $x^-(x)$ denote the seller's new reputation profile following the receipt of a positive and negative rating respectively, while $V(x) = G(x) + U(x)$ denotes the seller's expected future payoff at the *beginning* of a period where the seller's profile is equal to x .

A strategy $s(x)$ is an equilibrium strategy if and only if it satisfies the following incentive compatibility constraints for all x , given correct buyer beliefs:

$$\begin{aligned} s(x) = 0 &\Rightarrow U_{coop}(x) \leq U_{cheat}(x) \\ 0 < s(x) < 1 &\Rightarrow U_{coop}(x) = U_{cheat}(x) \\ s(x) = 1 &\Rightarrow U_{coop}(x) \geq U_{cheat}(x) \end{aligned} \quad (5)$$

The form of the resulting equilibria depends on the format of the public reputation profile x and the transition functions $x^+(x)$, $x^-(x)$. The most general case is one where the reputation mechanism makes the entire past history of buyer reports publicly available to future buyers. A seller's reputation profile x is then an ordered set of positive and negative reports, $x^+(x) = (x, +)$, and $x^-(x) = (x, -)$. All equilibria of reputation mechanisms that publish truncations or statistics of the history of buyer reports are also equilibria of this more general setting (because, given the entire history, players can construct these truncations or statistics and condition their play on them). Our first result establishes upper bounds on the seller payoffs that are achievable in such a setting.

Proposition 1: Let $\rho = \hat{w}_2/c$; ρ is a measure of the expected stage-game profit margin of a fully cooperating seller.

1. The set of perfect public equilibrium payoffs⁶ of a repeated game where the stage game

⁶A strategy for long-run player is public if at each period, it depends only on the public information and not on

structure is described in Table 1, the reputation mechanism has information structure (1), and the entire history of past reports is available to buyers, is bounded above by:

$$\begin{aligned}
 V^* &= 0 && \text{if } \rho < (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha) \\
 V^* &= \frac{1}{1-\delta} \left(\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right) && \text{if } \rho \geq (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)
 \end{aligned}$$

2. The above bounds remain unchanged if the seller is allowed to condition his actions on both public information and private information not available to buyers (e.g. his past actions, past bid amounts, etc.) and the full set of sequential equilibria of the repeated game is considered.

Proof: The proof makes use of the maximal score method, introduced by Fudenberg and Levine (1994) for computing the limiting set of payoffs of repeated games with long-run and short-run players. Note that, in such games, the Folk Theorem does not apply. Thus, the maximum attainable payoff is strictly bounded away from the first-best payoff $V_{first-best} = (\hat{w}_2 - c)/(1 - \delta)$ that would have been attainable in settings where sellers could credibly pre-commit to full cooperation. The details of the proof are given in the appendix.

Proposition 1 provides an upper bound on the maximum efficiency that is achievable through a binary reputation mechanism with a given information structure but does not provide specific guidelines on how such an equilibrium can be constructed. The following sections discuss two concrete implementations that achieve the maximum efficiency and clarify the intuitions behind this result.

3.2 Efficient implementation using a randomization device

A simple way to implement an equilibrium that attains payoffs equal to the upper bound of Proposition 1 is through the use of a two-state randomization device⁷. The idea is that a seller’s public reputation profile can be in one of two states: a “good” state and a “bad” state. New sellers always start from the good state. Whenever a seller receives a positive rating the center keeps his profile in the good state. Whenever he receives a negative rating, however, the center changes the seller’s profile to “bad” with probability π . The following theorem shows that, if ρ is high enough, π can

the private information of that player. A perfect public equilibrium is a profile of public strategies such that at every period and for every public history the strategies are a Nash equilibrium from that date on.

⁷Many thanks to Drew Fudenberg for suggesting this implementation method.

be chosen to construct an equilibrium where (i) the seller always cooperates as long as his profile is good, (ii) the seller always cheats (effectively placing himself out of the market) when his profile turns bad, (iii) the seller’s lifetime discounted payoff is equal to the maximum achievable payoff V^* predicted by Proposition 1.

Proposition 2:

1. If $\rho \geq (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ then a randomization device with $\pi = (1 - \delta)/\delta[\rho(\beta - \alpha) - \beta]$ induces the seller to cooperate as long as he remains in the “good” state. The present value of the seller’s discounted lifetime payoff in the “good” state is equal to $V_{good} = V^* = \frac{1}{1-\delta} \left(\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right)$.
2. If $\rho < (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ then no randomization device can provide sufficient incentives for cooperation. The only equilibrium is one where the seller always cheats, no buyers submit bids and $V_{good} = 0$

The above construction helps provide an intuitive understanding of the situation. Sanctioning reputation mechanisms induce current cooperation through the threat of punishment following bad news about the seller’s past behavior. In our setting, punishment is equal to the opportunity cost of foregone future profits when a seller’s profile turns bad (because the seller then effectively gets expelled from the market). The probability of a profile state transition is proportional to π times the probability of a negative rating. The latter is higher by $\beta - \alpha$ when the seller cheats. If π is set high enough so that the present value of the incremental opportunity cost associated with cheating exceeds the corresponding present gains c , rational sellers will find it preferable to cooperate. There is, thus, a threshold value π_0 such that, if we set $\pi \geq \pi_0$, the seller will be induced to cooperate. Unfortunately, in the presence of imperfect monitoring, even cooperating sellers will eventually accumulate negative ratings and will transition to the bad state with positive probability $\pi\alpha$. This property results in efficiency losses relative to the first-best case⁸. Lower values of π result in lower probabilities of “unfair” state transitions and, thus, in higher efficiency. To maximize efficiency the mechanism designer’s objective should be to set π equal to the minimum value that induces cooperation. Setting π just

⁸Because of imperfect monitoring of a seller’s actions, the center cannot be sure whether a negative rating is due to cheating behavior or bad luck. However, in order to sustain cooperation it must commit to punish the seller no matter what the reason. The effect is reminiscent of Green and Porter’s 1984 analysis of noncooperative collusion under imperfect price information.

high enough to make the seller indifferent between cooperation and cheating attains the maximum achievable payoff. This is equal to the first-best payoff $V_{first-best} = (\hat{w}_2 - c)/(1 - \delta)$ attainable in (hypothetical) settings where sellers can credibly pre-commit to full cooperation, minus a cost $\alpha c/(1 - \delta)(\beta - \alpha)$ associated with the mechanism's less than perfect ability to distinguish the seller's actions.

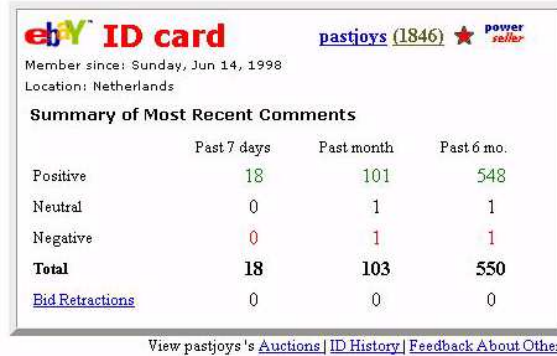
Since the mechanism does not punish sellers by levying actual penalties but simply by triggering their future exclusion from the market, the condition $\rho \geq (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ expresses the fact that, for the mechanism to provide incentives for cooperation, the stage-game profit margin of a cooperating seller must be high enough so that the present value of future profits from cooperation (that the seller jeopardizes by cheating) is higher than the current gains from cheating. This condition limits the applicability of sanctioning reputation mechanisms to settings where trusted sellers enjoy sufficiently high premiums (Klein and Leffler 1981; Shapiro 1983).

3.3 Efficient implementation using summary statistics

Though it attains the maximum theoretical efficiency, the use of a two-state randomization device has the undesirable property that cooperating sellers will eventually be expelled from the market with positive probability. Sellers may therefore consider such a mechanism unfair and avoid participating in communities that use it. Electronic markets might prefer to induce cooperation using measures of a less draconian nature. This section analyzes a reputation mechanism that publishes summary statistics of a seller's recent ratings. This mechanism has practical interest because it resembles important aspects of the feedback mechanism used by eBay. I show that this eBay-like mechanism attains the theoretically maximum efficiency of Proposition 1 and has the desirable property that it allows sellers to remain on the market indefinitely. A striking corollary of my analysis is that maximum efficiency is independent of the amount of seller activity summarized by the mechanism in the seller's public reputation profile. In fact, it turns out that summarizing fewer ratings broadens the range of settings where the mechanism can induce cooperation (because it places less stringent minimum stage-game profit margin requirements). This result is in direct contrast to equivalent results associated with mechanisms whose primary objective is social learning⁹ and shows that the

⁹The classic results of Blackwell (1951; 1953) show that, in such settings, more information always enables better decisions.

Figure 1: eBay’s ID Card



optimal architectures of the two classes of mechanisms can be very different.

Our mechanism hides the seller’s detailed rating history from buyers; instead at the beginning of each period it publishes the sum of positive and negative ratings posted for the seller in the most recent N transactions. N is a mechanism parameter chosen by the designer. A seller’s public reputation profile can then be modeled as a pair $p = (z, N)$, where N is the width of the time window and z is the sum of negative ratings accumulated in the most recent N transactions. If we ignore neutral ratings and interpret N as an estimate of the number of transactions performed by a seller during a six-month period, the above mechanism can be thought of as an abstract model of eBay’s “ID Card”, which summarizes ratings posted on a seller during the most recent six-month period (Figure 1).

If N remains constant for the duration of the game, a seller’s public reputation profile (the only information made available to buyers) is completely characterized by the number of negative ratings z in the current “time window”. On the other hand, a seller’s *private* reputation profile (known to the center and the seller) is equal to an ordered list of length N , consisting of the N most recent positive and negative ratings in chronological order (earlier ratings first). If we assign the digit 1 to a positive rating and 0 to a negative rating, each of the 2^N possible private profiles can be mapped to an N -digit binary number x between 0 and $2^N - 1$ (for example, if $N = 3$ then the profile $(+, -, +)$ maps to $x = 101 = 5$). In the rest of the section we will refer to x as the *algebraic representation* of the seller’s private reputation profile.

At the end of each period the center deletes the first (earliest) rating from the seller’s private profile and appends the most recent rating at the end of the profile. Algebraically, this operation is equivalent

to: (i) taking the $N - 1$ rightmost bits of the binary representation of the original x (they are equal to the remainder of the division of x by 2^{N-1}), (ii) multiplying them by 2 (this appends a zero at the end), and (iii) adding 1 if the most recent rating was positive (this changes the tail zero to one). The private profile transition functions can thus be expressed algebraically as:

$$x^-(x) = 2(x \oplus 2^{N-1}) \quad x^+(x) = 2(x \oplus 2^{N-1}) + 1$$

where \oplus denotes the *modulo* (division remainder) operator. I will show that, if the profit margin of a fully cooperating seller is high enough, the above mechanism induces a cooperative equilibrium that attains the maximum efficiency predicted by Proposition 1.

Let us assume that the seller follows a public strategy $s = \{s(z)|z = 0, \dots, N\}$ that only depends on the number of negative ratings on the seller's profile. The seller's maximization problem then becomes:

$$V(x) = \max_{s(z(x))} \begin{cases} s(z(x))\hat{w}_2 - c + \delta[(1 - \alpha)V(2(x \oplus 2^{N-1}) + 1) + \alpha V(2(x \oplus 2^{N-1}))] & \text{if } s(z(x)) > 0 \\ s(z(x))\hat{w}_2 + \delta[(1 - \beta)V(2(x \oplus 2^{N-1}) + 1) + \beta V(2(x \oplus 2^{N-1}))] & \text{if } s(z(x)) < 1 \end{cases} \quad (6)$$

subject to the incentive compatibility constraints:

$$\begin{aligned} s(z(x)) = 0 &\Rightarrow V(2(x \oplus 2^{N-1}) + 1) - V(2(x \oplus 2^{N-1})) \leq c/\delta(\beta - \alpha) \\ 0 < s(z(x)) < 1 &\Rightarrow V(2(x \oplus 2^{N-1}) + 1) - V(2(x \oplus 2^{N-1})) = c/\delta(\beta - \alpha) \\ s(z(x)) = 1 &\Rightarrow V(2(x \oplus 2^{N-1}) + 1) - V(2(x \oplus 2^{N-1})) \geq c/\delta(\beta - \alpha) \end{aligned} \quad (7)$$

for all $0 \leq x \leq 2^N - 1$. The following Proposition holds:

Proposition 3: Consider a repeated game with the stage game structure described in Table 1 and an eBay-like reputation mechanism with information structure (1). Each period the mechanism publishes the sum of negative ratings received by the seller during the most recent N transactions.

1. If $\rho = \frac{\hat{w}_2}{c} \geq \frac{N}{(\beta - \alpha) \sum_{i=1}^N \delta^i}$ then there exists a perfect public equilibrium where:

- (a) the seller’s probability of cooperation is a decreasing linear function of the current number z of negative ratings in his profile:

$$s(z) = 1 - z\theta \quad \text{where } \theta = \frac{c}{\hat{w}_2(\beta - \alpha) \sum_{i=1}^N \delta^i}$$

- (b) the present value of the seller’s discounted lifetime payoff is maximized during periods when the seller has zero negative ratings in his profile; it is then equal to the maximum payoff $V^* = \frac{1}{1-\delta} \left(\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right)$ attainable by any reputation mechanism with information structure (1).

- (c) the maximum attainable payoff is independent of the number N of transactions summarized in the seller’s public profile

2. If $\rho < \frac{N}{(\beta - \alpha) \sum_{i=1}^N \delta^i}$ then all perfect public equilibria achieve seller payoffs strictly below V^* .

Proposition 3 states that, if ρ is large enough, eBay-like reputation mechanisms induce an equilibrium where buyers (correctly) expect the seller to cooperate fully as long as he has no negative ratings in his profile and otherwise to follow a mixed strategy where the probability of cooperation is a decreasing linear function of the number of negative ratings in his profile. If sellers start the game with a “clean record” (zero negative ratings), this equilibrium achieves the maximum efficiency attainable by any reputation mechanism with information structure (1).

Observe that, even though the form of the above equilibrium has an intuitive flavor that is consistent with empirical results (buyers post lower bids for sellers with more negative ratings in their profile)¹⁰, the equilibrium does not arise because of buyer inferences regarding the seller’s type (i.e. it does not follow that sellers with more negatives are believed to be “bad” sellers with higher probability). It should rather be viewed as a sanctioning scheme, a *community norm* that maps reputation profiles to buyer beliefs about the seller’s *current* action (and thus prescribes how a seller with a given reputation profile should be treated) with the property that those beliefs are self-fulfilling (given the beliefs, the seller’s optimal strategy is to behave exactly according to them).

In common with the two-state randomization device of Section 3.2, the optimal “community norm” induced by eBay-like reputation mechanisms relies on the threat of punishment following a negative

¹⁰See Dellarocas (2003), Section 3 for a survey of empirical studies of eBay’s feedback mechanism.

rating to induce seller cooperation. Upon receipt of a negative rating, the randomization device either delivers “lethal” punishment (effective expulsion of the seller from the marketplace) or forgives the seller; eBay-like mechanisms, on the other hand, deliver a less severe per-period punishment that lasts for a finite number of periods: every time the seller receives a negative rating, the anticipated probability of seller cooperation is reduced by an amount θ for the N following cycles (the length of time it takes for the negative rating to propagate to the other end of the time window and be erased from the seller’s public profile). Given their anticipation, buyers lower their bids by θw_i during the following N -cycle period, reducing the present value of the seller’s expected revenues by $\sum_{i=1}^N \delta^i \theta \hat{w}_2$.

To induce cooperative behavior, the opportunity cost of cheating must be greater than or equal to the gains from cheating. The opportunity cost of cheating is equal to the incremental probability of a negative rating induced by cheating behavior $(\beta - \alpha)$, times the present value of lost profits $\sum_{i=1}^N \delta^i \theta \hat{w}_2$ following a negative rating. The gains from cheating on the other hand are equal to c . To minimize efficiency losses due to negative ratings “unfairly” posted for cooperating sellers, the opportunity cost of a negative rating must be as low as possible. Maximum efficiency is attained when the opportunity cost of cheating is exactly equal to the present value of the gains from cheating. This happens when $\theta = \frac{c}{\hat{w}_2(\beta - \alpha) \sum_{i=1}^N \delta^i}$.

Proposition 3 shows that eBay-like feedback mechanisms allow traders to maximize their payoffs by relying on simple stationary strategies of intuitive appeal; players need not even consider more complex strategies (or the possibility that other players may use such strategies). A further strength of the results is that they are robust to errors individual buyers may make in the calculation of $s(z)$ as long as these errors are not systematic (i.e. as long as they are completely random with an expected value of zero)¹¹. These properties make eBay’s mechanism especially well suited to large-scale online trading environments, as it allows even inexperienced traders to perform well by relying on simple, intuitive strategies.

A striking corollary of Proposition 3 is that the maximum efficiency attainable by an eBay-like mechanism is independent of the number of ratings N that are included in the published summary statistic. This means that a mechanism that, each period, publishes only the outcome of a seller’s single most recent transaction ($N=1$) is capable of inducing the same average levels of cooperation

¹¹Many thanks to an anonymous referee for pointing this out.

and total surplus as more sophisticated mechanisms that summarize arbitrarily large numbers of recent transactions.

At first this result might seem counterintuitive. Common wisdom in decision theory suggests that more information leads to better outcomes (Blackwell 1951; 1953). The key to understanding the result lies in the distinction between the signaling and sanctioning roles of reputation mechanisms. Reputation mechanisms whose primary objective is signaling (such as Epinions and Amazon reviews) improve market efficiency by enabling better consumer inferences based on *past* information about, initially unknown, but “fixed” product attributes. The efficiency of such mechanisms is related to the quality of the accumulated information *à la* Blackwell. Therefore, the larger the number of published reports, the better the outcomes. Sanctioning mechanisms, on the other hand, improve market efficiency by inducing higher cooperation levels through the threat of *future* punishment associated with spreading information about the outcome of the current transaction. The efficiency of a sanctioning mechanism is, therefore, related to the incremental present value of punishment associated with a “bad” action in the current transaction. In our setting, the incremental present value of punishment that maximizes efficiency is independent of N and equal to the stage-game gains from cheating. The parameter N simply determines the number of “installments” during which the total punishment will be “paid”¹². For that reason, mechanism efficiency is independent of the number of ratings N summarized in the seller’s public reputation profile.

A more subtle corollary of Proposition 3 is that, as the number of ratings N summarized in a seller’s public profile grows, the minimum stage-game cooperative profit margin $\rho = \hat{w}_2/c \geq N/(\beta - \alpha) \sum_{i=1}^N \delta^i$ that is required for the mechanism to attain maximum efficiency outcomes also grows¹³. To see why, observe that for the mechanism to attain maximum efficiency, it must be able to make sellers indifferent between cooperation and cheating even during periods when they are paying “punishment installments” for past negative ratings. During those periods, sellers are expected to cooperate with probability $1 - z\theta$ where z is the number of past negative ratings in their profile. Sellers will only be deterred from cheating during such periods if it is possible to threaten them with a further drop in their expected probability of cooperation if they get another negative rating. But this can only

¹²Observe that, as the width of the time window N increases, the nominal value of each punishment “installment” (proportional to the incremental drop in the stage-game probability of cooperation $\theta = c/\hat{w}_2(\beta - \alpha) \sum_{i=1}^N \delta^i$ following a negative rating) decreases, whereas the present value of all punishment installments (proportional to $\sum_{i=1}^N \delta^i \theta$) remains constant.

¹³For $\delta < 1$, $N/\sum_{i=1}^N \delta^i$ is an increasing function of N .

happen if $1 - (z + 1)\theta \geq 0$, that is, if it is possible to further reduce their expected probability of cooperation by θ and still be left with a non-negative probability of cooperation¹⁴. The worst case scenario arises in periods where the seller is “paying” for N past negative ratings (i.e. all N ratings summarized in his public profile are negative). To deter a seller from cheating in such periods, it must be $1 - N\theta \geq 0$ (in the next cycle, one negative rating will disappear from the profile; if the seller receives a new negative he will be left with N negatives once again) or, equivalently, $N\theta\hat{w}_2 \leq \hat{w}_2$: the sum of the nominal values of all N punishment installments associated with a bad rating must be less than or equal to a fully cooperating seller’s expected auction revenue. Making an analogy with loans, recall that, given a fixed loan present value, the larger the number of payment installments N , the larger the sum of the nominal values of all installments (due to the impact of interest rates). Similarly, in our setting, given a fixed present value of punishment $c/(\beta - \alpha)$, the sum of the nominal values of punishment installments $N\theta\hat{w}_2$ increases with N . This means that, for given α, β, c , the higher the N , the higher the minimum \hat{w}_2 that satisfies the constraint $N\theta\hat{w}_2 \leq \hat{w}_2$. This, in turn, implies that, as N grows, the minimum profit margin $\rho = \hat{w}_2/c$ that is needed for the reputation mechanism to induce maximum efficiency must also grow, making the mechanism applicable to *fewer* settings.

The conclusion of the above discussion is that the optimal design of sanctioning reputation mechanisms follows different principles to that of signaling reputation mechanisms. The length of time during which a rating persists in the seller’s reputation profile only affects the duration of punishment and not the ability of the mechanism to distinguish between good and bad seller actions. Swifter punishment that lasts for fewer periods is preferable to more lenient punishment that lasts longer, because it reduces the minimum required profit margins and broadens the mechanism’s effectiveness to more market settings. The analysis thus provides arguments for simpler mechanisms that each period publish only a seller’s very recent ratings.

3.4 Implications of costless identity 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

¹⁴If sellers reach a state where they are expected to cooperate with zero probability, they have nothing more to lose by cheating.

“cheap pseudonyms”. This property challenges the effectiveness of reputation mechanisms: online community members can build a reputation, milk it by cheating other members, then disappear and re-enter the community with a new identity and a clean record. I focus my attention to the extreme case where identity changes are cost-free. The following results extend readily to settings where identity changes involve a small positive cost.

In our setting, the ability to costlessly change a player’s identity has a simple but disruptive effect: rational sellers will disappear and re-enter the community with a new identity whenever their reputation profile transitions to a state whose lifetime discounted payoff is lower than the payoff of the initial state of a new seller. This has particularly severe consequences in the case of the two-state randomization device (Section 3.2). If sellers can costlessly change their identity and re-enter the market as soon as their profile transitions into the “bad” state, they have no incentive to avoid negative ratings. Therefore, sellers will always cheat, buyers will expect them to do so, and no trade will ever take place irrespective of the value of ρ . Costless identity changes thus completely destroy the effectiveness of a two-state randomization device.

The only way to prevent sellers from changing their identities is to perturb the mechanism so that new sellers start the game from a state that has the lowest possible payoff among all states. Sellers will then never find it optimal to disappear and re-enter under a new identity because the payoffs of all intermediate states will be at least as high as that of the initial state. Such a perturbation is not possible in the case of the two-state randomization device, because the state with the lowest payoff (the “bad” state) effectively excludes the seller from future participation in the market. On the other hand, if ρ is sufficiently high so that all states have positive payoffs, eBay-like mechanisms can be made robust to identity changes by charging sellers an entrance fee that makes a newcomer’s expected discounted lifetime payoff equal to the payoff of a seller with the “worst” possible reputation profile (i.e. a profile where all summarized ratings are negative). The following proposition provides the details.

Proposition 4: If $\rho = \frac{\hat{w}_2}{c} > \frac{N}{(\beta-\alpha)\sum_{i=1}^N \delta^i}$ then:

1. An eBay-like reputation mechanism with information structure (1) can be made robust to costless identity changes by charging sellers an entrance fee equal to $\left(\frac{N}{\delta(1-\delta^N)} - \frac{1}{1-\delta}\right) \frac{c}{\beta-\alpha}$.

This makes a new seller’s expected discounted lifetime payoff equal to the payoff of a seller whose reputation profile has N negative ratings:

$$V_{free-ids}(N) = \frac{1}{1-\delta} \left(w - c - \frac{\alpha c}{\beta - \alpha} \right) - \left(\frac{N}{\delta(1-\delta^N)} - \frac{1}{1-\delta} \right) \frac{c}{\beta - \alpha}$$

2. New seller payoffs are inversely proportional to the number of ratings N summarized in the mechanism’s public reputation profile; they attain their maximum value for $N=1$:

$$V_{free-ids}(1) = \frac{1}{1-\delta} \left(w_2 - c - \frac{\alpha c}{\beta - \alpha} \right) - \frac{c}{\delta(\beta - \alpha)}$$

Proposition 4 has a number of interesting implications. First, it establishes that costless identity changes impose an efficiency loss because new sellers have to “pay their dues” in order to be trusted to not disappear following the receipt of a bad rating¹⁵. Second, it shows that this efficiency loss grows with the number of ratings N summarized by the mechanism. The loss is minimized in the simplest case where each period the mechanism publishes the rating of the immediately preceding transaction only ($N=1$). This result is remarkable. In Section 3.3 I have shown that the maximum efficiency of an eBay-like reputation mechanism is independent of the number of summarized ratings and that, in fact, summarizing fewer ratings reduces the minimum stage-game profit margin that is required for the mechanism to achieve maximum efficiency. Proposition 4 strengthens this “simpler-is-better” result even more: in environments where players can costlessly change their identities, reputation mechanisms that summarize fewer ratings induce *strictly more efficient outcomes* than mechanisms that summarize larger numbers of ratings.

I conclude this section by showing that, in the presence of free identity changes, no reputation mechanism with information structure (1) can attain higher seller payoffs than an eBay-like mechanism with an entrance fee equal to $c/\delta(\beta - \alpha)$ and $N = 1$. Efficiency losses associated with easy name changes are thus inevitable and a simple two-state reputation mechanism with an entrance fee constitutes the best possible solution.

Proposition 5: If $\rho > (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ and players can costlessly change identities, the set of sequential equilibrium seller payoffs of a repeated game where the stage game structure is described

¹⁵See Friedman and Resnick (2001) for a similar result in the context of a repeated prisoner’s dilemma game.

in Table 1, the reputation mechanism has information structure (1) and the entire public history of buyer reports is available to short-run players is bounded above by:

$$V_{free-ids}^* = \frac{1}{1-\delta} \left(w_2 - c - \frac{\alpha c}{\beta - \alpha} \right) - \frac{c}{\delta(\beta - \alpha)}.$$

4 Multi-valued Reputation Mechanisms

Recall, from Section 2, that a reputation mechanism is defined by (i) the set of supported buyer reports, (ii) the format and transition rules of the public reputation profile, (iii) the initial state of the reputation profile and other rules of entry for new sellers. The analysis of Section 3 derived upper bounds on mechanism efficiency and design guidelines for public reputation profiles in the special case where the mechanism only supports two report types. This section extends the analysis to the general case where the mechanism supports arbitrary finite sets of report types. I also show how the study of the implications and optimal policies for treating missing feedback can be treated as a special case of the general problem of multi-valued reputation mechanism design.

4.1 Upper bounds on mechanism efficiency

The study of binary reputation mechanisms provided valuable intuitions that can help inform the analysis of the multi-valued case. Recall that the power of a binary reputation mechanism relies on the presence of a “bad” report type whose conditional probability of incidence is higher when the seller cheats (β) than when the seller cooperates (α). The mechanism punishes the seller every time a “bad” report is submitted by triggering a reduction of his expected future profits by a present amount p . If the incremental probability of a “bad” report when the seller cheats ($\beta - \alpha$), times the present value of punishment p is greater than or equal to the current gains from cheating (c) then a rational seller will prefer to cooperate. The minimum present value of punishment that induces cooperation is equal to $p = c/(\beta - \alpha)$. Since reputation mechanisms do not levy explicit penalties but simply trigger reductions of future profits, punishment cannot exceed the present value of profits from cooperation. This requires that the incremental probability of a “bad” report ($\beta - \alpha$) be high enough to ensure that the minimum effective punishment is lower than the present value of profits from cooperation.

The efficiency losses of the reputation mechanism relative to the first-best case are due to the less than perfect ability of the mechanism to determine a seller's true action: even cooperating sellers will occasionally receive bad reports and will thus be "unfairly" punished. The average per-period efficiency loss is equal to the conditional probability of a bad report if the seller cooperates (α) times the present value of punishment $c/(\beta - \alpha)$. For a fixed c , this efficiency loss is minimized the smaller the probability of a bad report if the seller cooperates and the higher the incremental probability of a bad report if the seller cheats.

The same logic applies when the mechanism supports multiple report types. The objective of the mechanism still is to deter the seller from choosing the "bad" action (cheat). Any report type, or subset of report types, whose cumulative conditional probability of incidence is higher when the seller cheats can be used as the mechanism's "bad" report set B . The mechanism punishes the seller every time he receives a report that belongs to the "bad" set by an amount equal to the current gains from cheating over the incremental probability of incidence of a bad report if the seller cheats. To minimize efficiency losses, the "bad" report set must consist of report types whose cumulative conditional probability of incidence if the seller cooperates is as small as possible and whose incremental probability of incidence if the seller cheats is as high as possible. The following proposition formalizes this intuition.

Proposition 6: The set of sequential equilibrium payoffs of a repeated game where the stage game structure is described in Table 1, the reputation mechanism supports a report set $R = \{R_i | i = 1, \dots, m\}$ with associated information structure $p(R_i|coop) = \alpha_i$, $p(R_i|cheat) = \beta_i$ and the entire public history of buyer reports is available to short-run players is bounded above by:

$$V^* = \frac{1}{1 - \delta} \left(\hat{w}_2 - c - \frac{a(B)c}{b(B) - a(B)} \right)$$

where $a(B) = \sum_{R_i \in B} \alpha_i$, $b(B) = \sum_{R_i \in B} \beta_i$ and $B \subset R$ is the solution of the following constrained optimization problem:

$$B = \underset{S \subset R, a(S) < b(S)}{\operatorname{argmin}} \frac{a(S)}{b(S)} \quad \text{subject to} \quad \rho \geq \frac{b(B) + \frac{1-\delta}{\delta}}{b(B) - a(B)} \quad (8)$$

If the above problem has no solution then $V^* = 0$.

According to Proposition 6, the "bad" report set B that maximizes efficiency is the subset of supported

reports that minimizes the likelihood ratio of punishment $p(R_i \in B|coop)/p(R_i \in B|cheat)$, while keeping the incremental probability of punishment $p(R_i \in B|cheat) - p(R_i \in B|coop)$ high enough to ensure that the minimum present value of punishment that deters cheating is lower than the present value of future gains from cooperation.

Proposition 6 has a remarkable implication. It says that, even if we allow buyers to use multiple types of ratings (for example, integer ratings between 1 and 5) to express their degree of satisfaction with a transaction, the most efficient way to interpret buyer reports is to classify them into exactly two categories (good/bad) and to treat report types within each category as equivalent (for example, interpret any rating equal to or less than 2 as bad and any rating higher than 2 as good, without attempting to make any further distinctions). Proposition 6 implies that this coarse-grained rule induces average levels of cooperation that are at least as high as those induced by more sophisticated rules that make finer distinctions among report types (see the proof for details)¹⁶.

The above analysis provides a basis for comparing the maximum efficiency of two reputation mechanisms that support different report sets. Specifically, the following corollary is an immediate consequence of Proposition 6:

Corollary 1: A reputation mechanism that supports a report set $R = \{R_i|i = 1, \dots, m\}$ with associated information structure $p(R_i|coop) = \alpha_i$, $p(R_i|cheat) = \beta_i$ induces higher maximum efficiency than a mechanism that supports a report set $R' = \{R'_i|i = 1, \dots, m'\}$ with associated information structure $p(R'_i|coop) = \alpha'_i$, $p(R'_i|cheat) = \beta'_i$ if and only if:

$$\frac{a(B)}{b(B)} < \frac{a'(B')}{b'(B')}$$

where $a(B) = \sum_{R_i \in B} \alpha_i$, $b(B) = \sum_{R_i \in B} \beta_i$, $a'(B') = \sum_{R'_i \in B'} \alpha'_i$, $b'(B') = \sum_{R'_i \in B'} \beta'_i$ and $B \subset R$, $B' \subset R'$ solve:

¹⁶In an academic setting the equivalent result would be that, in a setting where all students are equally capable, a properly designed pass-fail system can induce at least as high student effort as any system with letter grades. Dewatripont, Jewitt and Tirole (1999) report that, under certain conditions, a similar finding (coarser-grained information structures induce agent higher effort) holds even in contexts with adverse selection.

$$\begin{aligned}
B &= \operatorname{argmin}_{S \subset R} \frac{a(S)}{b(S)} & \text{subject to} & \quad \rho \geq (b(B) + \frac{1-\delta}{\delta}) / (b(B) - a(B)) \\
B' &= \operatorname{argmin}_{S' \subset R'} \frac{a(S')}{b(S')} & \text{subject to} & \quad \rho \geq (b'(B') + \frac{1-\delta}{\delta}) / (b'(B') - a'(B'))
\end{aligned} \tag{9}$$

Corollary 1 provides a useful condition that can be used to determine whether, say, the multi-valued reputation mechanism used by Amazon auctions induces higher cooperation levels than the binary mechanism used by eBay. Observe that the efficiency of a sanctioning reputation mechanism does not depend on the number of supported report types but rather on its associated minimum likelihood ratio of punishment. It is, therefore, not always the case that replacing a coarse-grained report set with a finer-grained one will increase efficiency. Once again, we see that the intuitions that are associated with reputation mechanisms whose primary objective is social learning do not apply to the case of sanctioning reputation mechanisms.

The intuitions that led to the above results indicate that maximally efficient implementations of multi-valued sanctioning reputation mechanisms can be constructed as straightforward generalizations of the mechanisms of Sections 3.2 and 3.3. Specifically, either a two-state randomization device that changes the seller’s profile to “bad” with probability π every time the seller receives a rating that belongs to the “bad” report set B , or an eBay-like mechanism that each period publishes the sum of a player’s recent ratings that belong to the “bad” set B , induce equilibria that attain the maximum theoretical efficiency predicted by Proposition 6.

4.2 Application to the study of incomplete feedback submission

Most reputation mechanisms rely on voluntary feedback submission from users. Since this incurs a cost associated with connecting to a website and filling the necessary feedback forms, in the absence of concrete incentives a fraction of users might submit no report to the system¹⁷. In this section I apply the results of the previous section to the study of equilibria induced by binary reputation mechanisms in the presence of incomplete reporting.

¹⁷eBay, for example, currently does not provide concrete incentives for feedback submission. According to recent empirical evidence, only 50% of all eBay transactions receive ratings (Resnick and Zeckhauser 2002). Given the lack of concrete incentives, even this percentage seems remarkably high and invites further study. Initial evidence suggests that the relatively high levels of voluntary participation are primarily driven by the expectation that posting feedback for one’s trading partner will increase the probability that the partner will reciprocate (Dellarocas, Fan and Wood 2003).

The key idea is that failure to provide feedback can be considered as a special type of buyer report. A binary reputation mechanism with incomplete feedback reporting can thus be modeled as a reputation mechanism that supports a 3-valued report set $R = \{+, -, \emptyset\}$, where the symbols $+$, $-$, \emptyset denote positive, negative and no report respectively. Consider the base reputation mechanism described by information structure (1). Let η_+, η_- denote the (exogenously given) probabilities that a buyer who experiences a satisfactory and unsatisfactory transaction respectively will submit feedback to the system¹⁸. The perturbed mechanism can be described by a new conditional probability distribution:

$$\begin{aligned} p(+|coop) &= \eta_+(1 - \alpha) & p(-|coop) &= \eta_-\alpha & p(\emptyset|coop) &= 1 - \eta_+(1 - \alpha) - \eta_-\alpha & 0 < \alpha < \beta < 1 \\ p(+|cheat) &= \eta_+(1 - \beta) & p(-|cheat) &= \eta_-\beta & p(\emptyset|cheat) &= 1 - \eta_+(1 - \beta) - \eta_-\beta \end{aligned} \quad (10)$$

The following theorem is the counterpart of Proposition 1 in the presence of incomplete reporting.

Proposition 7: The set of sequential equilibrium payoffs of a repeated game where the stage game structure is described in Table 1, the reputation mechanism has information structure (10) and the entire public history of buyer reports is available to short-run players is bounded above by:

$$\begin{aligned} V^* &= 0 & \text{if } \rho &< (\beta + \min(\frac{1}{\eta_+}, \frac{1}{\eta_-}) \frac{1-\delta}{\delta}) \frac{1}{\beta-\alpha} \\ V^* &= \frac{1}{1-\delta} \left(\hat{w}_2 - c - \left(\frac{1}{\eta_+} - 1 + \alpha \right) \frac{c}{\beta-\alpha} \right) & \text{if } (\beta + \frac{1-\delta}{\eta_-\delta}) \frac{1}{\beta-\alpha} > \rho \geq \left(\frac{1-\eta_+(1-\beta)}{\eta_+} + \frac{1-\delta}{\eta_+\delta} \right) \frac{1}{\beta-\alpha} \\ V^* &= \frac{1}{1-\delta} \left(\hat{w}_2 - c - \frac{\alpha c}{\beta-\alpha} \right) & \text{if } \rho \geq (\beta + \frac{1-\delta}{\eta_-\delta}) \frac{1}{\beta-\alpha} \end{aligned}$$

We see that, for large ρ , incomplete feedback reporting does not change the maximum achievable payoff of a binary reputation mechanism. However it increases the minimum stage-game profit margin that is required for the mechanism to induce cooperation. Incomplete reporting thus restricts the applicability of reputation mechanisms to environments with higher premiums for trust.

In the rest of the section I will show that equilibria that achieve the maximum payoff prescribed by Proposition 7 can be constructed by augmenting the mechanisms of Sections 3.2 and 3.3 with a policy for treating missing feedback. There are two possible policies:

¹⁸The idea that the probabilities of feedback submission are conditional on the experienced level of satisfaction is consistent with empirical research. For example, Anderson (1998) has shown that consumers are more likely to engage in word of mouth about a product if they had extreme experiences than if they had average experiences.

- Policy 1: Treat missing feedback as positive feedback.
- Policy 2: Treat missing feedback as negative feedback.

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

Proposition 8: Let V_1^* , V_2^* be the maximum seller payoffs attainable in a repeated game where the stage game structure is described in Table 1, the reputation mechanism is either a two-state randomization device or an eBay-like mechanism, the mechanism’s information structure is given by (10) and missing feedback is treated according to Policies 1 and 2 respectively. If $\rho \geq (\beta + (\frac{1}{\eta_-})\frac{1-\delta}{\delta})/(\beta - \alpha)$ the following statement is true:

$$V_1^* = V^* \geq V_2^* \quad \text{Equality holds if and only if } \eta_+ = 1$$

Proposition 8 shows that the most efficient policy for treating missing ratings is Policy 1, an “optimistic” policy that treats missing feedback as positive feedback (“no news is good news”). Pessimistic Policy 2 (“no news is bad news”) is strictly less efficient than Policy 1 if there is incomplete reporting of satisfactory outcomes ($\eta_+ < 1$). The intuition behind this result is that, if $\eta_+ < 1$, some of the missing ratings correspond to satisfactory transactions (transactions that would have received positive feedback). Missing ratings, thus, constitute a more noisy statistic of a transaction’s outcome than negative ratings. Adding them into the “bad” report set can only dilute the ability of the mechanism to distinguish the seller’s action and reduce efficiency because cooperating sellers will then receive higher average unfair punishment. Note that Proposition 8 holds irrespective of the relative magnitudes of η_+ and η_- .

5 Conclusions

This paper offers a theoretical analysis of reputation mechanism design in trading environments where the mechanism’s primary aim is to induce sellers of known ability levels to engage in cooperative behavior in the presence of imperfect monitoring of actions and short-term incentives to cheat. Such

a use of reputation mechanisms is qualitatively distinct from the aggregation of consumer feedback for the purpose of informing a community about a product's or firm's (initially privately known, but fixed) quality attributes. Throughout the paper I use the term *sanctioning* to refer to the former role of reputation mechanisms and *signaling* to refer to the latter.

My analysis derives several important results related to the optimal design of sanctioning reputation mechanisms. These results are often in sharp contrast with the results that apply to the design of signaling reputation mechanisms. Reputation mechanism designers should, therefore, be conscious of their mechanism's primary objective (sanctioning or signaling) and carefully make design choices that maximize the resulting market efficiency, given that objective.

The following is a summary of the paper's principal results.

- *Efficiency loss* (Section 3.1). In environments where seller actions do not completely determine a buyer's rating (because of buyer errors, exogenous factors, etc.) sanctioning reputation mechanisms incur an efficiency loss relative to a hypothetical first-best setting where sellers can credibly pre-commit to full cooperation. Depending on the mechanism's implementation, this efficiency loss arises either from periods of partial cooperation, or from the possibility of "unfair" expulsion of cooperating sellers from the market. The efficiency loss is proportional to the average "unfair" punishment incurred by cooperating sellers due to spurious negative ratings.
- *Minimum profit margin* (Section 3.2). Sanctioning reputation mechanisms can only punish sellers by triggering a reduction of their future profits following bad news about their current behavior. Such mechanisms can, thus, only induce cooperative behavior if the stage-game profit margin of a fully cooperating seller is high enough so that the incremental future profits from cooperation exceed the current gains from cheating.
- *Format of public reputation profile* (Section 3.3). The maximum efficiency induced by a sanctioning reputation mechanism is independent of the length of past history summarized in a seller's public reputation profile. The length of history captured in a seller's profile affects the duration of punishment associated with a bad rating but not the expected present value of

punishment or the mechanism's ability to distinguish a seller's current action. Swifter punishment that lasts for fewer periods is preferable to more lenient punishment that lasts longer, because it reduces the minimum stage-game profit margin required for the mechanism to be effective. A profile that only publishes the single most recent rating minimizes the profit margin requirement and thus maximizes the applicability of the mechanism.

- *Buyer report set* (Section 4.1). Maximally efficient sanctioning reputation mechanisms that allow multiple buyer report types can be constructed by designating a subset of those types as “bad” reports and punishing the seller whenever he receives a report that belongs to the “bad” set. Given a set of reports, the mechanism's maximum efficiency is inversely proportional to the likelihood ratio of punishment if the seller cooperates vs. if he cheats. The “bad” report set should thus be chosen to minimize that likelihood ratio.
- *Missing feedback* (Section 4.2). In settings where feedback submission is voluntary the optimal policy for dealing with missing feedback is to treat it as positive feedback.
- *Free identity changes* (Section 3.4). The ability to freely change identities challenges the effectiveness of reputation mechanisms because sellers can disappear and re-appear with new identities whenever they receive bad ratings. To discourage sellers from such behavior, a reputation mechanism can charge sellers an entrance fee equal to the difference between the payoff of a seller with a clean record and the payoff of a seller with the worst possible reputation profile. Entrance fees reduce a seller's lifetime payoffs. This efficiency loss is an inevitable consequence of free identity changes and cannot be improved by any other approach.

As an application of the above results, the paper offers a theoretical model of eBay's feedback mechanism. The principal conclusion is that eBay's simple mechanism augmented with a summary of a seller's sums of recent positive and negative ratings (such as an eBay member's “ID card”) is capable of inducing the maximum possible efficiency. Moreover, it allows traders to maximize their payoffs by relying on a simple stationary strategy (community norm) according to which a seller's anticipated probability of cooperation in the current period is a decreasing linear function of the sum of negative ratings in his current eBay “ID card”. This property makes eBay's mechanism especially well suited to large-scale online trading environments, as it allows even inexperienced traders to perform well by relying on simple, intuitive behavior rules.

In terms of future work, the analysis of this paper can be extended in a number of dimensions. First, I have assumed only two possible seller actions. An interesting extension is to consider environments with multiple, or continuous, seller actions. Second, I have assumed that seller costs and conditional probabilities of outcomes given effort are completely known to buyers. In some online environments it is plausible that there might be different seller types with different cost structures and/or conditional probabilities of outcomes, initially unknown to buyers. For example, in online marketplaces for professional services, such as *eLance.com*, professionals of varying (and privately known) ability levels advertise their services. In such settings, reputation mechanisms play a hybrid (signaling and sanctioning) role: in addition to eliciting “good conduct”, past feedback should, ideally, also help buyers learn something about the unknown properties (type) of the seller they are facing. The design principles of such hybrid reputation mechanisms constitute an intriguing area for future research.

References

- Anderson, E.W.** (1998) Customer satisfaction and word of mouth. *Journal of Service* 1 (1), 5-17.
- Blackwell, D.** (1951) Comparison of experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press.
- Blackwell, D.** (1953) Equivalent comparison of experiments. *Annals of Mathematical Statistics* 24, 265-272.
- Dellarocas, C.** (2003) The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science* 49 (10), 1407-1424.
- Dellarocas, C., Fan, M., and Wood, C.** (2003) Self-Interest, Reciprocity, and Participation in Online Reputation Systems. 2003 *Workshop in Information Systems and Economics (WISE)*, Seattle, WA.
- Dewatripont, M. Jewitt, I. and Tirole, J.** (1999) The Economics of Career Concerns, Part I: Comparing Information Structures. *Review of Economic Studies* 66 (1), 183-198.
- Friedman, E. and Resnick, P.** (2001) The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy* 10 (1), 173-199.

- Fudenberg, D. and Levine, D. K.** (1994) Efficiency and Observability with Long-Run and Short-Run Players. *Journal of Economic Theory*, 62, 103-135.
- Green, E. J. and R. H. Porter** (1984) Noncooperative Collusion Under Imperfect Price Information. *Econometrica* 52, 87-100.
- Greif, A.** (1993) Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition. *American Economic Review* 83, 525-548.
- Kandori, M. and Matsushima, H.** (1998) Private Observation, Communication and Collusion. *Econometrica*, 66 (3), May 1998, 627-652.
- Klein, B. and Leffler, K.** (1981) The Role of Market Forces in Assuring Contractual Performance. *Journal of Political Economy*. 89 (4), 615-641.
- Klein, D.** ed. (1997) *Reputation: Studies in the Voluntary Elicitation of Good Conduct*. University of Michigan Press, Ann Arbor, MI.
- Kuwabara, K.** (2003) Decomposing Reputation Effects: Sanctioning or Signaling? Working Paper.
- 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, 1-23.
- Miller, N., Resnick, P., and Zeckhauser, R.** (2002) Eliciting Honest Feedback in Electronic Markets, Harvard Kennedy School Research Working Paper RWP02-039.
- Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K.** (2000) Reputation Systems. *Communications of the ACM*, 43 (12), 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. JAI Press.
- Shapiro, C.** (1983) Premiums for High Quality Products as Returns to Reputations. *The Quarterly Journal of Economics*, 98 (4), 659-680.

Appendix

Proof of Proposition 1

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 equilibria 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 K the long-run player's pure action set, ΔK the corresponding space of mixed actions, R the set of (publicly observable) stage-game outcomes and $\mathbf{B} : \Delta K \rightarrow \Delta K_1 \times \dots \times \Delta K_n$ 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 $h(k, \kappa_{SR})$ denote the long-run player's stage-game payoff and $p(r|k, \kappa_{SR})$ denote the probability that the stage-game outcome will be $r \in R$ if the long-run player plays $k \in K$ and the short run players play a mixed action $\kappa_{SR} \in \mathbf{B}(k)$ for some $\kappa \in \Delta K$. Let $\Pi(\kappa_{SR})$ be the matrix with rows corresponding to actions k , columns to outcomes r and with the (k, r) component equal to $p(r|k, \kappa_{SR})$. If a game has a product structure and, in addition, has the property that $\Pi(\kappa_{SR})$ 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:

$$\begin{aligned}
 & \max_{\kappa \in \Delta K, u(r) \in \mathbb{R}} v && \text{subject to} \\
 & v = h(k, \kappa_{SR}) + \delta \sum_{r \in R} p(r|k, \kappa_{SR})u(r) && \text{for } k \in K \text{ such that } \kappa(k) > 0 \\
 & v \geq h(k, \kappa_{SR}) + \delta \sum_{r \in R} p(r|k, \kappa_{SR})u(r) && \text{for } k \in K \text{ such that } \kappa(k) = 0 \\
 & v \geq u(r) && \text{for } r \in R
 \end{aligned}$$

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. This set includes all equilibria where the seller conditions his actions on both public information and private information.

The bilateral exchange game that forms the basis of this paper has $K = \{H, L\}$ (high effort, low effort), $R = \{+, -\}$ (positive report, negative report) 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 $G_i = sw_i$. This results in expected stage-game seller payoffs $h(H, \kappa_{SR}) = s\hat{w}_2 - c$ and $h(L, \kappa_{SR}) = s\hat{w}_2$. Finally,

$$\Pi(\kappa_{SR}) = \begin{bmatrix} p(+|H, \kappa_{SR}) & p(-|H, \kappa_{SR}) \\ p(+|L, \kappa_{SR}) & p(-|L, \kappa_{SR}) \end{bmatrix} = \begin{bmatrix} 1 - \alpha & \alpha \\ 1 - \beta & \beta \end{bmatrix}$$

The above stage game satisfies the full row rank condition (since $Rank\Pi(\kappa_{SR}) = 2$) 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:

$$\begin{aligned} \max_{s \in [0,1], u(+), u(-)} v \quad & \text{subject to} \\ v = s\hat{w}_2 - c + \delta[(1 - \alpha)u(+)] + \alpha u(-) & \quad \text{for } k = H \text{ and } s > 0 \\ v \geq s\hat{w}_2 - c + \delta[(1 - \alpha)u(+)] + \alpha u(-) & \quad \text{for } k = H \text{ and } s = 0 \\ v = s\hat{w}_2 + \delta[(1 - \beta)u(+)] + \beta u(-) & \quad \text{for } k = L \text{ and } s < 1 \\ v \geq s\hat{w}_2 + \delta[(1 - \beta)u(+)] + \beta u(-) & \quad \text{for } k = L \text{ and } s = 1 \\ v \geq u(+), v \geq u(-) & \end{aligned}$$

In addition, intermediate payoffs must satisfy $u(+)$ and $u(-)$ because no seller will stay in the game if he transitions to a state where the present value of future payoffs is negative.

For $\rho < (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ the above problem has solution $v = 0$, $s = 0$, $u(+)$ and $u(-) = 0$, whereas for $\rho \geq (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ the solution becomes:

$$\begin{aligned}
v &= \frac{1}{1-\delta} \left[\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right] \\
s &= 1 \\
u(+)&= v \\
u(-)&= v - \frac{c}{\delta(\beta - \alpha)}
\end{aligned} \tag{11}$$

Proof of Proposition 2

The present value of a cooperating seller's discounted payoff while the seller remains in the “good” state is described by the following Bellman equation:

$$V_{good} = \hat{w}_2 - c + \delta[(1 - \alpha)V_{good} + \alpha[(1 - \pi)V_{good} + \pi V_{bad}]]$$

where $0 \leq \pi \leq 1$, $V_{bad} = 0$. Rearranging, we get:

$$V_{good} = \frac{\hat{w}_2 - c}{1 - \delta(1 - \alpha\pi)} \tag{12}$$

The above is maximized for π equal to the minimum value that induces cooperation. The seller will cooperate if and only if the following incentive compatibility constraint holds:

$$-c + \delta[(1 - \alpha)V_{good} + \alpha(1 - \pi)V_{good}] \geq \delta[(1 - \beta)V_{good} + \beta(1 - \pi)V_{good}]$$

The constraint implies $\pi \geq c/\delta(\beta - \alpha)V_{good}$. Substituting $\pi = c/\delta(\beta - \alpha)V_{good}$ into (12) we obtain $V_{good} = \frac{1}{1-\delta} \left[\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right]$ and $\pi = (1 - \delta)/\delta[\rho(\beta - \alpha) - \beta]$. The requirement $0 \leq \pi \leq 1$ implies $\rho \geq (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$. For $\rho < (\beta + \frac{1-\delta}{\delta})/(\beta - \alpha)$ no $0 \leq \pi \leq 1$ can provide sufficient incentives for seller cooperation. The only equilibrium is one where the seller cheats, no buyers place bids, and $V_{good} = 0$.

Proof of Proposition 3

I will assume that the seller follows a mixed strategy conditioned on the number of negative ratings z in his public reputation profile. Specifically, let:

$$s(z(x)) = 1 - z(x)\theta \quad (13)$$

that is, the seller cooperates as long as there are no negative ratings in his profile and his probability of cooperation decreases by a fixed amount θ with every additional negative rating on his public reputation profile. I will show that, if $\rho = \hat{w}_2/c$ is sufficiently high, then there exists a θ such that the above strategy constitutes an equilibrium strategy that achieves the maximum attainable payoff predicted by Proposition 1.

Let x be the algebraic representation of a seller's private reputation profile that arises from mapping every positive report in his profile to the digit 1 and every negative report to the digit 0. For a mixed strategy, the incentive compatibility constraint (7) simplifies to:

$$V(2(x \oplus 2^{N-1}) + 1) - V(2(x \oplus 2^{N-1})) = c/\delta(\beta - \alpha)$$

or, equivalently, to:

$$V(x) - V(x - 1) = c/\delta(\beta - \alpha) \quad \text{for all odd } x \quad (14)$$

For $0 < s(z(x)) \leq 1$ the seller's Bellman equation (6) can be written as:

$$V(x) = (1 - z(x)\theta)\hat{w}_2 - c + \delta[(1 - \alpha)V(2(x \oplus 2^{N-1}) + 1) + \alpha V(2(x \oplus 2^{N-1}))] \quad \text{for } 0 \leq x \leq 2^N - 1 \quad (15)$$

In the rest of the proof, I will refer to the term $(1 - z(x)\theta)\hat{w}_2 - c$ above as the *stage-game payoff* and to the term $\delta[(1 - \alpha)V(2(x \oplus 2^{N-1}) + 1) + \alpha V(2(x \oplus 2^{N-1}))]$ as the *continuation payoff*. To show that

strategy (13) is an equilibrium strategy, it suffices to show that there exists θ , such that the solution $\{V(x)|x = 0, \dots, 2^N - 1\}$ of the system of Bellman equations (15) satisfies the incentive compatibility constraint (14) for all odd x . I proceed by deriving some key properties of the solution of equations (15):

1. For all $x \geq 2^{N-1}$, states x and $x - 2^{N-1}$ correspond to private profiles that differ from each other in the leading digit only (for example, for $N=4$, if $x=15=1111$ then $x-2^3=15-8=7=0111$). Specifically, state x has a 1 and state $x - 2^{N-1}$ has a 0 in that position. Thus, state $x - 2^{N-1}$ has one additional zero than state x (formally, $z(x - 2^{N-1}) = z(x) + 1$). The stage-game payoffs of the two states thus differ by $\theta\hat{w}_2$. Furthermore, $x \oplus 2^{N-1} = (x - 2^{N-1}) \oplus 2^{N-1}$, which means that the continuation payoffs of states x and $x - 2^{N-1}$ are identical. Intuitively, this situation arises because, following the current period, the leading digit on which the two states differ will be eliminated from the profile, and the rest of the two profiles are identical. By substituting the above into (15) we obtain:

$$V(x) - V(x - 2^{N-1}) = \theta\hat{w}_2 \quad \text{for all } x \geq 2^{N-1} \quad (16)$$

2. For all $x \geq 2^{N-1} + 2^{N-2}$, states x and $x - 2^{N-2}$ correspond to private profiles that differ from each other in the 2nd digit from the left only. Specifically, state x has a 1 and state $x - 2^{N-2}$ has a 0 in that position (for example, for $N=4$, if $x=15=1111$ then $x-2^{N-2}=15-4=11=1011$). Thus, state $x - 2^{N-2}$ has one additional zero than state x (formally, $z(x - 2^{N-2}) = z(x) + 1$). The current period payoffs of the two states thus differ by $\theta\hat{w}_2$. Furthermore, $2((x - 2^{N-2}) \oplus 2^{N-1}) = 2(x \oplus 2^{N-1}) - 2^{N-1}$, and thus, by equation (16), $V(2((x - 2^{N-2}) \oplus 2^{N-1})) = V(2(x \oplus 2^{N-1})) - \theta\hat{w}_2$. Substituting the above into (15) we obtain:

$$V(x) - V(x - 2^{N-2}) = (1 + \delta)\theta\hat{w}_2 \quad \text{for all } x \geq 2^{N-1} + 2^{N-2} \quad (17)$$

3. Continuing in this manner one can show that, for all $x \geq \sum_{i=1}^k 2^{N-i}$, states x and $x - 2^{N-k}$ correspond to private profiles that differ from each other in the k th digit from the left only,

and, furthermore, that:

$$V(x) - V(x - 2^{N-k}) = \left(\sum_{i=0}^{k-1} \delta^i \right) \theta \hat{w}_2 \quad \text{for all } x \geq \sum_{i=1}^k 2^{N-i} \quad (18)$$

4. For $k = N$ the above condition becomes:

$$V(2^N - 1) - V(2^N - 2) = \left(\sum_{i=0}^{N-1} \delta^i \right) \theta \hat{w}_2 \quad (19)$$

If we choose $\theta = c/\delta(\beta - \alpha) \left(\sum_{i=0}^{N-1} \delta^i \right) \hat{w}_2 = c/(\beta - \alpha) \sum_{i=1}^N \delta^i \hat{w}_2$ then comparison of (14) and (19) shows that the solution of the Bellman equations (15) satisfies the incentive compatibility constraint (14) for $x = 2^N - 1$. By using equation (18) one can show that the incentive compatibility constraint is also satisfied for all odd x . For example, to show that the incentive compatibility constraint is satisfied for $x = 2^N - 3$, substitute $k = N - 1$ into equation (18). The equation becomes: $V(2^N - 1) - V(2^N - 3) = V(2^N - 2) - V(2^N - 4) = \left(\sum_{i=0}^{N-2} \delta^i \right) \theta \hat{w}_2$. Rearranging, we obtain: $V(2^N - 3) - V(2^N - 4) = V(2^N - 1) - V(2^N - 2)$. Since we have already established that $V(2^N - 1) - V(2^N - 2) = c/\delta(\beta - \alpha)$ this implies that $V(2^N - 3) - V(2^N - 4) = c/\delta(\beta - \alpha)$, i.e the incentive compatibility constraint holds for $x = 2^N - 3$. Substituting progressively lower values of k into equation (18) and using a similar reasoning we can show that the incentive compatibility constraint is satisfied for all odd x .

Since $s(z) = 1 - z\theta$ represent probabilities, it must be $1 - z\theta \geq 0$ for all $z \leq N$. This requires that $\theta \leq 1/N$, or, equivalently, that $\rho = \hat{w}_2/c \geq N/(\beta - \alpha) \sum_{i=1}^N \delta^i$.

To summarize, if $\rho = \hat{w}_2/c \geq N/(\beta - \alpha) \sum_{i=1}^N \delta^i$, I have shown that there exists an equilibrium where the seller conditions his probability of cooperation on the number of negative ratings in his current profile according to the formula $s(z(x)) = 1 - z(x)\theta$, $\theta = c/(\beta - \alpha) \sum_{i=1}^N \delta^i \hat{w}_2$. The seller's maximum attainable payoff occurs when $x = 2^N - 1$, i.e. when the seller has no negative ratings in his profile. From (15) and (14) we obtain: $V(2^N - 1) = \frac{1}{1-\delta} \left[\hat{w}_2 - c - \frac{\alpha c}{\beta - \alpha} \right] = V^*$.

To show that for $\rho < N/(\beta - \alpha) \sum_{i=1}^N \delta^i$ all public perfect equilibria (PPE) attain efficiency lower than V^* , by Proposition 1 it suffices to show that for such ρ there can be no PPE that attains maximum efficiency equal to V^* . The sketch of the argument is as follows: Assume a general PPE

where $s(z) \geq 0$, $z = 0, \dots, N$. First, by examining Bellman equation (15) we see that the only way to attain $V(2^N - 1) = V^*$ is if: (i) $s(z = 0) = 1$ and (ii) the incentive compatibility constraint $V(2^N - 1) - V(2^N - 2) \geq c/\delta(\beta - \alpha)$ holds with equality for $x = 2^N - 1$. Second, by applying the reasoning steps 1-4 above for arbitrary $s(z)$ we derive the constraint $V(2^N - 1) - V(2^N - 2) = \sum_{i=0}^{N-1} \delta^i (s(0) - s(1)) \hat{w}_2$ equivalent to equation (19). Together with the above, this implies that $s(1) = 1 - \theta$. Continuing in this manner for lower values of x we show that the only PPE that attains $V(2^N - 1) = V^*$ must have $s(z) = 1 - z\theta \geq 0$, $z = 0, \dots, N$. Such an equilibrium cannot exist when $\rho < N/(\beta - \alpha) \sum_{i=1}^N \delta^i$, because then $s(z)$ will become negative for high z , which contradicts our assumption.

Proof of Proposition 4

To discourage identity changes, the entrance fee must make a new seller's expected discounted lifetime payoff smaller than or equal to the payoff of any intermediate state. From the proof of Proposition 3 (equation (18)) we have established that the presence of a negative rating in the k th position (from the left) of a seller's private profile reduces his remaining discounted lifetime payoff relative to a state where there is no negative rating in that position, by an amount equal to $\sum_{i=0}^{k-1} \delta^i \theta \hat{w}_2$. Therefore, the intermediate state with the lowest payoff is a state where the seller's private profile contains negative ratings in all N positions. The payoff difference between that state and a state where the seller's private profile contains no negative ratings is equal to:

$$\sum_{k=1}^N \sum_{i=0}^{k-1} \delta^i \theta \hat{w}_2 = \frac{\sum_{k=1}^N \sum_{i=0}^{k-1} \delta^i}{\sum_{i=1}^N \delta^i} \frac{c}{\beta - \alpha} = \left(\frac{N}{\delta(1 - \delta^N)} - \frac{1}{1 - \delta} \right) \frac{c}{\beta - \alpha}$$

If we assume that new sellers start with a clean profile and we set the entrance fee f equal to the above payoff difference, their expected discounted lifetime payoff will be equal to:

$$V_{free-ids}(N) = V(x = 2^N - 1) - f = \frac{1}{1 - \delta} \left(w - c - \frac{\alpha c}{\beta - \alpha} \right) - \left(\frac{N}{\delta(1 - \delta^N)} - \frac{1}{1 - \delta} \right) \frac{c}{\beta - \alpha} = V(x = 0)$$

Since N grows faster than $1 - \delta^N$ (for $\delta < 1$, it is $(1 - \delta^{N+1}) - (1 - \delta^N) < 1$), the entrance fee f and associated payoff losses grow with N . They attain their minimum value for $N=1$. For $N=1$ the entrance fee becomes equal to $c/\delta(\beta - \alpha)$.

Proof of Proposition 5

The proof is similar to the proof of Proposition 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 $v \geq u(r)$ for $r \in R$ of the linear programming problem of Proposition 1 must thus be replaced with the constraint $v \leq u(r)$ for $r \in R$.

Proof of Proposition 6

The proof is analogous to the proof of Proposition 1 and makes use of Fudenberg and Levine's maximal score method. The corresponding linear programming problem in the case where there are multiple possible buyer reports is:

$$\begin{aligned}
 & \max_{s \in [0,1], u(R_i)} v && \text{subject to} \\
 & v = s\hat{w}_2 - c + \delta \sum_{i=1}^m \alpha_i u(R_i) && \text{for } s > 0 \\
 & v \geq s\hat{w}_2 - c + \delta \sum_{i=1}^m \alpha_i u(R_i) && \text{for } s = 0 \\
 & v = s\hat{w}_2 + \delta \sum_{i=1}^m \beta_i u(R_i) && \text{for } s < 1 \\
 & v \geq s\hat{w}_2 + \delta \sum_{i=1}^m \beta_i u(R_i) && \text{for } s = 1 \\
 & v \geq u(R_i) \geq 0
 \end{aligned} \tag{20}$$

By referring to Proposition 1, we can see that one solution that satisfies the constraints of the problem (but that is not necessarily optimal) can be constructed by splitting the set of report types into two subsets, a bad set $B \subset R$ and a good set $G = R - B$, and constraining the continuation payoffs $u(R_i)$ so that the payoffs associated with any report that belongs to the same subset are equal. The problem then reduces to that of calculating the maximum efficiency of a binary reputation mechanism where the equivalent of the parameter pair α, β is the pair $a(B), b(B)$, defined by $a(B) = \sum_{R_i \in B} \alpha_i$,

$b(B) = \sum_{R_i \in B} \beta_i$. The solution of this problem is an exact counterpart of equation (11):

$$\begin{aligned}
v^* &= \frac{1}{1-\delta} \left[\hat{w}_2 - c - \frac{a(B)c}{b(B)-a(B)} \right] \\
s &= 1 \\
u(R_i) &= v^* && R_i \in B \\
u(R_j) &= v^* - \frac{c}{\delta(b(B)-a(B))} && R_j \in R - B
\end{aligned} \tag{21}$$

if $\rho \geq (b(B) + \frac{1-\delta}{\delta})/(b(B) - a(B))$, and $v^* = 0$ otherwise.

The above solution corresponds to an equilibrium where the seller is punished by an amount $\frac{c}{\delta(b(B)-a(B))}$ every time he receives any report whose type belongs to the set B (and is not punished otherwise). The set B that maximizes the resulting payoff v^* is the subset of R that *minimizes* the quantity $a(B)/(b(B) - a(B))$, or equivalently, the likelihood ratio of punishment $a(B)/b(B)$.

I will now show that we cannot improve upon the above solution by more fine-grained punishment rules that split report types into more than two subsets (e.g. good, bad and “intermediate” reports) and administer different levels of punishment when the seller receives a report that belongs to each subset. Observe that the optimization problem (20) reduces to choosing continuation payoffs $v \geq u(R_i) \geq 0$ that maximize:

$$v = \hat{w}_2 - c + \delta \sum_{i=1}^m \alpha_i u(R_i) \tag{22}$$

subject to the incentive compatibility constraint:

$$\delta \sum_{i=1}^m (\alpha_i - \beta_i) u(R_i) = c \tag{23}$$

Consider a partition of the set of report types R into 3 subsets: Good report types G with corresponding continuation payoffs $u(R_i|R_i \in G) = v$, bad report types B with continuation payoffs $u(R_i|R_i \in B) = v' < v$ and intermediate report types I with continuation payoffs $u(R_i|R_i \in I) = v'' < v$. As before, let $a(B) = \sum_{R_i \in B} \alpha_i$, $b(B) = \sum_{R_i \in B} \beta_i$. Equation (23) can be rewritten as:

$$\delta \left(\sum_{i|R_i \in G} (\alpha_i - \beta_i)v + \sum_{i|R_i \in I} (\alpha_i - \beta_i)v'' + (a(B) - b(B))v' \right) = \delta(b(B) - a(B))(v - v') = c$$

where $\hat{v} = (\sum_{i|R_i \in G} (\alpha_i - \beta_i)v + \sum_{i|R_i \in I} (\alpha_i - \beta_i)v'') / \sum_{i|R_i \in G \cup I} (\alpha_i - \beta_i) < v$ and $\sum_{i|R_i \in G \cup I} (\alpha_i - \beta_i) = (1 - a(B)) - (1 - b(B)) = b(B) - a(B)$. The above equation gives $v' = \hat{v} - c/\delta(b(B) - a(B))$ which shows that v' grows with \hat{v} . Therefore, if we replace the continuation payoff v'' associated with the intermediate report types with the payoff v associated with good report types (effectively collapsing the distinction of report types into good and bad only) then the bad report continuation payoff v' that satisfies the incentive compatibility constraint (23) grows, and, therefore, the value of the objective function (22) also grows because the continuation payoffs of all states are greater than or equal to the corresponding continuation payoffs of the original solution. Therefore, any equilibrium that involves more than two sets of continuation payoffs associated with different report types is dominated by an equilibrium where “intermediate” report types are considered as good report types. The maximum payoff (21) that is attainable by considering two sets of continuation payoffs (the payoff that arises when we choose the bad report set so that the likelihood ratio $a(B)/b(B)$ is minimized) is thus the maximum payoff attainable by any public perfect equilibrium of the game (and, by Theorem 5.2 of Fudenberg and Levine (1994), by any sequential equilibrium of the game).

Proof of Proposition 7

The proof is an application of Proposition 6 in the special case where the mechanism’s information structure is given by (10). The three cases of the resulting expression arise as follows:

For $B = \{+, -, \emptyset\}$, the set of non-empty proper subsets $S \subset R$ and their associated probabilities of incidence $a(S)$, $b(S)$ is summarized below:

S_i	$a(S_i)$	$b(S_i)$	$b(S_i) - a(S_i)$
$S_1 = \{+\}$	$\eta_+(1 - \alpha)$	$\eta_+(1 - \beta)$	$\eta_+(\alpha - \beta)$
$S_2 = \{-\}$	$\eta_- \alpha$	$\eta_- \beta$	$\eta_-(\beta - \alpha)$
$S_3 = \{\emptyset\}$	$(1 - \eta_+)(1 - \alpha) + (1 - \eta_-)\alpha$	$(1 - \eta_+)(1 - \beta) + (1 - \eta_-)\beta$	$(\eta_+ - \eta_-)(\beta - \alpha)$
$S_4 = \{+, \emptyset\}$	$1 - \eta_- \alpha$	$1 - \eta_- \beta$	$\eta_-(\alpha - \beta)$
$S_5 = \{+, -\}$	$\eta_+(1 - \alpha) + \eta_- \alpha$	$\eta_+(1 - \beta) + \eta_- \beta$	$(\eta_- - \eta_+)(\beta - \alpha)$
$S_6 = \{-, \emptyset\}$	$1 - \eta_+(1 - \alpha)$	$1 - \eta_+(1 - \beta)$	$\eta_+(\beta - \alpha)$

Define $l_i = a(S_i)/b(S_i)$ and $\rho_i = (b(S_i) + \frac{1-\delta}{\delta})/(b(S_i) - a(S_i))$. It is easy to see that if $l_i < l_j$ and

$b(S_i) - a(S_i) > b(S_j) - a(S_j)$ then $\underline{\rho}_i < \underline{\rho}_j$. I will now proceed to solve the constrained optimization problem (8) for the above set of parameters:

1. Subsets S_1 and S_4 are eliminated from further consideration because, for $\alpha < \beta$, it is $b(S_i) < a(S_i)$.
2. If $\eta_+ > \eta_-$ then subset S_5 is likewise eliminated. From among the remaining three subsets, $l_2 \leq l_6 \leq l_3$. Furthermore the minimum required stage-game profit margins satisfy $\underline{\rho}_3 \geq \underline{\rho}_6$. In addition, if η_+ is sufficiently higher than η_- , it is also $\underline{\rho}_2 \geq \underline{\rho}_6$. These relationships imply that (i) if $\rho \geq \underline{\rho}_2$ then the mechanism's maximum efficiency is determined by the likelihood ratio l_2 , (ii) if $\underline{\rho}_2 > \rho \geq \underline{\rho}_6$ then the mechanism's maximum efficiency is determined by the likelihood ratio l_6 , and (iii) if $\rho < \underline{\rho}_6$ then the constrained optimization problem that specifies the mechanism's efficiency has no solution, hence $V^* = 0$.
3. If $\eta_+ < \eta_-$ then subset S_3 is likewise eliminated. From among the remaining three subsets, $l_2 \leq l_6$ (equality iff $\eta_+ = 1$) and $l_2 \leq l_5$ (equality iff $\eta_+ = 0$). Furthermore the minimum required stage-game profit margins satisfy $\underline{\rho}_2 \leq \underline{\rho}_6$ and $\underline{\rho}_2 \leq \underline{\rho}_5$. These relationships imply that (i) if $\rho \geq \underline{\rho}_2$ then the mechanism's maximum efficiency is determined by the likelihood ratio l_2 , (ii) if $\rho < \underline{\rho}_2$ then the constrained optimization problem that specifies the mechanism's efficiency has no solution, hence $V^* = 0$.

If we replace the corresponding algebraic expressions for l_2 and l_6 into (8) we obtain the expressions listed in Proposition 7.

Proof of Proposition 8

From Corollary 1, comparison of the maximum efficiency of two reputation mechanisms reduces to the comparison of their corresponding likelihood ratios of punishment. Binary reputation mechanisms augmented with Policies 1 and 2 are equivalent to 3-valued reputation mechanisms with bad report sets equal to $B_1 = \{-\}$ and $B_2 = \{-, \emptyset\}$ respectively. Given information structure (10) the

corresponding likelihood ratios of punishment can be constructed by observing that:

$$\begin{aligned} a(B_1) &= \eta_- \alpha & a(B_2) &= 1 - \eta_+(1 - \alpha) \\ b(B_1) &= \eta_- \beta & b(B_2) &= 1 - \eta_+(1 - \beta) \end{aligned}$$

Therefore:

$$l_1 = \frac{a(B_1)}{b(B_1)} = \frac{\alpha}{\beta} \quad l_2 = \frac{a(B_2)}{b(B_2)} = \frac{1 - \eta_+(1 - \alpha)}{1 - \eta_+(1 - \beta)}$$

It is easy to see that $l_1 \leq l_2$ with equality if and only if $\eta_+ = 1$. According to Corollary 1, this implies that the maximum efficiencies induced by the two mechanisms satisfy $V_1^* \geq V_2^*$ with equality if and only if $\eta_+ = 1$. Finally, since the likelihood ratio l_1 is equal to the likelihood ratio of punishment in the baseline case where all buyers submit feedback, from Corollary 1 it follows that $V_1^* = V^*$.