Giving Gossips Their Due:
Information Provision in Games with Private Monitoring

Robert S. Gazzale*
rgazzale@williams.edu
Department of Economics
Williams College

January 2009

Abstract

The ability of a long-lived seller to maintain and profit from a good reputation may induce her to provide high quality or effort despite short-run incentives to the contrary. This incentive remains in place with private monitoring, provided that buyers share their information. However, this assumption is unrealistic in environments where information sharing is costly or the beneficiaries of a buyer's sharing are strangers. I study a simple mechanism that induces costly information provision, and may explain such behavior in environments where the incentives are not overt. Agents who possess information may share it with the community and acquire a reputation for *gossiping*. Reputations function in tandem: sellers provide high effort because they face agents with reputations for information sharing, and expect the outcome of their dealings will be made public, while information holders share their information as a reputation for doing so results in higher effort from sellers.

1 Introduction

The promise of electronic commerce is alluring. Taking advantage of increased interconnectivity and processing power, the Internet enables physically distant buyers and sellers to

*Department of Economics, Williams College, 24 Hopkins Hall Drive, Williamstown, MA 01267; rgazzale@williams.edu. I gratefully acknowledge the helpful comments and suggestions of Yan Chen, Emma Hutchinson, Kai-Uwe Kühn, Jeffrey MacKie-Mason, Katya Malinova, Yo Nagai, Andreas Pape, Paul Resnick, Barbara Tran and seminar participants at the University of Michigan, Amherst College and Williams College. I the Horace H. Rackham School of Graduate Studies for financial support.

find each other and transact, permitting otherwise infeasible mutually beneficial pairings.[1] The massive increase in potential trading partners brings to the fore an age-old question: Whom should we trust? Physical distance often makes pre-transaction inspection infeasible, and payment is often required before delivery. Distance, combined with the relative ease of building a credible electronic commerce facade, restricts the use of visible clues of a seller's trustworthiness.[2] Finally, many relationships are one-shot affairs, blunting the threat to withhold future business.

When interactions between particular buyers and sellers are infrequent, repeat interaction with a particular market may substitute for repeated interaction with a particular buyer. The ability of a seller to develop, and profit from, a good reputation may explain trustworthy behavior such as high post-transaction effort in these environments. Indeed, many on-line marketplaces have organized reputation mechanisms, the most notable being eBay's feedback system. Taking advantage of the same forces that facilitate the on-line search for a trading partner, a market can collect information about a particular seller, and, at nearly zero marginal cost, aggregate it in a meaningful way and transmit it to all interested parties. Thus, in the same way a local merchant's good standing in the community may serve as a bond of her trustworthiness, so may a firm's on-line reputation.

Previous work has shown that with sufficient information, reputations can be quite effective in inducing cooperative (or trustworthy) behavior when particular agents interact infrequently. For example, Kandori (1992) looks at games with perfect private monitoring. He shows that with minimal information requirements, as long as players voluntarily share their information, any feasible individually rational payoff is robustly supported despite the lack of repeat interactions between particular agents. In an experimental setting, Bolton et al. (2004) show that accurate information about the trustworthiness of sellers induces greater trust and trustworthy behavior than a market in which buyers do not have such information.[3] In their study, providing this information does not generate the same levels of trust and trustworthy behavior as repeated play with the same partners. In his study of eBay-like reputation mechanisms in environments in which a buyer's outcome is an imperfect signal of hidden seller actions, Dellarocas (2004) shows maximal efficiency can be attained by providing buyers with a simple statistic of recent seller behavior.

In order for a seller to develop a meaningful reputation, information about the seller's dealings must be made public. In this area, electronic commerce faces some unique challenges. For example, consider the proprietor of a wedding catering business in a community. While she may not expect too many future transactions with a particular buyer, she still has incentives to provide high quality services, as she reasonably believes that others in

[1] For example, eBay reports that 34.7 million users participated in at least one auction in the twelve months ending September 30, 2003.

[2] Klein and Leffler (1981) highlight the ability of firm-specific sunk costs such as advertising to help consumers decide whom to trust. More generally, Kihlstrom and Riordan (1984) present a model in which a firm may use advertising to signal high quality. See Bagwell (2001) for a thorough overview of the literature on advertising and signalling.

[3] However, in Bolton et al. (2004), the levels of trust and trustworthy behavior induced by the introduction of a reputation mechanism were lower than achieved in a treatment with repeated play by the same partners.

the community will become aware of her efforts. First, in game theoretic terms, her relationship with a particular couple has aspects of public monitoring: attendees at the affair observe a signal of her effort level. This contrasts with the private monitoring that marks the typical electronic commerce transaction: the actual buyer may be the only recipient of signals of the seller's efforts. Second, our local caterer may expect to be the subject of local *gossip*. Prospective customers may know and ask the caterer's previous clients about their experiences. More subtly, information may be passed along in the normal course of community interaction, as Mr. Jones raves to his neighbor about how fabulous last night's Martinez wedding was. An electronic marketplace can of course take steps to facilitate information sharing. However, without positive incentives, information is still likely to be under-provided. By paying the costs associated with providing information, the buyer is contributing to a public good whose beneficiaries are likely to be strangers to him.

Most theoretical work studying reputations assumes that games are either public monitoring,[4] or if private monitoring, that all buyers contribute their private information about sellers to the reputation mechanism, without specifying explicit incentives for them to do so. There have been a few papers that study mechanisms or institutions that elicit an agent's information about trading partners. Avery et al. (1999) consider the case in which a product's true quality is unknown. They propose an intuitive payment-based mechanism. Those who benefit most from the community's information (i.e., those who try a product later) compensate those who take the bigger risks (i.e., those who try a product first). Klein (1992) studies the information sharing role of credit bureaus. In his model, a firm decides whether or not to join a credit bureau. If it joins, it must provide information about the creditworthiness of its customers to the bureau.[5] In return, the firm receives information about the creditworthiness of new customers, which enables it to avoid extending credit to those who will not repay. Finally, Milgrom et al. (1990) study the role of law merchants in assuring performance in medieval trade where contracts were unenforceable. In this model, agents supply information because only by doing so can they expect restitution if they are cheated. Those who cheat pay restitution to reinstate their good reputation.

I propose a simple mechanism, feasible in many on-line environments, in which agents voluntarily submit their private information to some central reputation mechanism. Just as a seller may provide costly effort in order to develop a reputation for doing so, a buyer may take the costly action of information sharing in order to develop a reputation for doing so. The reputations function in tandem: a seller provides effort because she faces agents with reputations for information sharing, and expects the outcomes of her dealings will be made public.[6] A buyer shares his information because the resulting reputation for doing so results in more cooperation from sellers.[7]

[4]The ability of an incumbent to develop a reputation for fighting entry in Kreps and Wilson's 1982 study of the chain-store paradox is a perfect example of public monitoring.

[5]In this model, there are no costs to providing information, but rather a cost to join the bureau.

[6]I use feminine pronouns for sellers, and masculine for buyers.

[7]This information structure contrasts with the model studied by Klein (1992). He assumes that firms do not develop a reputation for sharing credit information with the credit bureau.

In this paper, I make several contributions to the study of games with private monitoring. First, I offer guidance to those charged with the design of reputation mechanisms in on-line settings. In environments in which altruistically provided information is insufficient and payment-based schemes are infeasible or undesirable, allowing the development of information-sharing reputations can be quite effective. Second, I help explain the provision of costly evaluations in environments in which the incentives to do so are not overt. After all, it would be foolish for a firm not to provide promised high effort to a known gossip, and knowing this, a community member might desire to cultivate that reputation.[8] Finally, this paper makes a contribution to the theoretical understanding of repeated games of private monitoring. I show that information can be structured so that players do provide their private information even when it is costly to do so. In games of perfect private monitoring, if players are sufficiently patient, the basic mechanism can be modified such that equilibrium evaluation costs can be made arbitrarily small. Therefore the set of feasible and individually rational payoffs that are attainable under costly evaluations can be made arbitrarily close in size to the set attainable were evaluation costs zero. In games in which a buyer's outcome does not perfectly reveal the seller's effort level, however, there is a trade-off. While the costs of information sharing are increasing in its frequency, the benefits, in the form of overall seller effort, are increasing as well.

2 The Model

Consider an environment in which there are $2N$ players, indexed $n = \{1, 2, \ldots, 2N\}$. Each player has a role ρ. N players are buyers ($\rho = b$), while N players are sellers ($\rho = s$). I let N_ρ represent the set of role ρ players. In each period, each seller $\hat{s}$ is randomly matched with a buyer denoted by $\hat{b} = \nu(\hat{s}, \tau)$ to play a stage game, where $\nu(n, \tau)$ is some function that determines each player's match in each period, τ.[9] Each player is matched with exactly 1 other player in each period.

Each stage game consists of two distinct subgames. The first is the *trading game*. If both players agree to transact, the buyer pays an amount P to the seller, where P is exogenously determined.[10] I let t_n represent player n's decision about whether or not to transact: $t_n = 1$ if he agrees to transact, $t_n = 0$ if not. If the players agree to transact, the seller provides either high effort ($h_{\hat{s}} = 1$) at a cost to her of $c > 0$, or low effort ($h_{\hat{s}} = 0$) at a cost normalized to zero. The effort level selected by the seller affects the outcome experienced by the buyer, denoted $z_{\hat{b}}$. If the seller exerts high effort, the buyer has a good outcome ($z_{\hat{b}} = 1$) with probability λ, and a poor outcome ($z_{\hat{b}} = 0$) with probability $1-\lambda$. If the seller exerts low effort, the probabilities of good and poor outcomes are μ and $1-\mu$ respectively. To capture the fact that buyers value effort, I assume that good outcomes are more likely if the seller

[8]If I know that the local restaurant is aware of John's reputation as a truthful gossip, I must take John's positive review with a grain of salt. This is why many restaurant reviewers guard their anonymity. Of course, it would be a very strong signal if John got a poor meal despite his reputation.

[9]Throughout the paper, I let $\hat{s}$ refer to an arbitrary seller and $\hat{b}$ to an arbitrary buyer.

[10]I assume that contracts that specify outcome-based payments are unenforceable.

provides high effort ($\lambda > \mu$). I normalize buyer payoffs such that the value to the buyer of a good outcome is 1, while the value of a poor outcome is 0.

I assume private monitoring. That is, the signal received by a player is visible only to that player.[11] Players may, however, decide to reveal their information in the evaluation game that directly follows the trading game. In the evaluation game, player n may provide an evaluation to the community ($e_n = 1$) at a cost of $k > 0$. Alternatively, if he decides not to provide an evaluation ($e_n = 0$), he avoids this cost. It is most natural to think of buyers evaluating sellers. However, I consider cases in which an equilibrium strategy calls for certain buyers to allow themselves to be punished. In such cases, a community may need the seller to make public the buyer's actions.

I represent the stage game actions of seller $\hat{s}$ with $a_{\hat{s}} = \{t_{\hat{s}}, h_{\hat{s}}, e_{\hat{s}}\}$, buyer $\hat{b}$'s with $a_{\hat{b}} = \{t_{\hat{b}}, e_{\hat{b}}\}$, the actions of both matched players with $a = \{a_{\hat{s}}, a_{\hat{b}}\}$, and the set of possible actions with A. Player n's mixed action is denoted by α_n. I define $BR(\alpha_n)$ as player $\nu(n, \tau)$'s stage-game best response to α_n.

Expected *trading* game payoffs for a pair, $f_n(a), n = \hat{b}, \hat{s}$, are:

$$f_{\hat{s}}(a) = t_{\hat{b}} t_{\hat{s}} (P - h_{\hat{s}} c);$$
$$f_{\hat{b}}(a) = t_{\hat{b}} t_{\hat{s}} \left(-P + h_{\hat{s}} \lambda + (1 - h_{\hat{s}}) \mu\right).$$

Expected payoffs for the entire stage game (trading and evaluation games), $g_n(a), n = b, s$, are:

$$g_{\hat{s}}(a) = t_{\hat{b}} t_{\hat{s}} (P - h_{\hat{s}} c) - e_{\hat{s}} k; \tag{1}$$
$$g_{\hat{b}}(a) = t_{\hat{b}} t_{\hat{s}} \left(-P + h_{\hat{s}} \lambda + (1 - h_{\hat{s}}) \mu\right) - e_{\hat{b}} k. \tag{2}$$

A player can guarantee himself a stage-game payoff of 0 by neither transacting nor evaluating. Let $g_n(a^m) = 0$ represent player n's minimax payoffs. I consider two sets of payoffs that are feasible and individually rational in the sense that they are at least as large as the minimax payoffs. The first set is feasible and individually rational *trading* game payoffs: $V = \{v \in \text{co } f(A) | v \gg 0\}$ where co $f(A)$ is the convex hull of the set $f(A)$. I represent the action profile that achieves a particular $v^* \in V$ with a^*. In order to achieve payoffs v in the interior of this set, players may need to play a correlated strategy.[12] I assume a correlation device is available to the players. Second, I define the set of feasible and individually rational *stage*-game payoffs, which includes both the trading and the evaluation games. Note that if a correlated strategy profile calls for the buyer to not provide

[11] Of course, if an action always produces a particular outcome ($\lambda = 1$ or $\mu = 0$), then the seller knows the buyer's outcome if she selects that action.

[12] For example, they may need to *jointly* play $a' = \{a'_{\hat{b}}, a'_{\hat{s}}\}$ with a certain probability and $a'' = \{a''_{\hat{b}}, a''_{\hat{s}}\}$ rest of the time in order to receive v in expectation. Note that this is not the same as mixed strategies, where buyers independently mix between $a'_{\hat{b}}$ and $a''_{\hat{b}}$ while sellers are independently mixing between $a'_{\hat{s}}$ and $a''_{\hat{s}}$.

an evaluation, he then knows that the seller will not provide effort, and will therefore not trade at a price above μ. Therefore, the set of feasible and individually rational *stage* game payoffs is $V^{\dagger} = \{v \in \text{co } g(a) | v \gg 0, a_n \in BR(a_{-n}) \text{ if } e_{\hat{b}} = 0 \text{ for } n = \{b, s\}\}$ where co $g(A)$ is the convex hull of the set $g(A)$. The two sets coincide under the assumption of costless evaluation provision, as $f_n(a) = g_n(a)$. In this study, I compare the two sets under costly evaluations.

The stage game, including the random matching and both subgames, is repeated infinitely. For a sequence of payoffs $\{g_n(\tau)\}$, a player's average discounted per-period payoffs are

$$(1 - \delta) \sum_{\tau=1}^{\infty} \delta^{\tau-1} g_n(\tau),$$

where $\delta \in (0, 1)$ is the common discount factor.

I make the following assumptions about the provision of evaluations.

Assumption 1. *The individual direct costs or benefits associated with providing an evaluation are independent of the content of the evaluation.*

Assumption 2. *If a player provides an evaluation, he does so truthfully.*

Assumption 1 states that the content of an evaluation in no way affects the incentives to provide an evaluation. One can imagine plausible scenarios in which this might not hold. For example, one might receive a psychic benefit for providing either a positive or negative evaluation. Providing a negative evaluation after a bad experience might be particularly satisfying. In other cases, however, particularly in cases in which the distinction between positive and negative outcomes is more ambiguous, a buyer might be hesitant to leave a negative evaluation, heeding his mother's advice about not saying anything unless he has something good to say. In the mechanisms I detail below, future sellers know only whether or not a buyer provides evaluations, and not the content of these evaluations. Buyers therefore have no incentive to report dishonestly, and Assumption 2 states that if a buyer is indeed indifferent, he reports honestly.[13]

A community will have some convention for turning information available about an agent into her *reputation*. For example, eBay gives each participant a *feedback score* which is the summation of 1 for every user who leaves a positive review and -1 for every negative.[14] In order to study the effects of reputations on equilibrium payoffs and evaluation provision, I adapt the simple yet powerful reputation framework first introduced by Okuno-Fujiwara and Postlewaite (1995) and further studied by Kandori (1992).[15]

[13] Miller et al. (2005) present a mechanism for eliciting honest feedback in environments in which the evaluated object's type is unknown, i.e., environments of adverse selection.

[14] In addition, eBay also presents a summary table of feedbacks in the previous 6 months, as well as a listing of every feedback, including comments.

[15] Kandori (1992) cites an earlier working paper version of the Okuno-Fujiwara and Postlewaite (1995) paper.

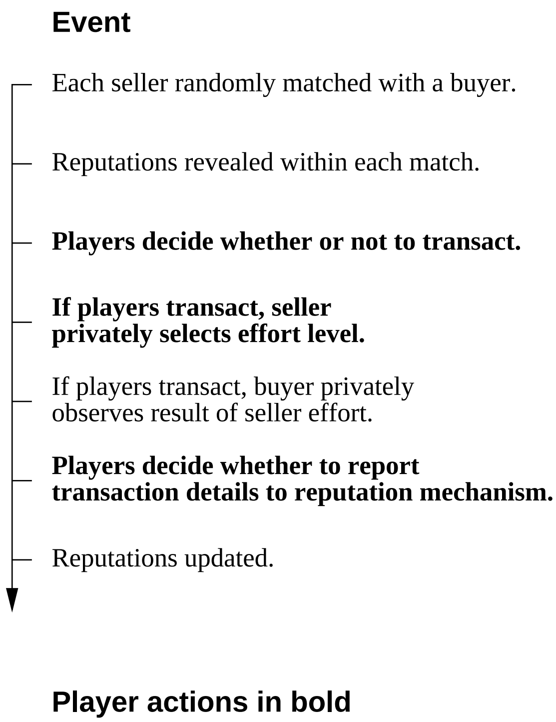

Figure 1: Timing of the stage game.

Definition 1. A matching game with **extended local information processing** has the following information structure.

1. An updating phase $x_n \in \{1, 2, \ldots, T_\rho\}$ is permanently assigned to each player $n \in N_\rho$, $\rho = b, s$.
2. A reputation $r_n(\tau) \in R_\rho$ is assigned to each $n \in N_\rho$, $\rho = b, s$, at time τ.
3. When player $\hat{s}$ and $\hat{b}$ meet at time τ and take actions $(a_{\hat{s}}(\tau), a_{\hat{b}}(\tau))$, their next states are determined by

$$r_n(\tau+1) = \begin{cases} r_n(\tau) & \text{if } (\tau + x_n) \bmod T_\rho \neq 0 \\ \phi_\rho(\cdot) & \text{if } (\tau + x_n) \bmod T_\rho = 0, \end{cases}$$

where

$$\phi_\rho(\cdot) = \phi_\rho(r_n(\tau), a_n(\tau - T_\rho + 1), \ldots, a_n(\tau), a_{\nu(n,t-T_\rho+1)}(\tau - T_\rho + 1), \ldots, a_{\nu(n,\tau)}(\tau))$$

4. At time τ, player n can observe at least $(r_n(\tau), r_{\nu(n,\tau)}(\tau))$, but not $x_{\nu(n,\tau)}$, before choosing her action.

A feature of extended local information processing is that players may keep the same reputation over many periods. His next reputation is then based on his and his partners' actions over these periods. This contrasts with Kandori (1992), who assumed that each player's reputation cycle lasts one period ($T_\rho = 1$). Second, I endogenize the choice of evaluating, but do assume that evaluations are truthful if provided. This means that the mechanism must be able to assign reputations when, for example, a buyer does not provide an evaluation of the seller. I present an outline of the stage game in Figure 1.

In deciding upon a strategy, a player may, in certain environments, have at his disposal information in addition to the reputation of his trading partner. I consider equilibria that are robust to details of information structure. In particular, I look at *straightforward* equilibria, as introduced by Kandori (1992). In a straightforward equilibrium, player actions are determined solely by the *labels* associated with each player in the pair, which in this setting will be their reputations. Strategies, a function of reputations, are thus given by the following functions:

$$\alpha_n(t) = \sigma_\rho(r_n(t), r_{\nu(n,t)}(t)) \text{ for } n \in N_\rho, \rho = b, s.$$

It is also desirable that an equilibrium be robust. Players, particularly newcomers to a market or community, might make mistakes. An equilibrium achieving payoffs v is *globally stable* if for any finite history of actions $\{a\}$,

$$\lim_{\tau \to \infty} E(v_n(\tau)|\{a\}) = v_\rho \quad \forall n \in N_\rho, \rho = b, s.$$

If an equilibrium is globally stable, then for any "initial" distribution of reputations, payoffs converge to v for all players when players play equilibrium strategies from that point on. This requirement rules out, for example, equilibria that are supported by punishing deviators for an infinite number of period.

Finally, in certain cases, I shall consider different player types. With probability γ, a buyer is an *altruist.* I assume that an altruist always provides an evaluation. All others, both buyers and sellers, are risk-neutral *strategic* players whose objective is to maximize the average discounted value of per-period payoffs. Consistent with the adverse-selection literature, a player's type is only observable to himself.

3 Games with Perfect Private Monitoring

In games with perfect monitoring, the signal that each player receives perfectly informs him about the action of the other player. This means that high effort always produces a good outcome ($\lambda = 1$), while low effort always produces a poor outcome ($\mu = 0$).[16] While the buyer knows with certainty which effort level the seller selected based on his outcome, the private monitoring assumption means that only the matched buyer and seller know the outcome.

[16]Perfect monitoring has the same effect as observable actions.

As a benchmark, I start by considering outcomes possible if all buyers provide costless evaluations. This is the case considered by Kandori (1992). As demonstrated by the next proposition, he shows that all individually rational payoffs are possible.

Proposition 1 (Kandori (1992))**.** *Under the assumption of costless truthful evaluation provision, every point $v^* \in V$ is sustained by a straightforward and globally stable equilibrium with extended local information processing if $\delta \in (\delta^*, 1)$ for some δ^*.*

Proof. By individual rationality, it must be the case that $P \in [c, 1]$ for $v \gg 0$. Sellers would never provide effort if the price were lower, buyers would never agree to transact were it higher. As evaluation costs are assumed to be zero, for all $P \in [c, 1]$, if buyers believe that sellers will provide high effort, buyers agree to transact. I now show that sellers will provide high effort.

Let $\{a^*, P\}$ be such that $v = v^*$. Strategies are given as follows:

$$\sigma(r) = \begin{cases} a^* & \text{if } r_{\hat{s}} = 0 \\ a^m & \text{if } r_{\hat{s}} \neq 0, \end{cases}$$

where a^m is any strategy in which the buyer does not transact and thus gives the seller her minimax payoff. Set each buyer's reputation equal to 0, and let each seller's reputation be determined by the following:

$$\phi_s(r, a) = \begin{cases} 0 & \text{if } r_{\hat{s}} = 0 \text{ and } a_{\hat{s}} = a^* \\ 1 & \text{if } a_{\hat{s}} \neq \sigma_s(r) \\ (r_{\hat{s}} + 1) \bmod (X + 1) & \text{if } r_{\hat{s}} \neq 1 \text{ and } a_{\hat{s}} = \sigma_s(r) \end{cases}$$

For a given δ sufficiently large, set $X < \infty$ such that the maximal gains from deviating, namely not incurring effort cost c, are offset by X periods of zero payoffs. As v is positive by individual rationality, such an $X < \infty$ exists for δ sufficiently large. As $X < \infty$, after X periods, all players receive v^* regardless of starting history, therefore the equilibrium is globally stable. ■

Proposition 1 highlights the ability of local information processing to sustain high levels of "cooperative" effort. In equilibrium, if a seller cheats, she gets a bad reputation. Buyers punish her by holding her to her minimax payoff for the next X periods. As long as she cares enough about these future periods, this punishment is severe enough to outweigh the gains from not providing high effort, c. Thus, even without repeated interactions between particular agents, and with very little information transmitted between them, folk theorem-like results hold: any individually rational payoff possible in the stage game can be an equilibrium outcome of the infinitely repeated game if players are sufficiently patient.[17] The next proposition confirms the intuition that, in games of private monitoring with costly information sharing, strategic buyers need to be given some incentives in order to provide the information necessary for high effort equilibria.

[17] Fudenberg and Maskin (1986) show this is the case when the same agents play the infinitely repeated game.

Proposition 2. *If all buyers are strategic buyers and* $\phi(\{a_s^*, t_{\hat{b}}^*, e_{\hat{b}}\!=\!1\}, r, \tau) = \phi(\{a_s^*, t_{\hat{b}}^*, e_{\hat{b}}\!=\!0\}, r, \tau)$ *for* $\sigma(r) = a^*$ *and all* τ*, then the only straightforward equilibrium is the stage-game Nash equilibrium of no seller high effort.*

Proof. Assume not. Let a^* be the action profile that achieves $v^* \in V^\dagger$ and r^* a reputation pair in which this strategy is played.

First, if a buyer is required to provide an evaluation in r^* he can profitably deviate by not providing an evaluation. He earns k by not providing an evaluation. His future payoffs are not affected as, by assumption, both his and his seller's reputation is independent of whether or not he provides an evaluation. Therefore, buyers will not provide evaluations.

Under no evaluation provision, a seller's reputation next period is independent of whether she provides effort this period. As strategies are based only on reputations, a seller's best response is to provide no effort. As this is true for an arbitrary state, it is true for all. ■

Proposition 2 says that if reputations cannot signal the fact that a player did not provide an evaluation, then, if all buyers are strategic, the only equilibrium is one in which sellers either provide no effort or do not transact. It highlights two forces at work. The first is the classic free-rider problem. If reputations do not track whether or not evaluations are provided, buyers have no incentives to provide them. Second, a community must share some information in order to achieve cooperative outcomes.

Before considering the ability of buyers to develop a reputation for information sharing, I consider alternative motivations for sharing information. In certain communities, there may be unmodelled benefits to providing information to the community. If certain buyers care about the well-being of fellow community members, they may desire to provide costly evaluations. The existence of altruistic behavior has been well documented in the experimental literature.[18] In an environment similar to this model without reputations for information sharing, Resnick and Zeckhauser (2002) report that buyers rate sellers in approximately 50% of all eBay transactions, a finding consistent with altruistic behavior.

With this in mind, I assume that with probability equal to $\gamma \in (0, 1]$, a buyer is a commitment type who always provides evaluations (an altruist).[19] The remaining buyers are strategic types whose payoffs are as previously specified. A player's type is hidden from his match. As the next result shows, the folk theorem still holds even when strategic buyers do not provide evaluations.

[18] See, for example, Eckel and Grossman (1996), Palfrey and Prisbey (1996) and Andreoni and Miller (2002).

[19] Under the assumption that the provision of evaluations does not affect buyer reputations, it does not matter whether γ refers to the percentage of players who always play altruistically or whether all buyers play "altruistically" with probability γ. The important assumption is that sellers believe that the probability of an evaluation following any particular transaction equals γ.

Proposition 3. *If there are buyers who always provide evaluations, there exists a straightforward and globally stable equilibrium under extended local information processing for all* $v^* \in V$ *if* $\delta \in (\delta^*, 1)$ *for some* δ^*.

Proof of Proposition 3. Assume for altruistic buyers who always leave evaluations, $k = 0$.[20] Let $\{a^*, P\}$ be such that strategic buyers never provide evaluations ($e_b = 0$) and $v = v^*$. Set each buyer's reputation equal to 0, and let each seller's reputation be determined by the following:

$$\phi_s(r, a) = \begin{cases} 0 & \text{if } r_{\hat{s}} = 0 \text{ and } (a_{\hat{s}} = a^* \text{ if } e_{\hat{b}} = 1) \\ 1 & \text{if } a_{\hat{s}} \neq \sigma_b(r) \text{ and } e_{\hat{b}} = 1 \\ (r_{\hat{s}} + 1) \bmod (X + 1) & \text{if } r_{\hat{s}} \neq 0 \text{ and not } (a_{\hat{s}} \neq \sigma_b(r) \text{ and } e_{\hat{b}} = 1). \end{cases}$$

Note that seller punishment only (re)starts if, when she deviates, she happens to be faced with an altruistic evaluator. The remainder of the proof follows the proof of Proposition 1. ■

Provided that at least some evaluations are provided altruistically, then the efficient outcome of always providing effort is an equilibrium outcome if players are sufficiently patient. The degree of patience required will depend on the proportion of altruistic evaluators. The fewer altruists, the more likely that a seller can get away with low effort without getting caught. This increase in the expected benefits of deviating from a cooperative equilibrium must be matched with an increase in the punishment conditional on getting caught. As the punishment takes the form of periods of holding the caught deviator to her minimax payoff, an increase in punishment takes the form of extending the punishment period. This threat has deterrence power only insofar as the seller sufficiently cares about these periods far in the future. The next result demonstrates this link between the percent of altruistic buyers who always share their information with the community and the rate at which players discount future earnings.

Proposition 4. *If*

$$\delta > \frac{c}{\gamma P + (1 - \gamma)c,} \tag{3}$$

then there exists a straightforward and globally stable equilibrium under extended local information processing in which sellers always provide high effort.

Proof. Proposed action profiles are $a_s = \{1, 1, 0\}$ for sellers, $a_b = \{1, 0\}$ for strategic buyers and $a_b = \{1, 1\}$ for commitment buyers. Let this profile be a^*. If sellers always provide effort, then by Equation (2), the best response for is for buyers to transact.

[20] If evaluations costs were negative for altruists, total payoffs would actually be greater than compared with the base case.

The following specification gives sellers the greatest incentive to provide high effort. Namely, players transact only if the seller has a good reputation:

$$\sigma(r) = \begin{cases} a^* & \text{if } r_{\hat{s}} = 0 \\ m & \text{if } r_{\hat{s}} = 1. \end{cases}$$

Set $r_{\hat{b}} = 0 \; \forall \; j \in N_\rho$. Seller reputations are:

$$\phi_s(r, a) = \begin{cases} 0 & \text{if } r_{\hat{s}} = 0 \text{ and } (a_s = \sigma_s(r) \text{ if } e_{\hat{b}} = 1) \\ 1 & \text{otherwise.} \end{cases}$$

As buyers are always playing a best response, I only need to check the incentive constraints for sellers. Values to following the strategy and a deviation are given as below:

$$\begin{aligned} v(a_{\hat{s}} = 1) &= (1-\delta)(P-c) + \delta(P-c) \\ v(a_{\hat{s}} = 0) &= (1-\delta)(P) + \gamma\delta 0 + (1-\gamma)\delta(p-c) \end{aligned}$$

Thus, players do not deviate from the equilibrium as long as $\delta > \frac{c}{P\gamma + (1-\gamma)c,}$. ■

This result highlights the need for rents from investing in a reputation. Looking at the case in which all players are altruists who provide evaluations ($\gamma = 1$), we see that in order for sellers to provide high effort, it must be the case that $\delta < \frac{c}{P}$. Thus, as the seller gets more of the surplus created by her effort, i.e., a larger P, less patience is required to support the efficient outcome. This result has been noted in Klein and Leffler (1981) and Shapiro (1983).

Altruists perform a public service. Given the fact that a seller has no way of knowing the buyer's type, she is not able to differentiate effort levels between them. Assuming that buyers are sufficiently patient, the possibility that a current match *might* provide his information to the community is sufficient to induce the seller to act in a trustworthy manner. Of course, the fewer altruists, the more tempting it is for a seller to defect from the high-effort equilibrium because it is less likely that she will get caught. It is thus not surprising that environments that depend on the voluntary provision of private information, such as eBay, encourage its provision.[21] From the Web site of eBay Inc. (2004):

> "Leave feedback after completing any purchase or sale on eBay. Your honest feedback shapes the eBay community and impacts the success and behavior of other eBay members."

Notwithstanding eBay's commercial success, sufficient altruism may be difficult to foster in on-line environments similar to eBay. Due to the sheer number of buyers, the probability that a particular provision of information will be of benefit to those a particular buyer

[21] eBay has also started to provide traders with information concerning a member's history in leaving feedback. I discuss eBay's new system in relation to the mechanism I propose in Section 5.

knows, much less cares about, will be low. Hoffman et al. (1996) find that social distance between "giver" and "recipient" or observer may be an important determinant of altruistic behavior.

When buyers are not sufficiently patient relative to the number of altruistic evaluators, other mechanisms provide incentives for private information provision. I consider a community or market that tracks whether or not an information holder provides his information to the community, and thus allows him to develop a reputation for information provision. It also allows sellers to distinguish between those who do and those who do not provide evaluations, and potentially treat them differently.

Both the buyer and the seller have an incentive problem. Just as the seller has a myopic incentive not to provide high effort after payment has been received, the buyer has a myopic incentive not to provide an evaluation after the good or service has been received. This creates an additional problem, as illustrated by the following mental exercise. Imagine a community that has been playing a two-sided reputation game for an arbitrarily large number of periods, with each player using an arbitrary strategy. After period N, a central planner gathers all the community members together in a room and proposes that starting with period $N+1$, everyone play a specified "equilibrium" strategy. The players, each of whom has had a different history of play and thus different beliefs about the distribution of reputations, go home and contemplate this strategy. The specified strategy is only an equilibrium if every player, regardless of history, and believing everyone else will play the strategy, desires to play the strategy at $N + 1$. Consider, however a player with a good reputation who believes that he will encounter many players with bad reputations in the future.[22] If punishment is costly to the player with a good reputation, the belief that he will need to punish many in the future may induce him to "defect" now. The future with a positive reputation is not very valuable to him, and he might well desire to not invest in a positive reputation today. I therefore define the following *repent* action a_ρ^Υ for each role.

Definition 2. A buyer's **repent action** is $a_b^\Upsilon = \{t_b = 1, e_b = 0\}$ while the seller plays $a_s^{m^*} = \{t_s = 1, h_s = 0, e_s = 1\}$ for some $P > k$. A seller's repent action is $a_s^\Upsilon = \{t_s = 1, h_s = 1, e_s = 0\}$ while the buyer plays $a_b^{m^*} = \{t_b = 1, e_b = 1\}$ for $P = 0$.

The repent action for an agent is one that, when her partner is playing a minimax strategy in the trading game, increases the partner's payoffs while not increasing the repenting agent's payoffs. In this game, the seller repents by providing high quality even though the buyer pays nothing. The buyer repents by paying $P > k$ with the knowledge that the seller with whom he is matched will not provide high effort.[23] In either event, the "payment" needs to be at least k so that the partner can provide the information to the community that the match did, indeed, repent, while at the same time receiving a payoff better than

[22] A player with a good or positive reputation is one who does not need to be "punished." In this context, a seller with a good reputation is one who provides high effort, while a buyer with a good reputation is one who provides evaluations.

[23] If the buyer does provide an evaluation, it does not affect the seller's reputation as the community knows that she is supposed to punish the buyer.

the minimax value. If the buyer deviates from a cooperative equilibrium when he has a good reputation, the signal will be the *lack* of an evaluation. However, when a buyer is "repenting," it needs to be the seller who provides the information to the community about whether or not he repented.

The equilibrium strategies have the following flavor. When two players with a positive reputation are matched, they play a^* achieving $v^* \in V^\dagger$. When two players with poor reputations are matched, they play the mutual minimax strategies. When a player with a poor reputation is matched with a player with a positive reputation, the latter minimaxes the former, while the player with a poor reputation plays the repent action. The next result demonstrates the range of payoffs available under local information processing when evaluation provision is costly and reputations are updated every period.

Proposition 5. *If reputations are updated every period ($T_\rho = 1,\ \rho = b, s$), every point $v \in V^\dagger$ is sustained by a straightforward and globally stable equilibrium with local information processing, if $\delta \in (\delta^*, 1)$ for some δ^*.*

Proof of Proposition 5. The proposition follows from Kandori (1992, Theorem 2), the modified proof of which I present here.

Let $T_b = T_s = 1$. Let a^* be the action profile, possibly in correlated strategies that achieves payoffs v^*. Strategies are given by:

$$\sigma(r) = \begin{cases} a^* & \text{if } r = (0,0) \\ (a_s^m, a_b^\curlyvee) & \text{if } r_{\hat{s}} = 0 \text{ and } r_{\hat{b}} \neq 0 \\ (a_s^\curlyvee, a_b^m) & \text{if } r_{\hat{s}} \neq 0 \text{ and } r_{\hat{b}} = 0 \\ a^m & \text{if } r_{\hat{s}}, r_{\hat{b}} \neq 0. \end{cases} \tag{4}$$

Reputations for sellers are given by

$$\phi_s(r, a) = \begin{cases} 0 & \text{if } r_{\hat{s}} = 0 \text{ and } a_{\hat{s}} = \sigma_s(r) \\ 1 & \text{if } a_{\hat{s}} \neq \sigma_s(r) \\ (r_{\hat{s}}+1) \bmod (X+1) & \text{if } r_{\hat{s}} \neq 0 \text{ and } a_{\hat{s}} = \sigma_s(r), \end{cases}$$

where X is the number of punishment periods. Buyer reputations are determined analogously.

I first check that there is some $(\delta^* < 1, X < \infty)$ such that a seller in the punishment phase follows the proposed strategies. If she complies she gets

$$x(1) + \delta x(2) + \cdots + \delta^{X-1} x(t) + (\delta^X + \delta^{X+1} + \cdots) v_s \tag{5}$$

where $x(\tau)$ is either $g(a^m) = 0$ or $g_s(a_s^\curlyvee, a_b^{m^*}) \leq 0$, depending on her match in period τ. If she defects from her punishment, the best she can do is

$$0 + \delta x(2) + \cdots + \delta^{X-1} x(\tau) + \delta^X g_s(a_s^\curlyvee, a_b^{m^*}) + (\delta^{X+1} + \delta^{X+2} + \cdots) v_s. \tag{6}$$

If she deviates, she will definitely meet a buyer with a positive reputation in period $(X+1)$, as all other players have followed equilibrium strategies. Subtracting equation (6) from (5) reduces to $x(1) - \delta^X g_s(a_s^{\curlyvee}, a_b^{m^*}) + \delta^X v_s$, which will be positive for δ^X close to one.[24]

I now check the incentives of a seller with a positive reputation to comply with the strategy. She will assuredly comply as long as

$$(1-\delta^X) \min \left\{ g_s(a_s^{\curlyvee}, a_b^{m^*}), v_s \right\} + \delta^X v_s \geq (1-\delta)\hat{v}_s + \delta \left((1-\delta^X) g_s(a^m) + \delta^X v_s \right), \qquad (7)$$

where the left-hand side is the minimum value of complying with σ, the right-hand side the best outcome if she deviates, and $\hat{v}_s$ is the seller's most profitable deviation. Increasing δ while holding δ^X constant, the right-hand side of equation (7) approaches $(1-\delta^X)g_s(a^m) + \delta^X v_s$, which is strictly less than the left-hand side.

The incentive constraints for the buyer are analogous. As all players have a positive reputation after at most X periods, the equilibrium is globally stable. ■

Proposition 5 shows that without the altruistic provision of evaluations, all individually rational stage game payoffs can be sustained in equilibrium. I now investigate what must be true in order to get players to provide their private information to the community. The value to buyers of an equilibrium in which they provide evaluations and sellers provide high effort is given by $(1-P-k)$. The harshest punishment is one in which sellers never trade with a buyer with a bad reputation for information sharing. The value of deviating is at least $(1-\delta)(1-P)+\delta 0$. Therefore, information provision is rational if and only if

$$\delta > \frac{k}{1-P}. \qquad (8)$$

Just as sufficient rents available to the seller are necessary in order for her to invest in a reputation for providing high effort, sufficient rents available to the information holder $(1-p)$ are also necessary in order for him to invest in a reputation for information sharing.

Combining Equation (3) given that all buyers provide evaluations ($\gamma = 1$) with Equation (8) shows that if

$$\delta > c + k,$$

then there exists a P such that the efficient outcome is obtainable under local information processing in a community that tracks information sharing reputations.

That the set of equilibrium payoffs is smaller when evaluations are costly is not surprising: someone needs to fund these costs out of trading-game profits. When buyer reputations are updated each period, as assumed, these costs need to be paid each period in which there is effort by the seller. However, as shown in Proposition 4 in the context of altruistically provided evaluations, evaluation provision in each period may not be necessary to provide

[24]To see this, note that the difference is at least as large as $(1-\delta^X)f_s(a_s^{\curlyvee}, a_b^{m^*}) + \delta^X v_s$.

incentives for high effort. If a seller believes that her probability of being caught is large enough relative to her discount factor, she will not chance punishment.

Taken together, these observations imply that communities may be able to reclaim much of the lost surplus by requiring that buyers provide an evaluation in some fraction of periods in order to maintain a reputation for information sharing. Consider a community that assigns a positive reputation for information sharing as long as a buyer provides an evaluation once every "cycle" of 7 periods. Due to discounting, if the buyer desires to maintain a positive reputation, he will wait until the final period of each cycle to provide an evaluation. That is, if his cycle ends on Tuesday, he will wait until Tuesday to provide an evaluation. Assume that on each day of the week, one out of every seven buyers starts a new cycle. As long as the seller does not know on which day of the week a particular buyer's cycle ends, she faces the same problem as she does when $\frac{1}{7}$ of all buyers are altruists.[25]

I now present the main result of this section. Under perfect private monitoring, even if evaluations are costly, all payoffs in the interior of the set of feasible and individually rational payoffs can be supported as long as the discount factor is high enough. In these equilibria, if a buyer has a positive reputation, his reputation is updated once every T_b periods. If he provided an evaluation at least once in the last T_b periods, he maintains a positive reputation. If not, he enters an X-period punishment phase: he plays the repent action if he meets a seller with a positive reputation, and minimaxes a seller with a negative reputation. Each seller's reputation cycle is one period long ($T_s = 1$): any time she is "caught" deviating, she enters the punishment phase.

First, I formally define $V^{\ddagger}$, the set of all feasible and individually rational payoffs except for the outer boundary of feasible trading game payoffs.

Definition 3. $V^{\ddagger} = \{v \in \text{co } f(A) | v \gg 0, v_b + v_s < 1 - c\}$ where co $f(A)$ is the convex hull of the set $f(A)$.

Theorem 1. *Let $V(\delta, N)$ be the set of equilibrium payoffs when the discount factor is δ and there are N players of each type. Under extended local information processing,*

$$\lim_{\delta \to 1, N \to \infty} V(\delta, N) = V^{\ddagger}.$$

Proof. By Proposition 5, every $v \in V^{\dagger}$ is sustained by a straightforward and globally stable equilibrium under extended local information processing. Select any $v^* \in V^{\ddagger} \setminus V^{\dagger}$. Letting $\Pr[e_b = 1] = \frac{1}{T_b} = \gamma$, consider, in particular the minimum positive integer T_b such that

$$\Pr[h_s t_b t_s = 1 | a^*] - \frac{1}{T_b} k = v_b^* + v_s^*,$$

[25] I thank Paul Resnick for pointing out an alternative version of the mechanism, which may be more feasible in an online environment. Rather than requiring the buyer to provide an evaluation once every 7 periods, this mechanism would, with probability $\frac{1}{7}$ after any transaction, inform the buyer that he needs to provide an evaluation this period in order to maintain a positive reputation. The outcomes of either mechanism are qualitatively similar.

with $\Pr[h_s t_b t_s = 1|a^*] \leq 1$. Given a total value of $v_b^* + v_s^*$, select P such that buyer receives v_b^* and seller receives v_s^*.

Let a seller's reputation be updated every period, and a buyer's every T_b periods. Assign each buyer and updating phase $x_{\hat{b}} \in \{1, 2, \ldots, T_b\}$ such that there exists at least one buyer with $x_{\hat{b}} = x'$ for all $x' \in \{1, 2, \ldots, T_b\}$. This requires that $N \geq T_b$. Therefore, for all τ, there is at least 1 buyer such that $(\tau + x_{\hat{b}}) \bmod T_b = 0$.

Strategies are as specified in equation (4). Seller reputations are given by

$$\phi_s(r, a) = \begin{cases} 0 & \text{if } r_{\hat{s}} = 0 \text{ and } (a_{\hat{s}} = \sigma_s(r) \text{ if } a_{\hat{b}} = 1) \\ 1 & \text{if } (a_{\hat{s}} \neq \sigma_s(r) \text{ if } a_{\hat{b}} = 1) \\ (r_{\hat{s}} + 1) \bmod (T+1) & \text{if } r_{\hat{s}} \neq 0 \text{ and } (a_{\hat{s}} = \sigma_s(r) \text{ if } a_{\hat{b}} = 1). \end{cases}$$

Buyer reputations are given by:

$$\phi_b(r, a, \tau) = \begin{cases} 0 & \text{if } r_{\hat{b}} = 0 \text{ and } (\text{if } (t + x_{\hat{b}}) \bmod N^* = 0, a_{\hat{b}} = \sigma_b(r)) \\ 1 & \text{if } (a_{\hat{b}} \neq \sigma_b(r) \text{ if } r_{\hat{b}} \neq 0) \text{ or } (\text{if } r_{\hat{b}} = 0, (t + x_{\hat{b}}) \bmod N^* = 0 \\ & \text{and } a_{\hat{b}} \neq \sigma_b(r)) \\ (r_{\hat{b}} + 1) \bmod (T+1) & \text{if } r_{\hat{b}} \neq 0 \text{ and } a_{\hat{b}} = \sigma_b(r). \end{cases}$$

Note first that in the punishment phase, there are always evaluations, so satisfaction of the incentive constraint follows from the proof of Proposition 5.

Note that the incentive constraint of a buyer with a positive reputation is unchanged as well. The incentive constraint of a seller with a positive reputation is more subtle. This is due to the fact that if she defects, she might not get caught. At the time of deciding whether or not to deviate, the reputations of her future partners are independent of whether or not she gets caught. Define $\zeta(\tau)$ as a buyer's assessment of the probability of being partnered with a buyer with a poor reputation at time τ. Further define $\Upsilon_\tau^X = \sum_{t=1}^{X} \delta^{t-1}(\zeta(t+\tau) \cdot g_s(a_s^{m^*}, a_b^\Upsilon) + (1 - \zeta(t+\tau)) \cdot v_s)$, the expected discounted payoffs over X periods conditional on having a positive reputation, starting τ periods in the future.

The seller's constraint becomes

$$(1-\delta^X)\Upsilon_0^X + \delta^X v_s \geq (1-\delta)\hat{v}_s + \delta\left((1-\delta^X)(\gamma g_s(a^m) + (1-\gamma)\Upsilon_1^X) + \delta^X v_s\right). \tag{9}$$

Holding δ^T constant, as δ increases, the right-hand side approaches $(1-\delta^T)(\gamma g_s(m) + (1-\gamma)\Upsilon_1^T) + \delta^T v_s$, which is strictly less than the right-hand size as $g_s(m)$ is less than either $g_s(a_s^m, a_b^\Upsilon)$ or v_s .

As the punishment phase ends in X periods, the equilibrium is globally stable. ■

Remark: To appreciate the need for a sufficient number of players ($N \geq N^*$) consider the following. Assume there are 10 buyers, and a community decides that maintaining a positive reputation requires one evaluation every 100 periods. A seller knows that no buyer is providing an evaluation in at least 90 out of every 100 periods. She may decide not to provide high effort for if she is caught, she has learned information about the period of the cycle in which detection is more likely. If there is a positive probability of getting caught in each period (i.e., there are at least as many players as there are periods in a cycle), then by Proposition 3, there is threshold δ^* such that if the discount rate is above δ^*, the seller does not deviate even if she knows which period has the lowest (non-zero) probability of being detected. If there are players who provide evaluations altruistically in each period, then there is no longer a need for a sufficient number of buyers.

Theorem 1 shows that as players become increasingly patient, a community can require less frequent evaluations—they can increase the buyer's reputation cycle T_b. Total evaluation costs, even for a large cost of providing evaluations, tend toward 0 as T_b increases, and therefore, in the limit, the set of payoffs available to the community are arbitrarily close to those available if evaluations were costless. This result is similar in spirit to the main result presented by Ben-Porath and Kahneman (2003). In both their work and this study, evaluations are provided only periodically. They, however, consider games of private monitoring played repeatedly by the same group of players. Their mechanism induces truthful announcements, but relies on a more complicated mechanism with additional information assumptions that require both the evaluator and the evaluatee to announce the evaluatee's action.

4 Games with Imperfect Private Monitoring

In games with imperfect monitoring, the signal a player receives does not perfectly inform him about the actions of other players in the game. In the trading game, this means that either outcome is possible regardless of seller effort level. While a good outcome for the buyer is more likely if the seller exerts high effort, it is also possible that the seller exerted low effort.[26] As a result, even if all outcomes are publicly revealed, a seller always providing high effort cannot be an equilibrium.

The trading game is a *moral hazard mixing game*. Fudenberg and Levine (1994) show that in these games, maximal long-run payoffs are strictly less than those possible if actions rather than outcomes were observable. This result holds even if all outcomes are publicly observable. Bad outcomes are inevitable, and buyers must punish bad outcomes in order to provide incentives for the seller to exert effort. The seller strategy "always provide high effort" is not credible because if the buyers believed it, the seller would not provide high effort and blame the fates for poor outcomes.

Thus, games of imperfect monitoring have the feature of putting punishments on the equilibrium path. This contrasts with the game under the assumption of perfect monitoring:

[26]Weakening the assumption that evaluations are truthful (Assumption 2) in games of perfect monitoring yields a game analogous to the game with imperfect monitoring. In either case, there is a positive probability that a seller exerts high effort but the evaluation is negative.

punishment is off the equilibrium path, and therefore increasing the severity of the punishment does not alter the equilibrium "value" of the game. When punishments are on the equilibrium path, the severity of any punishments affect equilibrium payoffs and efficiency. In this section, I study the relationship between the amount of information elicited (i.e., the evaluation frequency) and required seller punishment, and thus overall efficiency.

As "moral hazard mixing game" suggests, players must randomize in order to maintain incentives for the seller to provide high effort with positive probability. Consider the trading game in which all outcomes are public. The seller's profits must be higher after a good outcome than after a bad outcome, otherwise she would never desire to provide costly high effort.

Assume first that prices are fixed. In order for expected profits to be greater with a better reputation, the probability of a transaction must be higher when the seller has a good reputation. Buyers can accomplish this by agreeing to transact with a higher probability when a seller has a good reputation. A buyer must therefore be indifferent between transacting and not. To make buyers indifferent, sellers must select the probability of high effort that leaves the buyer indifferent to transacting given the price. This has two implications. First, the probability of high effort is independent of the seller's reputation. A reputation is not an indicator of outcomes, but rather a coordination device that indicates to the indifferent buyer the probability with which he ought to transact. Second, transactions do not occur in all periods, complicating the reputation updating process.

If prices change according to the reputations of the agents in a pair, the pricing mechanism may work to increase good-reputation profits. I therefore assume that the price a buyer pays, if both parties agree to transact is equal to a fixed proportion, p, of the expected value of the item to the buyer.

Assumption 3. *The price in the trading game is* $P = p \cdot \tilde{z}_{\hat{s}\hat{b}}$, *where* $\tilde{z}_{\hat{s}\hat{b}} = \mu + Prob[h_b = 1 | r_{\hat{s}}, r_{\hat{b}}](\lambda - \mu)$ *(i.e., the expected outcome given strategies induced by given seller and buyer reputations) and* $p \in [0, 1]$ *is exogenously determined.*

This assumption leads to a couple of desirable characteristics in the equilibria described below. First, transactions rationally occur within each match, simplifying reputation transitions and allowing focus on the incentives to develop reputations. Second, reputations will be meaningful in the sense that a better reputation will be associated with an increased probability of high effort.

As a first step, I compare achievable outcomes under imperfect *public* monitoring with those available under perfect monitoring. I then derive the conditions under which, if we assume evaluations are provided with a specified probability, extended local information processing delivers these outcomes in games of imperfect *private* monitoring. Finally, I return to the question of buyer incentives, and derive necessary and sufficient conditions for buyers to provide evaluations in equilibrium under extended local information processing. The next lemma specifies maximal seller payoffs under public monitoring.

Lemma 1. *Under public monitoring, the maximal sequential equilibrium payoff for the seller is*

$$v = \begin{cases} \mu p & \textit{if } \frac{p}{c} < \frac{1-\mu}{(\lambda-\mu)^2} \\ \lambda p - c - \frac{c(1-\lambda)}{\lambda-\mu} & \textit{if } \frac{p}{c} \geq \frac{1-\mu}{(\lambda-\mu)^2}. \end{cases}$$

Proof. I use the *maximal score method* presented by Fudenberg and Levine (1994) in order to determine the greatest perfect public equilibrium payoffs available to the seller in the trading game under the assumption of perfect public monitoring. First note that the game has a *product structure*. That is, the probability distribution of the signals that seller $\hat{s}$ sends to her buyer (i.e., buyer $\nu(\hat{s},\tau)$'s outcome) is independent of the actions of other sellers. Denote public signals by $z_1 = (z_{\hat{b}} = 1)$ and $z_0 = (z_{\hat{b}} = 0)$, and the outcome probability matrix $\Pi_{\hat{s}}(\alpha_{\hat{b}})$ where rows correspond to each strategy $a_{\hat{b}}$ and columns to $z_{\hat{b}}$, where an entry corresponds to $Pr[z_{\hat{b}}|a_{\hat{s}}]$. A game with a product structure has *sufficient rank* at α if every left null vector of $\Pi_{\hat{s}}(\alpha_{\hat{b}})$ is orthogonal to the vector $f_{\hat{s}}(\alpha_{-i})$. For $\alpha_{\hat{b}} \in BR(\alpha_s)$, the outcome probability matrix is

$$\Pi_{\hat{s}}(\alpha_{\hat{b}}) = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \lambda & 1-\lambda \\ 0 & \mu & 1-\mu \end{pmatrix},$$

where the columns correspond to no trade, positive outcomes and negative outcomes respectively. The first two rows correspond the case in which the seller does not transact. As $\lambda \neq \mu$ by assumption, the left null vector of $\Pi_{\hat{s}}(\alpha_{\hat{b}})$ at $\alpha_{\hat{b}} \in BR(\alpha_s)$ is 0, and is thus orthogonal to the vector $f_{\hat{s}}(\alpha_{-i})$.

Given satisfaction of the full-rank condition Fudenberg and Levine (1994, Theorem 5.2) show that the maximal seller sequential equilibrium payoff v is the solution to the following linear programming problem:

$$\begin{aligned}
&\max_{v_{\hat{s}}, w_{\hat{s}}} v_{\hat{s}} \text{ subject to} \\
v_{\hat{s}} &= (1-\delta) f_s(a_{\hat{s}}, \alpha_{\hat{b}}) + \delta \left(\Pr[z_1|a_{\hat{s}}] w_{\hat{s}}(z_1) + (1 - \Pr[z_1|a_{\hat{s}}]) w_{\hat{s}}(z_0)\right) \\
&\text{for } a_{\hat{s}} \text{ such that } \alpha(a_{\hat{s}}) > 0 \\
v_{\hat{s}} &\geq (1-\delta) f_s(a_{\hat{s}}, \alpha_{\hat{b}}) + \delta \left(\Pr[z_1|a_{\hat{s}}] w_{\hat{s}}(z_1) + (1 - \Pr[z_1|a_{\hat{s}}]) w_{\hat{s}}(z_0)\right) \\
&\text{for } a_{\hat{s}} \text{ such that } \alpha(a_{\hat{s}}) = 0 \\
v_{\hat{s}} &\geq w(z) \text{ for } z = \{z_0, z_1\}.
\end{aligned}$$

When $\frac{(1-\mu)c}{(\mu-\lambda)^2} \geq \frac{p}{c}$, the solution to the linear programming problem is:

$$\begin{aligned}
v_{\hat{s}} &= p * \mu \\
h_s &= 0 \\
w_{\hat{s}}(\cdot) &= v_{\hat{s}}.
\end{aligned}$$

When $\frac{(1-\mu)c}{(\mu-\lambda)^2} < \frac{p}{c}$, the solution to the linear programming problem is:

$$\begin{aligned} v_{\hat{s}} &= \lambda p - \frac{c(1-\mu)}{\lambda-\mu} \\ h_s &= 1 \\ w(z_1) &= v_{\hat{s}} \\ w(z_0) &= v_{\hat{s}} - \frac{c(1-\delta)}{\delta\gamma(\lambda-\mu)}. \end{aligned}$$

■

A seller will desire to provide high effort when its value is sufficiently large. Lemma 1 identifies the two major determinants of the value of high effort. First, the value of high effort, and thus the incentives to provide it, are increasing in the direct returns to effort ($\frac{p}{c}$). Second, as high effort becomes more highly correlated with good outcomes relative to low effort ($(\lambda-\mu)^2$ increasing), effort becomes more desirable. When the value of effort is not sufficiently large, the best the seller can do is never provide effort and receive μp every period. When the seller does desire to provide high effort, her payoffs are lower than would be the case under perfect monitoring. The final term of her maximal value represents this loss due to imperfect monitoring.

I now turn back to the game with private monitoring. As outcomes signal effort level, it is natural to let seller reputations depend on outcomes as well as any observable actions. I use seller reputations based on those studied by Dellarocas (2004). A seller's reputation is positive if her previous evaluation was positive, and negative if it was negative.[27] Seller reputations are given by:

$$\phi_s(r,a,z) = \begin{cases} 0 & \text{if } (r_{\hat{s}}=0 \text{ and } e_{\hat{b}}=0) \text{ or } (z_{\hat{b}}\!=\!1 \text{ and } e_{\hat{b}}=1) \\ 1 & \text{if } (r_{\hat{s}}=1 \text{ and } e_{\hat{b}}=0) \text{ or } (z_{\hat{b}}\!=\!0 \text{ and } e_{\hat{b}}=1). \end{cases} \tag{10}$$

I first look at the case in which the community does not track buyer reputations. I assume that an evaluation is provided altruistically with probability γ. Due to the fact that buyers cannot develop reputations for information sharing, there is no incentive for strategic buyers to provide evaluations. Therefore, sellers assume that the probability of an evaluation after any transaction is γ. As the next proposition shows, maximal seller payoffs are a globally stable equilibrium under local information processing for a wide range of parameters.

Proposition 6. *Assume evaluations are provided with probability γ in every stage game.*

1. *If $\frac{p}{c} < \frac{1-\mu}{(\lambda-\mu)^2}$, then the maximal seller payoff $v=\mu p$ is obtainable as a globally stable equilibrium under extended local information processing with $a^*_{\hat{s}} = \{t_{\hat{s}}\!=\!1, h_{\hat{s}}\!=\!0, e_{\hat{s}}\!=\!0\}$ and $\sigma_s(r)\!=\!a^*_{\hat{s}}$ for all r.*

[27] In his base model, Dellarocas (2004) assumed that all buyers provide evaluations, and thus did not need to specify a seller reputation when a seller is not evaluated.

2. *If $\frac{p}{c} \geq \frac{1}{\delta(\lambda-\mu)^2}$, then the maximal seller payoff*

$$v_{\hat{s}} = p\lambda - c - \frac{c(1-\lambda)}{\lambda-\mu}$$

is obtainable as a globally stable equilibrium under extended local information processing for $\gamma \in [\frac{c}{p\delta(\lambda-\mu)^2}, 1]$. The seller strategy that achieves this payoff is

$$\sigma_s(r) = \begin{cases} \{t_s, h_s, e_s\} = \{1, 1, 0\} & \text{if } r_{\hat{s}} = 0 \\ \{t_s, h_s, e_s\} = \{1, 1 - \frac{c}{\gamma p(\lambda-\mu)^2\delta}, 0\} & \text{if } r_{\hat{s}} = 1. \end{cases}$$

Proof. Let $\pi_{rr'}(a_{\hat{b}})$ be the probability that a seller with reputation r has reputation r' next period conditional on selecting action $a_{\hat{b}}$. Under the assumption that a seller maintains the same reputation if the buyer does not provide an evaluation, these transition probabilities are given by:

$$\begin{aligned} \pi_{00}(h_{\hat{s}}=1) &= 1-\gamma(1-\lambda) \\ \pi_{00}(h_{\hat{s}}=0) &= 1-\gamma(1-\mu) \\ \pi_{10}(h_{\hat{s}}=1) &= \gamma\lambda \\ \pi_{10}(h_{\hat{s}}=0) &= \gamma\mu, \end{aligned} \tag{11}$$

where the probability of transitioning to a negative reputation ($r_{\hat{s}} = 1$) is one minus the probability of transitioning to a positive one.

I look for an equilibrium strategy a_s^* such that a seller always selects high effort when she has a good reputation ($r_{\hat{s}}=0$). The value of this strategy to a seller with a positive reputation is:

$$\begin{aligned} v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) = {} & (1-\delta)(p\lambda - c) + \\ & + \delta\left(\pi_{00}(h_{\hat{s}}=1)\cdot v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) + (1-\pi_{00}(h_{\hat{s}}=1))\cdot v_{\hat{s}}(r_{\hat{s}}=1, a_s^*)\right) \end{aligned} \tag{12}$$

Rearranging, in order for the seller to exert high effort, it must be the case that:

$$v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) - v_{\hat{s}}(r_{\hat{s}}=1, a_s^*) = \frac{c(1-\delta)}{(\lambda-\mu)\delta\gamma}. \tag{13}$$

Let $\tilde{h}_1$ be the probability that a seller provides high effort when she has a negative reputation. The value to her of a_s^* when she has a poor reputation is

$$\begin{aligned} v_{\hat{s}}(r_{\hat{s}}=1, a_s^*) = {} & (1-\delta)(p(\tilde{h}_1\lambda + (1-\tilde{h}_1)\mu) - c) + \\ & + \delta\left(\pi_{10}(h_{\hat{s}}=1)\cdot v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) + (1-\pi_{10}(h_{\hat{s}}=1))\cdot v_{\hat{s}}(r_{\hat{s}}=1, a_s^*)\right). \end{aligned} \tag{14}$$

Plugging equations (11), (12), and (13) into equation (14), I solve for the probability of high effort conditional on the seller having a poor reputation

$$\tilde{h}_1 = 1 - \frac{c}{(\mu-\lambda)^2\delta\gamma p}.$$

Letting $\gamma=1$ reveals that as long as

$$\frac{p}{c} \geq \frac{1}{\delta(\lambda-\mu)^2},$$

there exists a $\gamma^* \in (0,1]$ of high effort given poor seller reputation $(\tilde{h}_1)$ is non-negative.

In this case, the equilibrium is as follows

$$\sigma_s(r) = \begin{cases} (t_{\hat{s}}=1, h_{\hat{s}}=1, e_{\hat{s}}=0) & \text{if } r_{\hat{s}}=0 \\ (t_{\hat{s}}=1, h_{\hat{s}}=1-\frac{c}{(\mu-\lambda)^2\delta\gamma p}, e_{\hat{s}}=0) & \text{if } r_{\hat{s}}=1 \end{cases}$$

$$v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) = \lambda p - \frac{c(1-\mu)}{\lambda-\mu} \tag{15}$$

$$v_{\hat{s}}(r_{\hat{s}}=1, a_s^*) = v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) - \frac{c(1-\delta)}{\delta\gamma(\lambda-\mu)}.$$

For $p \in [0,1]$, buyers receive non-negative payoffs from the trading game at p, and therefore have no incentive to deviate. I now show that the seller is indifferent between high and low effort, and therefore does not desire to deviate from the strategy profile. If the seller is indifferent between effort levels, then

$$\begin{aligned} -c + \delta\left(\pi_{r0}(h_{\hat{s}}=1)\cdot v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) + (1-\pi_{r0}(h_{\hat{s}}=1))\cdot v_{\hat{s}}(r_{\hat{s}}=1, a_s^*)\right) = \\ = \delta\left(\pi_{r0}(h_{\hat{s}}=0)\cdot v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) + (1-\pi_{r0}(h_{\hat{s}}=0))\cdot v_{\hat{s}}(r_{\hat{s}}=1, a_s^*)\right), \end{aligned} \tag{16}$$

where the left-hand side captures the payoffs from providing high effort, the left from low, and $\pi_{r0}(h_{\hat{s}})$ the probability of having a positive reputation next period $(r=0)$ given effort $h_{\hat{s}}$ and current reputation r. Equation (16) reduces to

$$v_{\hat{s}}(r_{\hat{s}}=0, a_s^*) - v_{\hat{s}}(r_{\hat{s}}=1, a_s^*) = \frac{c(1-\delta)}{(\pi_{r0}(h_{\hat{s}}=1)-\pi_{r0}(h_{\hat{s}}=0))\delta\gamma}. \tag{17}$$

As $(\pi_{r0}(h_{\hat{s}}=1)-\pi_{r0}(h_{\hat{s}}=0)=(\lambda-\mu)$ regardless of current reputation $(r_{\hat{s}}=\{0,1\})$, equation (17) must hold as equation (13) holds.

In the case in which $\frac{p}{c} < \frac{1-\mu}{(\lambda-\mu)^2}$, never providing effort results in

$$\sigma_s(r) = (t_{\hat{s}}=1, h_{\hat{s}}=0)\ \forall\ r$$

$$v_{\hat{s}}(r, a_s^*) = \mu p.$$

For $p \in [0,1]$, buyers receive non-negative payoffs from the trading game, and therefore have no incentive to deviate. Likewise, the seller receives non-negative payoffs as well, and therefore transacting dominates not transacting. As both current and continuation payoffs are independent of whether the seller provides effort, providing costly high effort is dominated by not providing costly high effort.

As sellers play a^* regardless of history, the equilibrium is globally stable. ■

These results are similar to those presented by Dellarocas (2004) in his study of this seller reputation mechanism in an auction setting.[28] In the optimal high-effort equilibrium, a seller with a positive reputation always provides high effort. This promise is credible because she provides high effort with a lower probability when she has a negative reputation. Buyers pay less for a lower probability of high effort, which means that a seller's profits are less when she has a negative reputation.

Interestingly, assuming that the percentage of altruistic evaluators is sufficient to support the high effort equilibrium, the maximal seller equilibrium value does not depend on the actual percentage of altruists. However, even if altruistically provided evaluations are sufficient to support a high-effort equilibrium, I now show that the efficiency achieved in the trading game does depend on the frequency of evaluation provision.

Proposition 7. *If the maximal effort equilibrium is possible under extended local information processing, then expected equilibrium effort in the maximal effort equilibrium is equal to*

$$\tilde{h} = 1 - \frac{c(1-\lambda)}{(\lambda-\mu)(p(\lambda-\mu)\gamma\delta - c)}, \tag{18}$$

and therefore is increasing in γ*.*

Proof of Proposition 7. The probability that a seller has a positive reputation in part determines the probability of high effort on the equilibrium path. To calculate this probability, note that the seller's reputation progression is a Markov process as her reputation next period depends only her current period state and actions. Letting $\tilde{h}_1$ be the probability that a seller exerts high effort when she has a negative reputation, and recalling that sellers with a positive reputation always provide high effort, the Markov transition matrix is :

$$\mathbf{M} = \begin{pmatrix} 1-\gamma(\tilde{h}_1\lambda + (1-\tilde{h}_1)\mu) & \gamma(1-\lambda) \\ \gamma(\tilde{h}_1\lambda + (1-\tilde{h}_1)\mu) & 1-\gamma(1-\lambda) \end{pmatrix}.$$

The eigenvector associated with the eigenvalue equal to 1 is:

$$\vec{u} = \begin{pmatrix} \frac{1-\lambda}{\tilde{h}_1(\lambda-\mu)+\mu} \\ 1 \end{pmatrix}$$

The long-run probability distribution over states, and therefore the expected probabilities of positive and negative sellers are:

$$\begin{pmatrix} \Pr[r_{\hat{s}} = 1] \\ \Pr[r_{\hat{s}} = 0] \end{pmatrix} = \frac{1}{\vec{u}(1)+\vec{u}(2)} \cdot \vec{u} = \begin{pmatrix} \frac{1-\lambda}{1-(\lambda-\mu)(1-\tilde{h}_1)} \\ \frac{1-\lambda}{1-(\lambda-\mu)(1+\tilde{h}_1)} \end{pmatrix}. \tag{19}$$

[28] Dellarocas (2004) presents a more general seller reputation mechanism that presents, roughly, the number of negative evaluations over preceding T periods. In the optimal seller strategy, the probability of high effort decreases linearly in the number of negative evaluations. For a given environment, there will be an maximum T such that the equilibrium holds ($T^{\max}$), and the maximal seller equilibrium value is independent of T for all T less than $T^{\max}$.

Expected seller effort, $\tilde{h}$, is equal to $\Pr[r_{\hat{s}}=0]+(1-\Pr[r_{\hat{s}}=0])\tilde{h}_1$. After substituting for reputation probabilities with equation (19) and for $\tilde{h}_1$ with equation (15), this reduces to:

$$\tilde{h} = 1 - \frac{c(1-\lambda)}{(\lambda-\mu)(p(\lambda-\mu)\gamma\delta - c)}.$$

Differentiating with respect to γ

$$\frac{\partial\tilde{h}}{\partial\gamma} = \frac{c(1-\lambda)\gamma\delta}{(c-p(\lambda-\mu)\gamma\delta)^2} > 0.$$

■

Even though the maximal seller equilibrium payoffs are available under local information processing for any sufficiently large discount rate, the overall expected effort level, which measures efficiency, is increasing in the probability of an evaluation. To see why, recall that expected effort is equal to $\Pr[r_{\hat{s}}=0]+(1-\Pr[r_{\hat{s}}=0])\tilde{h}_1$. Note that by equation (19), evaluation probability has no *direct* effect on the long-run distribution of seller reputations ($\Pr[r_{\hat{s}}=0]$ and $\Pr[r_{\hat{s}}=1]$). This distribution does, however, depend on the the expected effort level from sellers with a poor reputation, which, by Proposition 6, is increasing in evaluation probability. Therefore, increasing the probability of evaluation increases *both* the probability of a seller having a positive reputation (and thus providing high effort with certainty) and the probability of high effort when the seller does find herself with a poor reputation.[29]

Therefore, even if altruistically provided evaluations are sufficient to support the high-effort equilibrium, it may be desirable to induce costly evaluation provision by strategic buyers. I now look at a community that allows buyers to develop a reputation for sharing information. I consider buyer reputations similar to those in the previous section: a buyer maintains a positive reputation as long as he provided an evaluation in at least one of his previous T_b transactions. Due to discounting, a strategic buyer will rate only in the final period of his cycle if he rates at all. Buyer reputations are given by:

$$\phi_b(a,r,\tau) = \begin{cases} 0 & \text{if } (r_{\hat{b}} = 0 \text{ and } (\tau + x_{\hat{b}}) \bmod T_b \neq 0) \\ & \text{or if } ((\tau + x_{\hat{b}}) \bmod T_b = 0 \text{ and } e_{\hat{b}} = 1) \\ 1 & \text{if } (r_{\hat{b}} = 1 \text{ and } (\tau + x_{\hat{b}}) \bmod T_b \neq 0) \\ & \text{or if } ((\tau + x_{\hat{b}}) \bmod T_b = 0 \text{ and } e_{\hat{b}} = 0). \end{cases} \tag{20}$$

The condition for the high effort equilibrium presented in Proposition 6 can only be satisfied with a positive probability of evaluation provision. In order for a high-effort equilibrium to

[29] Another way to see the link between the probability of evaluation and the value of a good reputation is to note that the value of a good reputation is equal to a reward today plus the discounted value of a convex combination of this good reputation value and the poor reputation value. Increasing the probability of evaluation decreases the value of this convex combination by increasing the probability of transitioning to a bad reputation. Therefore, in order to maintain the same good reputation value, the value of a poor reputation must increase. Increasing the probability of effort conditional on a poor reputation increases the value of this state.

be possible in the absence of altruists, sellers need to provide incentives to buyers to provide evaluations.

Whether these incentives are always sufficient to induce buyers to provide evaluations (i.e., even off the equilibrium path) depends on whether a buyer believes that he will be able to profit regardless of his previous history. Recall the previous discussion concerning player beliefs. In order for a given strategy profile to be an equilibrium, it must be rational for the buyer to provide evaluations even if he believes that *all* sellers currently have bad reputations (i.e., he has *pessimistic* beliefs). With this belief, even if everyone plays the equilibrium strategy from now on, the buyer believes that he is likely to encounter sellers who need to be punished. This belief lowers the expected returns to a good reputation.

I first solve for conditions which, starting from equilibrium beliefs, provide adequate incentives for buyers to provide evaluations. This will characterize conditions necessary for the existence of a high effort equilibrium in which incentives to provide evaluations to the community arise endogenously. I then return to the need to provide these incentives in face of pessimistic beliefs, and solve for sufficient conditions for high-effort equilibria with endogenous information sharing.

Maximal incentives for earning an information-sharing reputation are given when sellers provide no effort to buyers with a negative reputation. Let U_e be the expected payoffs when a buyer provides an evaluation in the final period of his reputation cycle, U_n his expected payoffs when he does not, and U^* the payoffs resulting from the optimal strategy. The returns to investing in a reputation for evaluation provision depend on the effort that a buyer expects from sellers given his positive reputation. Therefore, the distribution of seller reputations will be important. Letting $\tilde{h}_r$ represent the probability of high effort given reputation r, the overall probability of high effort, represented by $\tilde{h}$, is

$$\tilde{h} = \sum_{r_{\hat{s}}=0}^{1} \Pr[r_{\hat{s}}|\sigma_s(r)]\tilde{h}_{r_{\hat{s}}}.$$

Given that a buyer needs to evaluate once every T_b periods in order to maintain a positive reputation, in the last period of a buyer's reputation cycle, expected payoffs from providing an evaluation and not are:

$$U_e = -k + \delta\left((1-p)(\lambda\tilde{h} + (1-\tilde{h})\mu)(\sum_{\tau=1}^{T_b-1}\delta^{T_b-1}) + \delta^{T_b-1}U^*\right)$$

$$U_n = \delta\left((\mu(1-p)(\sum_{\tau=1}^{T_b-1}\delta^{T_b-1}) + \delta^{T_b-1}U^*\right).$$

The buyer provides an evaluation if and only if $U_e \geq U_n$, or

$$k \leq T(1-p)(\lambda - \mu)\tilde{h}\delta_b^T. \tag{21}$$

Expected seller effort for a buyer with a positive reputation is determined endogenously. As demonstrated in Proposition 6, seller effort will be a function of both seller reputation and the probability of evaluation. The next proposition details the necessary conditions for the existence of a high-effort equilibrium.

Proposition 8. *If a globally stable equilibrium in which sellers provide high effort exists, then it must be the case that*

$$k \leq \frac{(1-p)T_b\delta^{T_b}\left(cT_b(1-\mu) - p(\lambda-\mu)^2\delta\right)}{cT_b - p(\mu-\lambda)\delta}$$

for $N \geq T_b$.

Proof. I assume that evaluations are provided once every T_b periods, and show that in the maximal effort equilibrium, buyers provide evaluations once every T_b periods. To determine expected seller effort, $\tilde{h}$, substitute $\gamma = \frac{1}{T_b}$ into equation (18). Substituting this expression for $\tilde{h}$ into equation (21) reduces to the inequality in the proposition. ■

Proposition 8 highlights the two effects of decreasing evaluation frequency (increasing T_b) on the incentives for buyers to develop a positive reputation. The first-order effect is positive. Decreasing the evaluation frequency necessary to maintain a positive reputation increases the returns to a positive reputation as each evaluation buys more periods with high expected effort. This is partially mitigated by the second-order effect identified in Proposition 7: decreasing evaluation frequency lowers the effort that buyers with a positive reputation can expect from sellers.

Assuming the incentive compatibility condition in Proposition 8 is met, evaluation frequency has two overall efficiency effects. First, decreasing the frequency decreases the overall costs of information provision. However, decreasing evaluation frequency lowers trading game efficiency in games with imperfect monitoring. This contrasts with the case of perfect monitoring, in which evaluation frequency can be reduced to the incentive compatibility threshold without efficiency cost in the trading game.

The preceding analysis identifies conditions necessary for a high-effort equilibrium, and assumed "equilibrium" beliefs about seller effort. I now return to the issue of sufficient incentives to provide evaluations when a buyer believes that currently all sellers have negative reputations (i.e., pessimistic beliefs). I start by analyzing the case in which buyers must provide an evaluation each period in order to maintain a positive reputation for information sharing (i.e., $T_b = 1$). The next proposition provides conditions sufficient to ensure that a high effort equilibrium exists under extended local information processing.

Proposition 9. *If*

$$k \leq \frac{(1-p)\left(c - p(\lambda-\mu)^2\delta\right)}{p(\mu-\lambda)} \tag{22}$$

and $\frac{p}{c} \geq \frac{1}{\delta(\lambda-\mu)^2}$*, then the maximal seller payoff*

$$v_{\hat{s}} = p\lambda - c - \frac{c(1-\lambda)}{\lambda-\mu}$$

is obtainable as a globally stable equilibrium under extended local information processing.

Proof. Let the seller strategy be

$$\sigma_s(r) = \begin{cases} \{t_s, h_s, e_s\} = \{1, 1, 0\} & \text{if } r_{\hat{s}} = 0 \text{ and } r_{\hat{b}} = 0 \\ \{t_s, h_s, e_s\} = \{1, 1 - \frac{c}{p(\lambda - \mu)^2 \delta}, 0\} & \text{if } r_{\hat{s}} = 1 \text{ and } r_{\hat{b}} = 0 \\ \{t_s, h_s, e_s\} = \{1, 0, 0\} & \text{if } r_{\hat{b}} = 1, \end{cases}$$

and the buyer strategy be $\sigma_b = \{t_b, e_b\} = \{1, 1\}$. Letting $T_b = 1$ in equation (20), buyer reputations are given by:

$$\phi_b(r, a, z) = \begin{cases} 0 & \text{if } e_{\hat{b}} = 1 \\ 1 & \text{if } e_{\hat{b}} = 0. \end{cases}$$

In the proposed equilibrium, buyers provide an evaluation each period. To see that the seller strategy is an equilibrium strategy, note that buyer punishments last only one period. Therefore, in equilibrium, all buyers have positive reputations next period. Therefore, regardless of current beliefs, a seller believes that all future buyers will have a positive reputation. Proof that seller strategies are equilibrium strategies follows from the Proof of Proposition 6. Even if the current buyer has a negative reputation, U^* remains unchanged as all future buyers have positive reputations. Punishing the current buyer is a best response as the seller is indifferent between providing high and low effort.

I now show that buyers choose to always provide an evaluation. I show that even if a buyer believes his next partner will have a negative reputation, he still desires to provide an evaluation. Letting $\hat{h}$ be equilibrium effort level given $r_b = 0$ and $r_s = 1$, expected payoffs from providing an evaluation and not are:

$$U_e = -k + \delta \left((1-p)(\lambda \hat{h} + (1 - \hat{h}\mu) + U^* \right)$$
$$U_n = \delta \left((\mu(1-p) + U^* \right).$$

Substituting $1 - \frac{c}{p(\lambda - \mu)^2 \delta}$ for $\hat{h}$ and solving for k reduces to the condition in the Proposition. ■

The results of Proposition 9 provide sufficient conditions for the maximal effort equilibrium to be available under local information processing and imperfect personal monitoring when information sharing is costly. Even if an agent believes that all potential trading partners currently have a negative reputation, the specified strategies are an equilibrium. The condition in equation (22) states that even if a buyer is certain that his next partner will have a negative reputation, the increase in surplus that he realizes as a result of the seller exerting effort with a positive probability is greater than the cost of an evaluation.

Even if equation (22) is not satisfied, the high-effort equilibrium may still be feasible. First, consider an electronic commerce environment in which "the market" can credibly commit to providing aggregate statistics on the distribution of seller reputations. In such a case, pessimistic buyer beliefs may be untenable. Second, note that the derivation of equation (22) assumed that a buyer needs to refresh his reputation each period. Decreasing the evaluation frequency (increasing T_b) required to maintain a reputation for information sharing increases the incentives for a buyer to develop a positive reputation. The analysis is complicated,[30] but note that the longer buyer reputations last, the more likely a buyer will encounter a seller with a positive reputation with his current reputation, even if he believes all his future matches *currently* have a negative reputation. This increases the spread between the surplus received when future sellers exert no effort and the surplus received when more and more sellers have positive reputation and exert high effort with a high probability.

5 Discussion

The ability to earn and profit from a good reputation can be an effective mechanism for eliciting costly effort when an agent has myopic incentives to not incur the cost. This fact has been long appreciated in the context of seller reputations for high effort. My main contribution in this paper is to show that reputations for *gossiping* can be effective in inducing the provision of information upon which a seller's reputation is based. This finding is particularly salient in electronic commerce environments, where the market designer can control the information available to agents and level of information sharing sufficient to earn the reputation for gossiping.

In January 2004, eBay changed the information that it provides about traders. It now enables a member to view all of the feedback that a trader has left for her partners. eBay's information structure differs from the mechanisms in this paper in two important ways. First, by listing only the feedback left, and not listing those instances when a trader did *not* provide feedback, a trader's partner can only make inferences about the frequency with which a trader has provided feedback.[31] A buyer who has left feedback ten times in a thousand opportunities is quite different than the buyer who has left feedback ten times in ten opportunities. Second, the content of each evaluation is listed. This opens the possibility that a buyer may strategically adjust the content of his evaluations. Always leaving positive feedback may signal an unwillingness to leave negative feedback, and tempt a seller to provide poor service. Likewise, always leaving negative feedback may signal unrealistic expectations about good quality, and thus the seller may provide low effort at she believes that she will not be rewarded for exerting high effort. A buyer may try to find the optimal mix of positive and negative evaluations in order to elicit maximal seller effort. Ely et al. (2002) and Ely and Välimäki (2003) consider such games in which an agent has an incentive to change his actions in order to strategically mold his reputation, and derive conditions under which the reputation system can no longer ensure trustworthy behavior

[30] As we increase T_b, a seller with pessimistic beliefs can no longer be assured that her next partner will have a positive reputation. Therefore, the specified strategies for providing high effort are no longer best responses on the path to equilibrium.

[31] Resnick and Zeckhauser (2002) show that traditionally, sellers leave feedback more frequently than buyers. A seller can thus get get a rough idea of how frequently a buyer leaves feedback by comparing the number of feedbacks the buyer has received with the number of feedbacks he has left.

as incentives to choose an action to make his history consistent with that of a trustworthy person overwhelm the incentives to choose a trustworthy action in the current period. It remains to be seen what effects eBay's information structure will have on incentives to provide evaluations in a trustworthy manner.

In the mechanisms I propose, the incentives for buyers to obtain a reputation for information sharing work just like those for seller reputations: there need to be sufficient returns to building this reputation. I have identified three fundamental conditions that permit sufficient returns. First, sellers must be able to distinguish between buyers who do and do not have reputations for information sharing. Second, sellers must be able to treat buyers with a positive reputation differently from those with a negative reputation. Finally, buyers must retain a sufficient portion of the value created by this differential treatment. In the perfect monitoring case, I assumed a fixed transaction price. In order for buyers to desire a positive reputation for information sharing, this price needs to be low enough to ensure sufficient buyer surplus in the trading game for a buyer with a positive reputation. In the imperfect monitoring case, I made the assumption that a buyer paid a fixed fraction of the expected value. The important implication of this assumption is that buyer surplus is increasing in expected effort, and therefore he has incentives to invest in a reputation that promises extra effort.

The results of this paper points out a difference between reputations for effort and information sharing. Sellers provide high effort in order to avoid lower payoffs, and the community desires high effort as its cost is less than its value. While effort is valuable in and of itself, information is not in this context. Rather, its value lies in the seller effort that it enables. The question for a community or market designer is: "How much costly information do we need to induce desired outcomes?"

In the case of perfect monitoring, the desired outcome of "high effort always" is feasible with sufficient information if players are patient enough. Reputations for information provision can induce sufficient information. The level of information that is sufficient, and thus the costs that must be borne to provide sufficient information, is decreasing as players become more patient. This contrasts with the case of imperfect monitoring, in which "high effort always" is not feasible. Furthermore, there is a trade-off. More information increases information costs, but enables a community to get closer to maximal efficiency in the trading game.

References

Andreoni, J., Miller, J., March 2002. Giving according to garp: An experimental test of the consistency of preferences for altruism. Econometrica 70 (2), 737–753.

Avery, C., Resnick, P., Zeckhauser, R., June 1999. The market for evaluations. American Economic Review 89 (3), 564–584.

Bagwell, K. (Ed.), 2001. The Economics of Advertising. No. 136 in International Library of Critical Writings in Economics. Edward Elgar Publising, Cheltenham, U.K.

Ben-Porath, E., Kahneman, M., August 2003. Communication in repeated games with costly monitoring. Games and Economic Behavior 44 (2), 227–50.

Bolton, G. E., Katok, E., Ockenfels, A., November 2004. How effective are electronic reputation mechanisms? an experimental investigation. Management Science 50 (11), 1587–1602.

Dellarocas, C., July 2004. Sanctioning reputation mechanisms in online trading environments with moral hazard. Tech. Rep. 4297-03, MIT Sloan, available at: `http://ccs.mit.edu/dell/site2004.pdf`.

eBay Inc., 2004. Feedback system frequently asked questions. Available at `http://pages.ebay.com/help/feedback/feedback_faqs.html` accessed August 20, 2004.

Eckel, C. C., Grossman, P. J., October 1996. Altruism in anonymous dictator games. Games and Economic Behavior 16 (2), 181–191.

Ely, J., Välimäki, J., August 2003. Bad reputation. Quarterly Journal of Economics 118 (3), 785–814.

Ely, J. C., Fudenberg, D., Levine, D. K., 2002. When is reputation bad? Tech. Rep. 1962, Harvard Institute Research Working Paper.

Fudenberg, D., Levine, D. K., February 1994. Efficiency and observability with long-run and short-run players. Journal of Economic Theory 62 (1), 103–35.

Fudenberg, D., Maskin, E., May 1986. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 54 (3), 533–554.

Hoffman, E., McCabe, K. A., Smith, V. L., June 1996. Social distance and other-regarding behavior in dictator games. The American Economic Review 86 (3), 653–660.

Kandori, M., January 1992. Social norms and community enforcement. Review of Economic Studies 59 (1), 63–80.

Kihlstrom, R. E., Riordan, M. H., June 1984. Advertising as a signal. Journal of Political Economy 3 (3), 427–450.

Klein, B., Leffler, K. B., August 1981. The role of market forces in assuring contractual performance. The Journal of Political Economy 89 (4), 615–641.

Klein, D. B., July 1992. Promise keeping in the great society: A model of credit information sharing. Economics and Politics 4 (2), 117–36.

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

Milgrom, P. R., North, D. C., Weingast, B. R., March 1990. The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs. Economics and Politics 2 (1), 1–23.

Miller, N., Resnick, P., Zeckhauser, R., September 2005. Eliciting honest feedback: The peer-prediction method. Management Science 51 (9).

Okuno-Fujiwara, M., Postlewaite, A., April 1995. Social norms and random matching games. Games and Economic Behavior 9 (1), 79–109.

Palfrey, T. R., Prisbey, J. E., September 1996. Altruism, reputation and noise in linear public goods experiments. Journal of Public Economics 61 (3), 409–27.

Resnick, P., Zeckhauser, R., 2002. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In: Baye, M. R. (Ed.), The Economics of the Internet and E-Commerce. Vol. 11 of Advances in Applied Microeconomics. Elsevier Science, pp. 127–157.

Shapiro, C., November 1983. Premiums for high quality products as returns to reputations. The Quarterly Journal of Economics 98 (4), 659–680.