HARVARD UNIVERSITY

JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Faculty Research Working Papers Series

Eliciting Honest Feedback in Electronic Markets

Nolan Miller, Paul Resnick and Richard Zeckhauser

August 2002

RWP02-039

The views expressed in the KSG Faculty Research Working Paper Series are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or Harvard University. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Eliciting Honest Feedback in Electronic Markets

Nolan Miller, Paul Resnick, and Richard Zeckhauser[1]

August 30, 2002[2]

[1]Miller, Kennedy School of Government, Harvard University; Resnick, School of Information, University of Michigan; Zeckhauser, Kennedy School of Government and Visiting Professor at Harvard Business School.
[2]We thank Chris Avery, John Pratt, Lones Smith, Ennio Stachetti, and Hal Varian for helpful coments.

Abstract

Recommender and reputation systems seek to inform potential customers by securing current consumers' feedback about products and sellers. This paper proposes a payment-based system to induce honest reporting of feedback. The system applies proper scoring rules to each buyer's report, looking to how well it predicts the report of a later buyer. Honest reporting proves to be a Nash Equilibrium. To balance the budget, the incentive payment to each buyer is charged to someone other than the one whose report that buyer is asked to predict. In addition, payment schemes can be scaled to induce appropriate effort by raters.

1 Introduction

Reputation building, a time honored business practice, reduces the inefficiencies created by asymmetric information. Using their theory scalpels, economists have recently laid bare the anatomy of this practice. Their models posit repeat participation in a marketplace, but usually assume few or no repeat interactions between the same partners. Some (perhaps imperfect) information about a player's type is revealed during an interaction. We call this information feedback, generalizing from the use of that term at eBay. The feedback is then spread to other participants, who can use it to determine with whom to transact, or what strategies to employ. For example, a buyer might adjust his reservation price based on the feedback about a potential seller. The equilibria (or lack thereof) of these models suggest conditions under which reputations convey information about who is trustworthy, and the extent to which reputations provide incentives for good behavior in transactions (i.e., deter moral hazard), and the extent to which they discourage low-quality or dishonest participants from joining the marketplace (i.e., deter adverse selection).

To illustrate, Kreps and Wilson (1982) consider why it might be beneficial for a player to take costly actions to develop a reputation as being of certain type, even though such reputation building cannot be an equilibrium of a finitely repeated game of perfect information (the so-called chain-store paradox of Selten (1978)). They show that with even a little bit of asymmetric uncertainty about the players' payoffs, costly posturing to develop a reputation may be part of equilibrium expectations. Shapiro (1982) characterizes the conditions under which deliberate cycles of quality (building up a reputation and then cashing in on it) will be avoided. Kandori (1992) shows that moral hazard can be deterred even if information is only distributed and processed locally – that is, if players see their partners' reputations, but do not have full information about the history of activity in the market. Most recently, Tadelis (2002) shows that if firms' names can be traded, then their ability to provide incentives becomes "ageless." Young firms build a reputation to enjoy during their lifetimes. Older firms maintain the reputation so that it will garner a higher price when sold.

Electronic markets present new challenges for reputation building. Transactions rarely involve face-to-face contact, payment normally is made before goods can be inspected, repeat transactions are unusual, and other aspects of reputation that might be held hostage – e.g., standing in the community – are rarely available. When information comes by wire, severe asymmetries must be expected; market failure threatens. Moreover, because electronic merchants serve a geographically diverse clientele, individual buyers may be forced to rely on information provided by strangers, which may be of questionable relevance and dubious validity.[1]

New mechanisms must be found to build effective reputations. Here, the internet holds some advantages. Once information about sellers is collected, it can be disseminated widely at nearly zero cost. And, extraordinary information processing and storage capabilities are immediately at hand to tally scores potentially based on thousands of transactions. Thus while electronic markets are particularly vulnerable to opportunistic behavior, they are also particularly well suited to harness the power of reputations.

Most conventional models of reputations treat the elicitation and distribution of feedback as exogenous to the models. However, elicitation and distribution of feedback raise important issues of externalities and strategic behavior. This paper proposes mechanisms that employ the capabilities of the Internet to provide monetary incentives for buyers to make appropriate efforts to evaluate sellers and to make their evaluations truthful. Once the availability of reliable feedback is assured, the economic theory of reputations can be brought to bear on the problem of inducing sellers to act responsibly.

The paper is organized as follows. Section 2 discusses how reputation systems currently operate in electronic markets, their challenges, and unrealized potential. Section 3 presents our main theoretical result, a mechanism that simultaneously elicits honest feedback and balances the budget. It then discusses the problem of securing appropriate effort from evaluators. Section 4 addresses alternative approaches and complications, and section 5 concludes.

2 Recommender and Reputation Systems in Electronic Markets

A recommender system collects, distributes, and aggregates product evaluations, to inform the choices of future consumers (Resnick and Varian, 1997). A reputation system does these things with evaluations of people's behavior, especially that of sellers. Though few of those who provide or use the ratings know each other, these systems help people decide whom to trust, they encourage trustworthy behavior and appropriate effort, and they deter participation by those who are unskilled or dishonest (Resnick et al, 2000).

[1] For the same reason, buyers may be reluctant to provide useful feedback about their experiences with a seller, since they are unlikely to know the potential beneficiaries of this information.

A number of websites have built formal reputation systems. The best known and most widely studied system is that of eBay. eBay is the largest person-to-person online auction site, with more than 4 million auctions open at a time. eBay does not warranty its auctions; it is only as a listing service, while the buyers and the sellers assume all the risks associated with transactions. There are fraudulent transactions to be sure. Nonetheless, the overall rate of successful transactions remains astonishingly high for a market as "ripe with the possibility of large-scale fraud and deceit" as eBay (Kollock, 1999).

eBay attributes its remarkable rate of successful transactions to its reputation system, the Feedback Forum. After a transaction is completed, the buyer and seller have the opportunity to rate each other as positive, neutral, or negative (1, 0, or -1). They can also leave comments (such as "Good transaction. Nice person to do business with! Would highly recommend," or typical shorthand "A++++ seller"). All participants have the running total of feedback points attached visibly to their screen names, which may be pseudonyms. Breakdowns of cumulative ratings into positives, neutrals, and negatives are available with a mouse click. The descriptive comments come up with a further click, though they are difficult to peruse. Yahoo! Auction, Amazon, and other auction sites feature reputation systems like eBay's, with variations such as a rating scale from 1-5, or using several measures (friendliness, prompt response, quality of product, etc.), or averaging rather than totaling feedback scores.

Recommender and reputation systems have spread far beyond the auction sites, where their application seems most natural. Bizrate.com rates registered retailers by asking consumers to complete a survey after each purchase. So-called "expert sites" (www.expertcentral.com, www.askme.com) provide Q&A forums in which experts provide answers for questions posted by clients in exchange for reputation points and comments. Product review sites (such as www.epinions.com) offer both product evaluations and rating services for product reviewers – the better the review, the more points the reviewer receives. iExchange.com tallies and displays reputations for stock market analysts based on the performance of their picks. Messages on the news and commentary site Slashdot (www.slashdot.com) can be sorted and filtered based on reader feedback, and users accumulate karma or reputation points based in part on the popularity of their posts; higher reputation scores give users successively greater privileges within the system, including the option of moderating other people's posts.

Although recommender and reputation systems could enable the Internet to achieve its potential as a grand electronic bazaar, many challenges must be met. The principal challenge addressed in this paper is to secure sufficient and honest feedback. For, if this can be accomplished, meaningful reputations become possible, and the theory of reputations can be brought to bear on the problem.

On or off the Internet, it is not always easy to get feedback from buyers. First, they may be reluctant to report evaluations or to do so honestly. They may withhold positive evaluations if a seller's capacity is limited, in which case the information is not a public good. For example, wise parents are reluctant to reveal the names of their favorite baby-sitters. Buyers who wish to be seen as "nice" may withhold negative evaluations. The evaluated party may influence the content. For example, sellers' threats of retaliation for negative feedback or explicit or implicit offers of rewards for positive feedback might lead buyers to submit reports that do not accurately reflect their perceptions. The extreme paucity of negative feedback on eBay – buyers give negative or neutral feedback in fewer than 1% of all transactions – suggests that people are hesitant to leave such evaluations (Resnick and Zeckhauser, 2001). Moreover, when an eBay buyer does give negative feedback, the seller gives negative feedback 34% of the time, suggesting that some retaliation may be occurring.[2]

Dishonest feedback can also come from "phantom" transactions undertaken solely to generate feedback (that will build the reputations of friends or stain the reputations of enemies). Dellarocas (2000) models this situation. If phantom feedback can be limited to a small percentage of the total, and if true feedback is expected to follow some known distribution, it is possible to undermine phantom feedback by dropping outlier reports from computations of reputations.

A second problem in eliciting feedback is that evaluations may be underprovided. Providing feedback about a completed transaction requires some effort, and to improve the evaluation takes more, yet the information only benefits other participants.[3] There may also be insufficient testing of

[2] It is a common practice, officially encouraged by eBay, for buyers to contact sellers to try to resolve disputes before giving negative feedback. Since having to resolve a problem is already a negative outcome, but one that goes unrecorded, this kind of leniency in reporting negative feedback means that it takes longer for individual sellers with frequent problems to be unmasked. Such leniency, however, may have some desirable properties. Dellarocas analyzes a model in which, so long as harshness in interpretation of feedback is appropriately tied with leniency in giving feedback, leniency has some advantages in deterring seller opportunism (Dellarocas 2001). The problem we are concerned with here is not systematic leniency, but the failure to report negative evaluations, whatever threshold is in use.

[3] Despite the rational incentive to free-ride, provision of feedback at eBay is quite common, occurring in more than 50% of transactions for a sample from 1999 (Resnick and Zeckhauser, 2001).

participants whose reputations are not well established. That is, participants may avoid interacting with a newcomer, or someone whose initial feedback is not positive. Avery, Resnick, and Zeckhauser (1999, hereafter ARZ) explore this problem in the context of product evaluations. In their model, items have some underlying quality, but individuals vary in the value they attach to high- and low-quality items, and they perceive the true quality with some error. Absent other incentives, too few people will try out an item. For example, a high-quality theater production may never be discovered if, by chance, its first reviews are negative. Such reviews will influence others not to attend, and evaluations may dwindle to zero before a fair reputation is established.

ARZ analyze mechanisms whereby the early evaluators, who give the most useful information, are paid to provide information, and later evaluators pay so as to balance the budget.[4] They conclude that any two of three desirable properties for such a mechanism can be achieved, but not all three; the three properties are voluntary participation, no price discrimination, and budget balance. Conceivably, a producer or seller who might profit from positive evaluations of an item could subsidize the cost of early evaluations. (This approach sacrifices budget balance among the evaluators.) For example, theater producers at small venues who were confident of the quality of their productions might pay a fee to get the major newspaper critics to review it, even if the fee would not influence whether the review was favorable. In the Internet context, a new high-quality seller might offer items at a reduced price in order to ensure that there would be transactions from which a reputation could develop. This would represent "dues paying" by newcomers as a means to establish a reputation. A real world analogue of this model is the free trial, where potential customers are given free use of a product for some time, in the hope that they will purchase it in the future themselves and will tell others about it.

A third problem in eliciting feedback may result from the introduction of provision incentives. If payments are not tied to the quality of the evaluations, buyers may submit uninformative evaluations. Performance-based payments might seem a solution, but they could induce herding or collusion by the raters.

What mechanism could motivate buyers to provide evaluations, to make them honestly, and with sufficient care? The remainder of this paper addresses those questions. Our proposed mech-

[4] There are interesting parallels with the two-armed bandit problem, where the optimal strategy always involves playing the inferior arm with some probability. In an infinite two-armed bandit game, with ever changing players but a kept historical record, efficiency could be achieved by randomizing appropriately, and paying "unlucky" individuals to play the inferior arm.

anism recognizes that some reward is necessary; it relies on monetary payments to raters. To focus on the honest elicitation problem, we begin with the case in which acquiring and reporting information is costless to buyers. Later, we show that our basic mechanism can be adapted to situations where buyers bear costs to acquire and report information.

Our mechanism for eliciting honest feedback is based on tying the payments to buyers to the informativeness of their evaluations. Since there is no gold standard for what constitutes a "good" seller and the center has no independent information, the payment for an evaluation needs to reflect the degree to which the evaluation agrees with the evaluations of others.

3 A Mechanism for Eliciting Honest Feedback

We consider a world where a number of buyers engage with a seller and then rate her (or her product) for quality.[5] We assume that the quality of a seller does not vary over time, and that all buyers attach the same value to a seller's quality. Quality is observed with some error. It could be represented, for example, as a binomial process, where the quality on any trial is good or bad. After receiving the product, each buyer sends a message to a common processing facility called the center. The center has four roles: First, it continually updates the aggregate assessment of each seller's quality. Second, it distributes these assessments to the large audience of web users. Third, it rewards or penalizes each buyer on the basis of his rating. Fourth, it plays the role of a bank, ensuring that the mechanism at least breaks even in the long run. The center has no independent information.[6] Thus, when rewarding individual raters it can only rely on the information provided by other raters.

The sellers (or products), which are already in the market, have an innate, predetermined quality level, and play no active role in our model.[7] The central question is whether we can develop a payment system that simultaneously produces honest ratings as a Nash equilibrium and breaks even. The answer, fortunately, is positive for a quite reasonable set of assumptions. Our mechanism bases payments to agents on a proper scoring rule, a payoff structure that induces decision makers

[5] For clarity, we refer to sellers as female and buyers as male.

[6] Given independent verifying power, a variant of the system outlined below would be easier to implement. It could simply pay buyers on how effectively they predicted the center's information. Utilizing information from other buyers as well as the center would increase the reliability of the mechanism, but would not affect the incentive to report honestly.

[7] Future work should allow sellers to respond to their developing reputations.

to reveal their true beliefs about the distribution of a random variable by rewarding them based on their announced distribution of the possible outcomes, and the actual realization of the random variable. Our mechanism capitalizes on the predictive power of one rater's information for the information of other raters.

Formulation

Following tradition in the economics of information, we refer to a seller's quality as her type. We refer to a buyer's perception of a seller's type as his signal.

Suppose there are a finite number of seller types, indexed by $t = 1, ..., T$. Let $p(t)$ be the commonly held prior probability assigned to the seller being type t.[8] Assume that $p(t) > 0$ for all t and $\sum_{t=1}^{T} p(t) = 1$.

Let I be the set of buyers, where $|I| \geq 3$. We allow for the possibility that I is countably infinite. Each buyer, judging from his own transaction, privately observes a signal of the seller's type. Conditional on the seller's type, buyers' signals are independent and identically distributed. Let S^i denote the random signal received by buyer i. Let $S = \{s_1, ..., s_M\}$ be the set of possible signals, and let $f(s_m|t) = \Pr\left(S^i = s_m|t\right)$, where $f(s_m|t) > 0$ for all s_m and t, and $\sum_{m=1}^{M} f(s_m|t) = 1$ for all t. Further, we assume that the conditional distribution of signals is different for different values of t. Let $s^i \in S$ denote a generic realization of S^i. Frequently, we will use s_m^i to denote the event $S^i = s_m$. We assume that buyers are risk neutral and seek to maximize expected wealth.

In the mechanism we propose, the center asks each buyer to announce his signal and then makes a transfer to the buyer that may depend on all buyers' announcements. Let $a^i \in S$ denote one such announcement, and $a = \left(a^1, ..., a^I\right)$ denote a vector of announcements, one by each buyer. Let $a_m^i \in S$ denote buyer i's announcement when his signal is s_m, and $\bar{a}^i = \left(a_1^i, ..., a_M^i\right) \in S^M$ denote buyer i's announcement strategy. Let $\bar{a} = \left(\bar{a}^1, ..., \bar{a}^I\right)$ denote a vector of announcement strategies. As is customary, let the superscript "$-i$" denote a vector leaving off buyer i's component.

Let $\tau_i(a)$ denote the transfer paid to buyer i when the buyers make announcements a, and let $\tau(a) = (\tau_1(a), ..., \tau_I(a))$ be the vector of transfers made to all agents. We are particularly interested in transfers that balance in the sense that $\sum_{i=1}^{I} \tau_i(a) = 0$.

An announcement strategy $\bar{a}^i$ is a best response to $\bar{a}^{-i}$ for player i if for each m:

$$E_{S^{-i}}\left[\tau_i\left(a_m^i, \bar{a}^{-i}\right)|s_m^i\right] \geq E_{S^{-i}}\left[\tau_i\left(\hat{a}^i, \bar{a}^{-i}|s_m^i\right)|s_m^i\right] \text{ for all } \hat{a}^i \in S. \quad (1)$$

[8] We briefly address the issue of non-common priors later.

That is, a strategy is a best response if, conditional on receiving signal s_m, the announcement specified by the strategy maximizes that buyer's expected transfer, where the expectation is taken with respect to the distribution of all other buyers' signals conditional on $S^i = s_m$. Given transfer scheme $\tau(a)$, a vector of announcement strategies $\bar{a}$ is a Nash Equilibrium of the reporting game if (1) holds for $i = 1, ..., I$, and a strict Nash Equilibrium if the inequality in (1) is strict for all $i = 1,..., I$.

Truthful revelation is a Nash Equilibrium of the reporting game if (1) holds for all i when $a^i_m = s_m$ for all i and all m, and a strict Nash Equilibrium if the inequality is strict. That is, if all the other players announce truthfully, truthful announcement is a strict best response. Since buyers do not receive any direct return from their announcement, if there were no transfers at all then any strategy vector, including truthful revelation, would be a Nash equilibrium. However, since players are indifferent between all strategies when there are no monetary transfers, these Nash equilibria are not strict.

The Mechanism

Our main result shows that truthful revelation is a strict Nash equilibrium. The analysis begins by noting that although S^i and S^j are conditionally independent, they are not necessarily independent. Because each signal is drawn from the same distribution with unknown parameter t, S^i and S^j are generally dependent.

Our results rely on a form of dependence, which we call stochastic relevance.

Definition: Random variable S^i is stochastically relevant for random variable S^j if and only if the distribution of S^j conditional on S^i is different for different realizations of S^i.

More technically, S^i is stochastically relevant for S^j if for any distinct realizations of S^i, call them s^i and $\hat{s}^i$, there exists at least one realization of S^j, call it $\hat{s}^j$, such that $\Pr\left(s^j|s^i\right) \neq \Pr\left(s^j|\hat{s}^i\right)$. Let $g\left(S^j|S^i\right)$ be the distribution of S^j conditional on S^i, and let $g\left(s^j|s^i\right)$ represent $\Pr\left(S^j = s^j|S^i = s^i\right)$.

Lemma 1: *For generic distributions* $f(s_m|t)$ and $p(t)$, S^i *is stochastically relevant for* S^j *for any two distinct players* i *and* j.[9]

[9] That is, the closure of the set of distributions for which S^i is not stochastically relevant for S^j has Lebesgue measure zero. See Mas-Colell, Whinston, and Green (1995, p. 595) for a discussion of generic conditions.

Proof of Lemma 1: We argue by contradiction. By Bayes Rule:

$$g\left(s^j|s^i\right) = \sum_{t=1}^{T} f\left(s^j|t\right) \frac{f\left(s^i|t\right)p\left(t\right)}{\Pr\left(s^i\right)},$$

where $\Pr\left(s^i\right) = \sum_{t=1}^{T} f\left(s^i|t\right)p\left(t\right)$.

If there exists $s^i$and $\hat{s}^i$ such that $g\left(s^j|s^i\right) = g\left(s^j|\hat{s}^i\right)$ for all s^j, then the following must hold for each s^j_m :

$$\begin{aligned}
g\left(s^j_m|s^i\right) - g\left(s^j_m|\hat{s}^i\right) &= 0 \\
\sum_{t=1}^{T} f\left(s^j_m|t\right) \frac{f\left(s^i|t\right)p\left(t\right)}{\Pr\left(s^i\right)} - \sum_{t=1}^{T} f\left(s^j_m|t\right) \frac{f\left(\hat{s}^i|t\right)p\left(t\right)}{\Pr\left(\hat{s}^i\right)} &= 0 \\
\sum_{t=1}^{T} f\left(s^j_m|t\right) \left(\frac{f\left(s^i|t\right)p\left(t\right)}{\Pr\left(s^i\right)} - \frac{f\left(\hat{s}^i|t\right)p\left(t\right)}{\Pr\left(\hat{s}^i\right)}\right) &= 0 \\
\sum_{t=1}^{T} f\left(s^j_m|t\right)\Delta\left(t\right) &= 0, \qquad (2)
\end{aligned}$$

where $\Delta\left(t\right) = \left(\frac{f\left(s^i|t\right)p(t)}{\Pr(s^i)} - \frac{f\left(\hat{s}^i|t\right)p(t)}{\Pr(\hat{s}^i)}\right)$. Since (2) is an inner product, the set of distributions $f\left(s_m|t\right)$ and $p\left(t\right)$ that satisfy it is closed and has Lebesgue measure zero.[10]■

Consider two buyers, call them i and j. If S^i is stochastically relevant for S^j, then buyer i's signal provides information about the distribution of buyer j's information. Thus, if it were known that buyer j will honestly report his signal, the problem of eliciting buyer i's information is reduced to eliciting his belief about the distribution of j's signal.

The elicitation of beliefs about the distribution of S^j from an agent who has observed S^i is exactly the problem considered in the statistical decision theory literature on strictly proper scoring rules.[11] Put simply, suppose agent i privately observes the realization of S^i, which is stochastically relevant for some publicly observable random variable S^j, and agent i is asked to reveal his private information. A scoring rule is a function $R\left(s^j|s^i\right)$ that, for each possible announcement of S^i, assigns a score to each possible realization of S^j. A convenient interpretation is that the scoring rule specifies the payment made to the (risk neutral) agent following each realization S^j. A scoring rule is strictly proper if the expected score is uniquely maximized at the true value of the parameter

[10]To show this, note that (2) is satisfied only if all $\Delta t = 0$, or, failing this, $f\left(s_m|t\right)$ satisfies a linear equation of the form $\sum_{t=1}^{K} x_t \Delta t = 0$. It is straightforward to show that these restrictions are only satisfied non-generically.

[11]See Cooke (1991) for an introduction to the use of scoring rules.

(i.e., if an agent paid according to the rule uniquely maximizes his expected wealth by truthfully revealing his private information).

There are a number of strictly proper scoring rules. The three best known are:[12]

1. Quadratic Scoring Rule: $R\left(s_n^j|s_m^i\right) = 2g\left(s_n^j|s_m^i\right) - \sum_{h=1}^{M} g\left(s_h^j|s_m^i\right)^2$.

2. Spherical Scoring Rule: $R\left(s_n^j|s_m^i\right) = \frac{g\left(s_n^i|s_m^j\right)}{\left(\sum_{h=1}^{M} g\left(s_h^j|s_m^i\right)^2\right)^{\frac{1}{2}}}$.

3. Logarithmic Scoring Rule: $R\left(s_n^j|s_m^i\right) = \ln g\left(s_n^j|s_m^i\right)$.

Although we could use any strictly proper scoring rule in our construction, we adopt the logarithmic scoring rule because of its intuitive appeal and notational simplicity. The proof that the logarithmic scoring rule is strictly proper follows immediately from Jensen's inequality.

Lemma 2: *Suppose that risk neutral agent i knows the realization of random variable S^i that is stochastically relevant for random variable S^j. If agent i is asked to announce a realization of S^i and is paid according to the logarithmic scoring rule, the expected payment is uniquely maximized by announcing the true realization of S^i.*

Proof of Lemma 2: Let $s^* \in S$ be the realization of S^i and let $a \in S$ be the realization agent i announces.

$$\begin{aligned}
\sum_{n=1}^{M} \ln g\left(s_n^j|a\right) g\left(s_n^j|s^*\right) - \sum_{n=1}^{M} \ln g\left(s_n^j|s^*\right) g\left(s_n^j|s^*\right) &= \sum_{n=1}^{M} \ln \frac{g\left(s_n^j|a\right)}{g\left(s_n^j|s^*\right)} g\left(s_n^j|s^*\right) \\
&< \ln\left(\sum_{n=1}^{M} \frac{g\left(s_n^j|a\right)}{g\left(s_n^j|s^*\right)} g\left(s_n^j|s^*\right)\right) = 0.
\end{aligned}$$

The inequality in the second line follows from Jensen's inequality, strict concavity of the natural logarithm, and stochastic relevance.[13] ■

Lemma 2 proves that log-likelihood transfers can be used to induce truthful revelation by agent i as long as his private information is stochastically relevant for some other publicly available signal.

[12] See Cooke, 1991, p. 139.

[13] Stochastic relevance is necessary only for uniqueness of the best response. Without stochastic relevance, the inequality in the last line of the derivation is weak, $\leq$. Hence truth-telling is a best response, but not necessarily the only best response.

However, in the case we consider, each buyer's signal is private information, and therefore we can only check players' announcements against other players' announcements, not their actual signals.

We now turn to showing how to construct self-financing transfers that make truthful reporting a Nash equilibrium. The essence of our mechanism can be illustrated using three agents. For example, if buyer 1's signal is stochastically relevant for buyer 2's signal, then paying buyer 1 based on the log-likelihood of the reported realization of buyer 2's signal induces buyer 1 to truthfully announce his type, assuming buyer 2 announces honestly. However, since player 2's announcement affects the transfer to buyer 1, buyer 2 cannot also pay the transfers to player 1. Instead, to balance the transfers, let buyer 3 pay to buyer 1 the transfer specified by the scoring rule. Such payments do not affect incentives, since the transfer buyer 3 pays to buyer 1 is independent of buyer 3's action. Thus for each player a second player's signal is used to determine incentive payments, and a third player makes the payments. We formalize this intuition in the following proposition.

Proposition 1: *For generic distributions $f(s_m|t)$ and $p(t)$, there exist balanced transfers under which truthful reporting is a strict Nash Equilibrium of the reporting game.*

Proof of Proposition 1: For each buyer, choose another buyer $r(i)$ whose announcement i will be asked to predict. Let

$$\tau_i^*\left(a^i, a^{r(i)}\right) = \ln g\left(a^{r(i)}|a^i\right). \tag{3}$$

Assume buyer $r(i)$ reports honestly: $a^{r(i)}(s_m) = s_m$ for all m. Since S^i is stochastically relevant for $S^{r(i)}$, and $r(i)$ reports honestly, S^i is stochastically relevant for $r(i)$'s report as well. Given that $S^i = s^*$, player i chooses $a^i \in S$ in order to maximize:

$$\sum_{n=1}^{M} \ln g\left(s_n^{r(i)}|a^i\right) g\left(s_n^{r(i)}|s^*\right). \tag{4}$$

By Lemma 2, (4) is uniquely maximized by announcing $a^i = s^*$, i.e. truthful announcement is a strict best response. Thus, given that player $r(i)$ announces truthfully, player i's best response is to announce truthfully as well.

To balance transfers, choose any agent or group of agents not including $r(i)$ and divide the payment to agent i among them.[14] For example, suppose each agent is chosen to make transfers

[14] We refer to τ_i^* as the payment to buyer i, even though, given the log scoring rule, the payment is always negative: $0 < g\left(s_m^j|s_n^i\right) < 1$, so $\ln g\left(s_m^j|s_n^i\right) < 0$.

to one other agent. Let $b(i)$ be the buyer to whom i pays transfers, and choose $b(i)$ such that $b(r(i)) \neq i$. Let the net transfer to buyer i be given by:

$$\tau_i(a) = \tau_i^*\left(a^i, a^{r(i)}\right) - \tau_{b(i)}^*\left(a^{b(i)}, a^{r(b(i))}\right). \tag{5}$$

These transfers induce truthful reporting and are balanced.■

Proposition 1 shows that it is generally possible to solve the problem of eliciting honest information from a number of buyers yet still meet a breakeven constraint. Each player's expected gain from honest reporting is 0, but in particular cases it may be negative. To assure ex-post voluntary participation, a bond can be collected in advance. Let $q = \min_{s_m, s_n \in S} (\ln g(s_m|s_n))$. If $-q$ is collected in advance from each player, each player will receive positive payments after the evaluations are reported.[15]

Although we illustrate using the log-likelihood scoring rule, any other strictly proper scoring rule could be employed in Proposition 1. For example, consider the quadratic scoring rule described above. Basing transfers on the quadratic scoring rule, i.e., $\tau_i^*\left(a^i, a^{r(i)}\right) = 2g\left(a^{r(i)}|a^i\right) - \sum_{h=1}^{M} g\left(s_h^{r(i)}|a^i\right)^2$, also induces truthful reporting, and these transfers can be balanced as in (5).[16] The logarithmic and quadratic scoring rules each have advantages.[17] The logarithmic scoring rule is attractive for its simplicity and intuitive appeal, and also for the fact that the payoff assigned to an outcome depends on the probability of that outcome only; the probabilities of outcomes that do not occur do not factor into the payoff. It has been shown that when there are more than 2 outcomes, the logarithmic scoring rule (up to a positive affine transformation) is the only strictly proper scoring rule with this property, which has been termed "relevance."[18]

Criticism of the logarithmic score has focused on how it deals with low-probability events. Small changes in assessments of small probabilities can translate into very large changes in the score, and in order for the logarithmic score to be strictly proper such changes must be properly assessed by the decision maker. Further, when the distribution in question involves low-probability events,

[15] Since $0 < g\left(s_m^j|s_n^i\right) < 1$, $\ln g\left(s_m^j|s_n^i\right) < 0$. However, adding q to the log of the likelihood ensures that all payments to agent i are positive.

[16] A proof that the quadratic score is strictly proper can be found in Selten (1998).

[17] The merits of the spherical scoring rule relative to the logarithmic scoring rule are similar to those of the quadratic.

[18] See Cooke (1991) for a statement of the result, Winkler (1969) for a discussion of relevance, and Shuford, Albert, and Massengil (1966) for the proof that the only relevant strictly proper scoring rules are based on the logarithmic score.

the necessary range of payments may become very large. Selten (1998) cites these argument in supporting the quadratic scoring rule as an alternative, although he was considering a very different problem. The quadratic score has the advantage of not needing to rely on threats of large penalties to induce truthful revelation. However, quadratic scores are more complex, and have the unappealing property that the payoff assigned to an outcome depends on the likelihood of all outcomes. For our purposes, either score works, as would the spherical score or any other strictly proper scoring rule. Which rule is appropriate depends on the application.

Although we have only addressed the finite-signal case, our methods easily generalize to the continuous-signal case. However, when signals are continuous, the logarithmic scoring rule can be especially problematic. For example, if signals are normally distributed on the real line, then the signal density becomes arbitrarily small as the signal approaches positive or negative infinity. Consequently, the logarithm of the density approaches negative infinity, and the range of payments needed to induce truthful revelation may become infinite as well. In cases such as these, in which the signal density is very small (or zero), the quadratic scoring rule is more appropriate.[19]

Sequential Interaction

In the argument above, a number of buyers simultaneously report their experiences with the seller. The mechanism adapts readily to situations where buyers interact sequentially.[20] For example, if an infinite sequence of buyers interacts with the seller, then incentives for buyer i can be provided by a log-likelihood payment based on the information reported by buyer $i+1$ (or any other later player for whom buyer i's signal is stochastically relevant). This payment can be funded by player $i+2$ (or any combination of later players not including $i+1$). In fact, when there is a choice about the timing of the game, sequential reporting is preferable, since it allows later buyers to make immediate use of the information provided by their predecessors.

Suppose the seller interacts with an infinite sequence of buyers, indexed by $i = 1, 2, ...$ Initially, the commonly held prior distribution for the seller's type is given by $p(t)$. Let $p_1\left(t|s^1\right)$ denote the

[19] In Johnson et al. (2002), it is shown that a lower truncation of the logarithmic score can be used to induce nearly truthful revelation.

[20] Hanson (2002) applies a scoring-rule based approach in a model in which a number of experts are sequentially asked their belief about the distribution of a random event, the realization of which is revealed after all experts have reported. In our model, the seller's type is never revealed, and therefore we must rely on other agents' reports to provide incentives.

posterior distribution when buyer 1 receives signal s^1. That is,

$$p_1\left(t|s^1\right) = \frac{f\left(s^1|t\right)p\left(t\right)}{\Pr\left(s^1\right)}, \tag{6}$$

where $\Pr\left(s^1\right) = \sum_{t=1}^{T} f\left(s^1|t\right)p\left(t\right)$. Buyer 1's belief about the probability that $S^2 = s^2$ is given by

$$g\left(s^2|s^1\right) = \sum_{t=1}^{T} \frac{f\left(s^2|t\right)f\left(s^1|t\right)p_1\left(t|s^1\right)}{\Pr\left(s^2|s^1\right)}, \tag{7}$$

where $\Pr\left(s^2|s^1\right) = \sum_{t=1}^{T} f\left(s^2|t\right)p_1\left(t|s^1\right)$.

Given the distribution specified in (7), buyer 1 can be induced to truthfully reveal s^1 using the scoring rule specified in Proposition 1. Buyer 1 is asked the distribution of buyer 2's announcements. Payments to buyer 1 can then be made by buyer 3 after buyer 2 announces.

By recursively updating beliefs about the seller using Bayes' rule (as in (6)), each player's announcement feeds into the information used by subsequent players. Let

$$p_i\left(t|s^1,...,s^i\right) = \frac{f\left(s^i|t\right)p_{i-1}\left(t|s^1,...,s^{i-1}\right)}{\Pr\left(s^i\right)}, \tag{8}$$

where $\Pr\left(s^i\right) = \sum_{t=1}^{T} f\left(s^i|t\right)p_{i-1}\left(t|s^1,...,s^{i-1}\right)$, and let

$$g\left(s^{i+1}|s^i\right) = \sum_{t=1}^{T} \frac{f\left(s^{i+1}|t\right)f\left(s^i|t\right)p_i\left(t|s^1,...,s^i\right)}{\Pr\left(s^{i+1}|s^i\right)}, \tag{9}$$

where $\Pr\left(s^{i+1}|s^i\right) = \sum_{t=1}^{T} f\left(s^{i+1}|t\right)p_i\left(t|s^1,...,s^i\right)$.

The sequential elicitation game proceeds as follows. Beginning with buyer 1, each buyer i observes s^i and is asked the distribution of the subsequent buyer's signal. Transfers constructed according to (3) using the conditional distribution specified in (9) elicit truthful announcement. This announcement then becomes common knowledge and is used to update beliefs about the seller according to (8). Incentives to buyer $i+1$ are then constructed using a scoring rule that incorporates these updated beliefs. Transfers to buyer i are paid by buyer $i+2$.[21]

When a finite string of buyers purchase sequentially from the seller, given the end point, the scheme of checking each buyer's announcement against that of a later buyer has the potential to unravel. Although the last buyer has no incentive to lie, neither has he any incentive to tell the truth, since there is no future signal upon which to base his reward. Given the unreliability of the final announcement, there is no way to effectively induce the previous buyer to report truthfully, and

[21]Note that buyer 1 does not pay anyone.

therefore no way to provide appropriate incentives to the buyer before that, and so on. However, this incentive problem can be solved if the center can control when buyers' announcements become public. By grouping some buyers together and treating the groups as if they report simultaneously, the center can create reporting "rings" that provide appropriate incentives for every buyer. For example, suppose $N = 10$; we will consider the last three players, 8, 9, and 10. Base payments to 8 on 9's announcement of his signal, payments to 9 on 10's announcement, and payments to 10 on 8's announcement. To balance the transfers, let 8 pay 9, 9 pay 10, and 10 pay 8. As long as the center can avoid revealing these three buyers' announcements until all three have announced, effective incentives can be provided, and the chain will not unravel.

Effective incentives must also be provided to earlier buyers. The last buyer before the ring – buyer 7 in this case – presents a particular challenge, especially when the budget must be balanced. One way to provide incentives to 7 would be to ask 7 to predict 8's announcement, and have 9 pay transfers to 7. However, then buyer 9 pays two different buyers but gets paid only once. A possible solution would be to have buyer 1 pay a lump sum to buyer 9. Another approach would be to have buyer 1, who currently does not pay anyone, pay transfers to 7. The problem with this is that buyer 1 may have an incentive to lie in order to affect the perceived distribution of later signals. However, the center could address this by not revealing 1's information, at least until after player 7 has announced. This has no effect on the incentives of players other than 1, since 1's announcement is not used to provide incentives to any other players. And, since 1's announcement is not revealed, it cannot affect other players' beliefs about likelihood of later players' signals. A final approach would be to construct additional rings. For example, the ten-player chain could be divided into rings consisting of buyers 1 through 4, 5 through 7, and 8-10. The information provided by each ring could then be announced after the entire ring has reported. Although information is released somewhat more slowly in this scheme, it also solves the incentive problem.

Costly Reporting

To this point we have assumed that reporting is costless to allow us to focus on the essence of the scoring-rule based mechanism. In this section, we consider two different problems associated with costly reporting. The first, which we call the fixed-cost problem, posits that the buyer must incur a cost in order to evaluate the product, which may include direct costs and the opportunity cost of being an early evaluator rather than waiting for better information from other evaluators before deciding whether it is safe to consume the product (as discussed in ARZ). If the expected benefit is less than this cost, buyers will not provide feedback. Hence the elicitation mechanism must pay buyers to allow them to recoup the fixed cost. Payments of this sort address the real-world problem that often only the most motivated buyers offer feedback, providing a biased impression of the seller. The second problem, which we call the effort-inducement problem, supposes that the buyer, by incurring a cost, can increase the precision of his information about the seller. Hence the effort-inducement problem recognizes that even though buyers may report feedback, they may put too little effort into securing reliable information. The center then faces the dual problems of eliciting the more complex information and inducing the buyer to obtain information with the proper level of precision.

As is apparent from the proof of Proposition 1, the truth-inducing incentives provided by log-likelihood payments (or any of the scoring rules mentioned above) are unaffected by either a positive rescaling of all payments or the addition of a constant to all payments. Thus, even in the realm of log-likelihood rules, there is significant leeway to adapt the transfers. We next argue that adding a constant to all payments can overcome the fixed-cost problem, while rescaling the payments can partially address the effort-inducement problem.

Suppose there is a fixed cost of evaluating and reporting given by $c > 0$. To induce truthful reporting, a truthful announcement must maximize the expected transfer (4) and offer an expected return of at least c. This can be accomplished simply by adding c to $\tau_i(a)$ from equation 5. Of course, the budget can no longer be balanced solely with transfer payments: the amount c will need to be subsidized by those who benefit from the evaluation, as discussed in ARZ and above in section 2.

A modified log-likelihood scheme can also be used when, in addition to inducing truthful reporting, the center must induce buyers to gather the appropriate amount of information. Since a positive multiple of the logarithmic scoring rule is also strictly proper, by rescaling all payments proportionately the center can affect incentives to gather information without affecting the buyer's incentive to tell the truth.

To illustrate how such a scheme might work, we must generalize the information structure of the problem to allow the buyer to obtain "more precise" information.[22] To do so, we assume that the buyer's experience with the seller is encoded not as a single outcome, but rather as an infinite

[22] Clemen (2001) undertakes a similar investigation in a principal-agent model.

sequence of outcomes generated by repeated, independent sampling from distribution $f(s_m|t)$. By the law of large numbers, if the buyer knew the entire sequence, the seller's true type would be revealed. However, in this version of our model, we assume that the buyer must put forth effort to learn about his experience, letting $c_i(x_i)$ be the cost of learning the first x components of his experience, where $c_i(x_i)$ is strictly positive, strictly increasing, and strictly convex, and assumed to be known by the center.[23] We will refer to x_i as the precision of agent i's information.

For a buyer who already knows the first x components of his experience, learning the $x+1^{st}$ component represents a further partition of the outcome space. Savage (1954) studies the value of a further partition of the outcome space, which he calls the "partition problem," and our analysis follows in a similar vein.

We begin by arguing that, holding fixed the precision of agents' information, the scoring-rule based information can elicit this information. We then turn to the question of how the mechanism can be used to induce agents to acquire more precise information, even though such acquisition is costly.

For any fixed x_i, the information content of two possible x_i component sequences depends only on the frequencies of the various outcomes and not on their order. Consequently, let $Y^i(x_i)$ be the M dimensional random variable whose m^{th} component counts the number of times outcome s_m occurs in the first x_i components of the agent's information.[24] Let $y^i = \left(y^i_1, ..., y^i_M\right)$ denote a generic realization of $Y^i(x_i)$, where y^i_m is the number of times out of x_i that signal s_m is received, and note that $\sum_{m=1}^{M} y^i_M = M$. Based on his observation of $Y^i(x_i)$, buyer i can compute his posterior beliefs about the seller's type, which are informative about the expected distribution of the other players' signals. Since different realizations of $Y^i(x_i)$ yield different posterior beliefs about the seller's type, we can extend Lemmas 1 and 2 to the multiple signal case. In the remainder of this section, we let $g\left(y^j(x_j)\,|y^i(x_i)\right)$ denote the distribution of $Y^j(x_j)$ conditional on $Y^i(x_i)$.

Lemma 3: *Consider distinct players i and j, and suppose x_i, $x_j \geq 0$ are commonly known. For generic distributions $f(s_m|t)$ and $p(t)$, $Y^i(x_i)$ is stochastically relevant for $Y^j(x_j)$. If agent i is asked to announce a realization of $Y^i(x_i)$ and is paid according to the logarithmic score of*

[23] In single-agent context, Clemen (2001) examines the incentive problem in the case where the center does not know $c_i(x_i)$.

[24] $Y^i(x_i)$ is a multinomial random variable with x_i trials and M possible outcomes. On any trial, the probability of the m^{th} is $f(s_m|t)$, where t is the seller's unknown type.

$Y^j(x_j)$, *given that announcement (i.e.,* $R\left(y^j(x_j)|y^i(x_i)\right) = \ln g\left(y^j(x_j)|y^i(x_i)\right)$*), then the expected payment is uniquely maximized by announcing the true realization of* $Y^i(x_j)$.

Proof: The proof is analogous to the proofs of Lemmas 1 and 2 above. ■

Proposition 2 restates Proposition 1 in the case where the precision of the buyers' information is fixed and possibly greater than 1, i.e., $x_i \geq 1$ for $i = 1,..., I$. It follows as an immediate consequence of Lemma 3.

Proposition 2: *Suppose buyer* i *collects* $x_i \geq 1$ *signals. For generic distributions* $f(s_m|t)$ *and* $p(t)$ *there exist balanced transfers under which truthful reporting is a strict Nash Equilibrium of the reporting game.*

Proof: The construction follows that in Proposition 1, using $Y^i(x_i)$ for the information received by buyer i and constructing transfers as in (3) and (5). Under the equilibrium hypothesis, $j = r(i)$ announces truthfully. Let a^i denote buyer i's announcement of the realization of $Y^i(x_i)$, and let transfers be given by:

$$\tau_i^*\left(y^j|a^i\right) = \ln g\left(y^j|a^i\right). \tag{10}$$

Under these transfers, truthful announcement is a strict best response. Defining $\tau_i(a)$ as in (5) using τ_i^* as specified in (10) balances payments without affecting incentives to tell the truth. ■

Proposition 2 establishes that truthful reporting remains an equilibrium when buyers can choose the precision of their information. We next turn to the questions of how and whether the center can induce a buyer to choose a particular x_i. Let j denote the buyer whose signal player i is asked to predict (i.e., let $r(i) = j$), and suppose buyer j's information has precision x_j and truthfully reports the realization of $Y^i(x_j)$. For simplicity, we omit argument x_j in what follows. Further, suppose that buyer i is paid according to the log-likelihood scheme described in (10) and (5). Since x_i affects these transfers only through buyer i's announcement, it is optimal for buyer i to truthfully announce $Y^i(x_i)$ regardless of x_i.

Since x_i is chosen before observing any information, buyer i's incentive to choose x_i depends on his ex ante expected payoff before learning his own signal, i.e., his payoff taking expectations over both $Y^i(x_i)$ and Y^j. This expectation is written as:

$$Z_i(x_i) = E_{Y^i}\left(E_{Y^j} \ln g\left(Y^j|Y^i(x_i)\right)\right).$$

Lemma 4 establishes that buyers benefit from better information, and is a restatement of the well-known result in decision theory that every decision maker benefits from a finer partition of the outcome space.[25]

Lemma 4: For generic distributions $f(s_m|t)$ and $p(t)$, $Z_i(x_i)$ is strictly increasing in x_i.

Proof: Fix x_i and let y^i be a generic realization of $Y^i(x_i)$. Conditional upon observing y^i, buyer i maximizes his expected transfer by announcing distribution $g\left(Y^j|y^i\right)$ for buyer j's information. Suppose buyer i observes the $x_i + 1^{st}$ component of his information. By Lemma 3, i's expected transfer is now strictly maximized by announcing distribution $g\left(Y^j|\left(y^i, s_m\right)\right)$, and buyer i increases his expected value by observing the additional information. Since this is true for every y^i, it is true in expectation, and $Z_i(x_i + 1) > Z_i(x_i)$.■

As is apparent in the proof, Lemma 4 applies to any strictly proper scoring rule, not merely the logarithmic scoring rule. The idea that observing additional information improves a decision maker's expected value is well known in the decision theory literature. For example, see Savage (1954), or Lavalle (1968).

Lemma 4 establishes that buyer i's information becomes more informative regarding buyer j's signal as x_i increases. An interesting aspect of this result is that the direct effect of gathering more information is to provide buyer i with better information about the seller, not buyer j. Nevertheless, as long as buyer i's information is stochastically relevant for that of buyer j, better information about the seller translates into better information about buyer j.

When transfers are given by (5), the expected net benefit to buyer i from collecting a sample of size x_i and truthfully reporting his observation is $Z_i(x_i) - c(x_i)$. Hence, transfers (5) induce buyer i to collect a sample of size $x_i^* \in \arg\max(Z_i(x_i) - cx_i)$. Since the second term of (5), i.e., the payment agent i makes to agent $b(i)$, does not affect buyer i's incentives, we suppress it for the remainder of the discussion and focus on (10), the part of the payment that is relevant for buyer i's incentives.

Buyer i's incentives to truthfully report are unaffected by a uniform scaling of all payments in (10). Therefore, by a judicious rescaling of the payments to buyer i, the center may be able to induce the agent to vary the precision of his information. Expression (11) extends the transfers

[25] See, for example, Savage (1954).

described in (10) to allow for multiple signals and a rescaling of all payments by multiplier $\alpha_i > 0$:

$$\tau_i^*\left(a^i, y^{r(i)}\right) = \alpha_i \ln g\left(y^{r(i)}|a^i\right). \tag{11}$$

Under transfers (11), the maximal expected benefit from a sample of size B_i is $\alpha_i Z_i(x_i)$. Hence the center can induce buyer i to select a particular precision, $\hat{x}_i$, if and only if there is some multiplier $\hat{\alpha} > 0$ such that $\hat{x}_i \in \arg\max \hat{\alpha} Z_i(x_i) - c(x_i)$. The simplest case is where $Z_i(x_i)$ is concave, i.e., where $Z_i(x_i+1) - Z_i(x_i)$ decreases in x_i.

Proposition 3: If $Z_i(x_i+1) - Z_i(x_i)$ decreases in x_i, then for any precision $x_i \geq 0$ there exists a scalar $\hat{\alpha}_i \geq 0$ such that when paid according to (11), buyer i chooses sample size $\hat{x}_i$.

Proof: Since $Z_i(x)$ is concave, sample size $\hat{x}_i$ is optimal if there exists $\hat{\alpha}_i$ satisfying

$$\begin{aligned}
\hat{\alpha}_i Z_i(\hat{x}_i) - c_i(\hat{x}_i) &\geq \hat{\alpha}_i Z_i(\hat{x}_i+1) - c_i(\hat{x}_i+1), \text{ and} \\
\hat{\alpha}_i Z_i(\hat{x}_i) - c_i(\hat{x}_i) &\geq \hat{\alpha}_i Z_i(\hat{x}_i-1) - c_i(\hat{x}_i-1).
\end{aligned}$$

Solving each condition for $\hat{\alpha}_i$,

$$\begin{aligned}
\hat{\alpha}_i &\leq \frac{c_i(\hat{x}_i+1) - c_i(\hat{x}_i)}{Z_i(\hat{x}_i+1) - Z_i(\hat{x}_i)}, \text{ and} \\
\hat{\alpha}_i &\geq \frac{c_i(\hat{x}_i) - c_i(\hat{x}_i-1)}{Z_i(\hat{x}_i) - Z_i(\hat{x}_i-1)}.
\end{aligned}$$

Such an $\hat{\alpha}_i$ exists if and only if $\frac{Z_i(\hat{x}_i) - Z_i(\hat{x}_i-1)}{Z_i(\hat{x}_i+1) - Z_i(\hat{x}_i)} \geq \frac{c_i(\hat{x}_i) - c_i(\hat{x}_i-1)}{c_i(\hat{x}_i+1) - c_i(\hat{x}_i)}$; by our assumptions, this expression is always true.■

If $Z_i(x_i+1) - Z_i(x_i)$ does not decrease in x_i, then there may be some levels of precision that are never optimal.[26] Nevertheless, increasing the scaling factor never decreases optimal precision, and so while the center may not be able to perfectly control the buyers' effort choices, it can always induce them to put forth greater effort if it wishes.

[26] While we are not aware of any general results pertaining to the shape of the $Z()$ function, Clemen (2001) provides a number of examples of cases in which $Z_i(x_i+1) - Z_i(x_i)$ does decrease in x_i.

The problem of finding the set of sample sizes that maximize $\alpha Z_i(x_i) - c(x_i)$ for some α when $Z_i(x_i)$ is not concave is isomorphic to the problem in production theory of finding the set of outputs that maximize profit for some output price when the production function is not concave. The solution involves finding the set of outputs that remain on the convex closure of the technology set. See, for example, Mas-Colell, Whinston, and Green (1985, Section 5.D).

In practice, the center will not know each individual's cost structure for extracting signals of greater precision. However, the center may be able to estimate the distribution of cost structures, and then pick a scaling factor that will, in expectation, cause each buyer to give signals of the optimal precision.[27]

Non-common prior information

We have addressed the case in which the center and buyers share common prior beliefs about the seller. Fortunately, the proper-scoring-rule mechanism easily adapts to the case of non-common priors. In particular, buyers' prior beliefs need not be commonly held. However, in order for the center to properly interpret the buyers' announcements, it must have some knowledge of the buyers' prior beliefs.

To briefly illustrate, suppose that at time 0 all parties share "original" prior beliefs about the seller $p(t)$. At time 1, each buyer privately observes a finite-valued random variable H^i with range $H = \{h_1, ..., h_W\}$. Let h^i be a typical realization. Information h^i, which we call buyer i's history, summarizes buyer i's information before his interaction with the current seller. The history could take into account past interactions with other sellers, other information about the current seller, or anything else that influences his beliefs about the seller's type before interacting with her.

Let $\eta\left(h^i|t\right) > 0$ be the probability that $H^i = h^i$ when the seller is type t, and suppose that the distribution of H^i differs for different values of t. Based on the original priors $p(t)$ and his history h^i, buyer i can compute new beliefs about the seller's type by employing Bayes' rule. Denote these beliefs as $p\left(t, h^i\right)$, and note that generically, different histories lead to different interim beliefs. Thus, buyers with non-common priors are modeled as having observed different histories.

At time 2, buyer i observes signal S^i, which characterizes his current interaction with the seller. Thus the buyer's information in this model consists of both his history and signal, $\left(H^i, S^i\right)$. Structurally, this problem is identical to our original one, and therefore analogs to Lemmas 1 and 2 follow. Generically, buyer i's information, $\left(H^i, S^i\right)$, is stochastically relevant for buyer j's, $\left(H^j, S^j\right)$. Hence, the center can elicit buyer i's information by asking buyer i to predict buyer j's announced information and paying him according to a strictly proper scoring rule.[28] The analog

[27] Of course, although the center chooses the scale that induces the optimal ex ante precision, buyers, who know their costs, will tend to choose lower precision if they are high cost and higher precision if they are low cost.

[28] Since $\left(H^i, S^i\right)$ is also stochastically relevant for S^j alone, buyer i could be asked to predict only buyer j's experience and not his history.

to Proposition 1 then shows that there are balanced payments that elicit truthful information from all buyers in a strict Nash Equilibrium.

The assumption of a common, original prior can be relaxed slightly. The arguments go through essentially unmodified if the original priors of the various buyers are not common but are known to the center. The key is that buyer i's belief about the distribution of buyer j's signal depends on buyer i's prior beliefs, but not on buyer j's.

If the center does not know a buyer's original prior, then it cannot differentiate between information contained in the buyer's current experience and information deriving from his history. That is because proper scoring rules can elicit only beliefs about the posterior distribution of the other buyer's signal. In the models considered in this paper, a one-to-one mapping between information and posterior distributions allows the information to be elicited. If the center does not know the buyer's original prior, then it cannot separate the influence of the buyer's prior and his experience with the seller. For example, a buyer with a good history who has a bad experience with the seller could have beliefs about the seller's quality similar to a buyer with a bad history but a good experience with the seller. Thus, while common knowledge is not crucial to our mechanism, knowledge of the buyers' prior is.

In certain circumstances, it may be possible for the center to elicit a buyer's prior information by asking him to reveal his beliefs about another agent's announcement before observing his own signal. For example, consider the case where each buyer's prior beliefs about the seller's type are private information, but the distribution of signals given types, $f(s_n|t)$, is commonly known. Suppose buyer 2's announcement is being used to provide incentives to buyer 1. Let $p' = (p_1, ..., p_T)$ be buyer 1's prior beliefs about the seller's type, where $p_t \geq 0$ and $\sum_{t=1}^{T} p_T = 1$, and suppose buyer 1 is asked, before observing his signal, to announce his beliefs about the distribution of buyer 2's announced signal. If buyer 1 is paid according to a strictly proper scoring rule based on his announced distribution, he will truthfully announce his beliefs about the distribution of buyer 2's signal.[29]

Let $q' = (q_1, ..., q_N)$ be buyer 1's initial belief about the distribution of buyer 2's announcement. Since q is a distribution, $q_n \geq 0$ and $\sum_{n=1}^{N} q_n = 1$. Under a strictly proper scoring rule, buyer 1 reveals q. Let $p' = (p_1, ..., p_T)$ be buyer 1's prior belief about the seller's type, where $p_t \geq 0$ and

[29] These payments could be made, for example, by buyer 3.

$\sum_{t=1}^{T} p_t = 1$. Distributions p and q are related according to:

$$q = Fp,$$

where F is the N by T matrix whose elements are $f(s_n|t)$.

The usefulness of this procedure depends on whether buyer 1's initial announcement effectively reveals his prior distribution. If $N = T$ and F is invertible, then $F^{-1}q = p$, and q reveals p. However, if $N \neq T$ or F is not invertible, then there may be multiple prior distributions p consistent with a particular announcement q.[30] If the center is unable to form beliefs about which of the possible priors the buyer holds, then it may not be possible to extract meaningful information about the buyer's prior beliefs. Nevertheless, for many classes of prior beliefs such extraction is possible.

When priors can be extracted , buyer 1's signal can be elicited using a second round of transfers. These transfers are based on the scoring rule employed above, using the prior distribution revealed by the first round of elicitation.

Collusion

While optimal efforts and honest reporting are a Nash equilibrium, they are not a dominant strategy. All the players may collude to avoid effort while still collecting participation fees.[31] Or a

[30] Since buyer 1 wants to truthfully reveal his information, there will always be some p consistent with his announced q.

[31] In herding models, a similar lack of information revelation occurs even without collusion, but for different reasons. There, players receive personal payoffs for actions. Others try to infer a players' private information from his choice of action. Eventually, however, public priors swamp private information for each individual, and further actions no longer reveal any private information (Bikhchandani, Hirshleifer, and Welch, 1989; Banerjee, 1992). Our publicly revealed evaluations are analogous to the publicly revealed actions in herding models, but with an important difference that makes full information revelation possible. With a proper scoring rule, the expected payoffs that a player receives from a report depend on the priors that the player perceives, whereas individual payoffs for actions are fixed in herding models regardless of the public information available. As public information accumulates in herding models it swamps the private information, but in our situation the scoring rule effectively factors out the public information from a players' choice incentives.

Related problems arise in corporate herding models, where players may not reveal their information. The "yes men" of Prendergast (1993) give reports that match their impressions of their bosses' private information. In our situation, however, the early evaluator has no private information about a future evaluator's private signal, other than the stochastic relevance of his own signal. Ottavianni and Sorensen (2002) model a situation where evaluators subset of the players may collude in a way that induces the remaining players to pay net transfers to the colluding group. Mechanisms built around proper scoring rules are not invulnerable to collusion, but they have features that deter it. For example, if there is real time phasing to evaluations, that will make it difficult for the evaluators to gather together to agree on a plan. If collusion requires side payments, that puts another burden on potential colluders.

More active defenses may also be available. Take a situation familiar to academic readers, grading. Say we told all professors that they would get paid by how closely their grades corresponded to the grades given the same students by other professors. If the payments were significant, many would give every student an A. Forcing professors to grade on a curve can defeat this behavior. If we told professors they must give out one third each of As, Bs and Cs, they would have a difficult time coordinating to whom to give each grade, particularly if exams were graded anonymously.

Collusion can also be deterred by capitalizing on the natural variability in evaluations. Even good items sometimes produce poor outcomes. Say this likelihood is 0.1. If we had 100 evaluators, and 99 said an item was good, the system would "know" that there had been collusion – since this outcome is three standard deviations from the mean – and punishment could be meted out. Indeed, a range of techniques for deterring collusion look for "excess agreement." More sophisticated collusion – such as agreeing who would be negative on which items – could be met by more sophisticated countermeasures, such as randomizing the order of the items.

Efforts to deter collusion in reputation systems may prove like anti-submarine warfare. The advantage will shift back and forth between the offense and the defense, depending on new technologies and how much each side gains if it triumphs.

4 Alternative Scoring Approaches

It is instructive to consider two logical alternative scoring approaches that work less well. The effective scoring rules presented above compare the posterior beliefs after a player reports his signal to the actual signal seen (and reported) by a future player.

are compensated based not on the honesty of the reports of their private signals, but on the inferred precision of the private signals, discounting, and expectations about dishonesty in reporting. In such a situation, they find that full revelation of private signals is not possible in equilibrium. In our model, we compensate for correct posteriors rather than for correct private signals.

		Type	
		G	B
Signal	a	1/2	1
	b	1/2	0

Figure 1: Distribution of signals for each type.

A much simpler approach would simply compare a player's report to the report of some future player and charge based on how close they are. The obvious problem is that buyers would not have an incentive to report their private signals in this case but rather their posterior beliefs about the next signal. The posterior beliefs would be informed not only by the private signal but also by the prior beliefs as well. The approach proposed in this paper avoids this problem because the payment scheme automatically takes into account the buyer's priors, so that simply reporting the private signal yields the optimal payoffs.

A somewhat better approach would be to compare posteriors. The most straightforward scheme would pay a rater based on how far the mean of the posterior distribution after his rating (the prior for the future) diverges from the mean of the posterior after some number of additional raters. Since the mean of the prior equals the expected value of the posterior mean after further information is secured, and given the mean's well known least-squared-error properties, it might seem that charging the rater on a squared distance measure would induce him to report his true posterior mean.

Scoring based on the squared distance of the posterior means would work if the first report can do nothing to shift the variance of the future mean.[32] However, for many important situations a player's report will affect that variance, which introduces an incentive to distort the report. Consider the following example, where there are two seller types, G and B (for good and bad), and two possible signals, a and b. Initially, types G and B are equally likely. Figure 1 gives the distribution of signals for each type.

There will be two buyers in order. The first buyer gets a signal a. Employing Bayes Theorem,

[32] Thus the variance could be known and constant, or shrinking at a known rate with further evaluations.

		Posterior Given Signal	
Buyer 2's Signal	Frequency of Signal	Type G	Type B
a	5/6	1/5	4/5
b	1/6	1	0

Figure 2: Posterior distribution for each observed signal.

the posterior for types G and B become 1/3 and 2/3. Assuming the second buyer reports honestly, Figure 2 gives his posterior beliefs after observing either signal.

The 5/6 probability of observing a is comprised of two events. With probability 1/6, the seller is type G and buyer 2 receives signal a, while with probability 2/3 the seller is type B and buyer 2 receives a. It is readily computed that if the first buyer honestly reports the signal a, the expected charge is $K\left[\frac{5}{6}(\frac{1}{3}-\frac{1}{5})^2+\frac{1}{6}\left(\frac{1}{3}-1\right)^2\right]=K\frac{4}{45}$, where K is the scaling constant and $\frac{4}{45}$ is the expected squared distance of the posterior probability of type G from the $\frac{1}{3}$ estimate.

However, buyer 1 has a superior strategy. If he simply lies and says he saw b, the posterior goes to Type G $=1$, Type B $=0$. Moreover, whatever the next buyer observes, the posterior will remain at Type G $=1$. Thus, his expected charge will go to 0.

The explanation for this apparent paradox is that the variance of the posterior mean after buyer 2's report depends on what buyer 1 says.[33] Given that buyer 1 can reduce the variance with one report but not the other, a quadratic loss function will no longer give him an incentive to report honestly.

[33] We thank Ennio Stachetti for alerting us to this problem. The problem is somewhat analogous to the one analyzed by Brandenburger and Polak (1996). In their formulation, a manager may choose sub-optimal actions given her private information, in order to suggest a more favorable state of the world. The stock market determines stock price based on its posterior assessment of the state of the world given the manager's actions, without knowing the true state. In our counterexample, a false report (analogous to a sub-optimal action by the manager) may lead everyone to believe that a particular state of the world exists, one in which there will be little variance in the future posteriors and hence a better payoff to the buyer making the report.

5 Conclusion

The grand bazaar of electronic markets produces many temptations: for sellers they are sloth and misrepresentation, for buyers the lure is to shirk from the collective endeavor of producing accurate seller reputations. Both these temptations could be overcome through the effective use of reputation systems.

We show that a mechanism that capitalizes on stochastic relevance between the reports of different buyers, and that employs proper scoring rules and monetary payments, can give buyers appropriate incentives to report honestly. Beyond this, it can balance the budget and induce the right amount of effort when effort is costly.

Any such mechanism would require the widespread collection and dissemination of information about reputation, and computationally intensive calibration and distribution of payments. The vast reach of the Internet and its ready connection to superfast computers suggest that electronic markets could implement such mechanisms. The reputation systems of the future compared to those of today, such as eBay's, will be as a modern jumbo jet is to a Model T Ford. They will accomplish much the same task, but will operate at many times the scale, with vastly greater sophistication, using a radically different technology.

References

[1] Avery, Chris, Paul Resnick, and Richard Zeckhauser (1999): "The Market for Evaluations," *American Economic Review*, Vol. 89(3) 564-584.

[2] Banerjee, Abhijit (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics,* Vol. 107(3) 797-818.

[3] Banerjee, Abhijit (1993): "The Economics of Rumours," *Review of Economic Studies,* Vol. 60(2) 309-327.

[4] Bikhchandani, Suchil, David Hirshleifer, and Ivo Welch (1989): "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy,* Vol. 100(5) 992-1026.

[5] Brandenburger, Adam, and Ben Polak (1996): "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See," *The RAND Journal of Economics,* Vol. 27(2) 523-541.

[6] Clemen, Robert (2001): "Incentive Contracts and Strictly Proper Scoring Rules," Working Paper, Faqua School of Business, Duke University.

[7] Cooke, Roger M. (1991): *Experts in Uncertainty: Opinion and Subjective Probability in Science,* Oxford University Press: New York.

[8] Dellarocas, Chrysanthos (2000): "Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior," ACM Conference on Electronic Commerce EC-00, Minneapolis, http://ccs.mit.edu/dell/ec00reputation.pdf.

[9] ________________________________ (2001): "Analyzing the Economic Efficiency of eBay-like Online Reputation Reporting Mechanisms." Proceedings of the 3rd ACM Conference on Electronic Commerce, Tampa, FL, October 14-16, 2001.

[10] Hanson, Robin (2002): "Logarithmic Market Scoring Rules for Modular Combinatoral Information Aggregation," Working Paper, George Mason University Department of Economics, http://hanson.gmu.edu/mktscore.pdf.

[11] Johnson, Scott, Nolan Miller, John Pratt, and Richard Zeckhauser (2002): "Efficient Design with Interdependent Valuations and an Informed Center," Kennedy School Working Paper, RWP02-025, http://ksgnotes1.harvard.edu/research/wpaper.nsf/rwp/RWP02-025.

[12] Kandori, Michihiro (1992): "Social Norms and Community Enforcement," *Review of Economic Studies,* Vol. 59(1) 63-80.

[13] Kollock, Peter (1999): "The Production of Trust in Online Markets," *Advances in Group Processes,* Vol. 16.

[14] Kreps, David, and Robert Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory,* Vol. 27(2) 253-279.

[15] Lavalle, I (1968): "On Cash Equivalents and Information Evaluation in Decisions Under Uncertainty: Part I: Basic Theory," *Journal of the American Statistical Association*, Vol. 63(321), 252-276.

[16] Mas-Colell, A., M. Whinston, and J. Green (1995): *Microeconomic Theory,* Oxford University Press, New York.

[17] Ottaviani, Marco, and Peter Sorensen (2002): "The Strategy of Professional Forecasting," Working Paper, London Business School.

[18] Prendergast, Canice (1993): "A Theory of Yes Men," *American Economic Review,* Vol. 83(4) 757-770.

[19] Resnick, Paul and Hal Varian (1997). "Recommender Systems," *Communications of the ACM,*Vol. 40(3): 56-58.

[20] Resnick, Paul, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara (2000). "Reputation Systems: Facilitating Trust in Internet Interactions." Communications of the ACM 43(12): 45-48.

[21] Resnick, Paul, and Richard Zeckhauser (2001): "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," http://www.si.umich.edu/~presnick/papers/ebayNBER/index.html

[22] Savage, L. (1954): *Foundations of Statistics,* Dover Publications, Inc., New York.

[23] Selten, Reinhart (1978): "The Chain-Store Paradox," *Theory and Decision,* Vol. 9, 127-159.

[24] Selten, Reinhart (1998): "Axiomatic Characterization of the Quadratic Scoring Rule," *Experimental Economics,* Vol. 1, 43-62.

[25] Shapiro, Carl (1982): "Consumer Information, Product Quality, and Seller Reputation," *The Bell Journal of Economics,* Vol. 13 (1) 20-35.

[26] Shuford, E., A. Albert, and H. Massengil (1966): "Admissible Probability Measurement Procedures," *Psychometrika,* Vol. 31, 125-145.

[27] Tadelis, Steven (2002): "The Market for Reputations as an Incentive Mechanism," *Journal of Political Economy*, Vol. 110(4), 854-882.

[28] Winkler, Robert (1969): "Scoring Rules and the Evaluation of Probability Assessors," *Journal of the American Statistical Association,* Vol. 64, 1073-1078.