Designing Reputation Mechanisms for Efficient Trade

Christina Aperjis
Social Computing Lab, HP Labs, christina.aperjis@hp.com

Ramesh Johari
Department of Management Science and Engineering, Stanford University, ramesh.johari@stanford.edu

A seller in an online marketplace with an effective reputation mechanism should expect that dishonest behavior results in higher payments now, while honest behavior results in higher reputation—and thus higher payments—in the future. We study two widely used classes of reputation mechanisms. First, we show that weighting all past ratings equally gives sellers an incentive to falsely advertise. This result supports eBay's recent decision to base the Positive Feedback percentage on the past 12 months of feedback, rather than the entire lifetime of the seller. We then study reputation mechanisms that weight recent ratings more heavily. We characterize conditions under which it is optimal for the seller to advertise truthfully, and relate seller truthfulness to returns to reputation. If there is no reputation premium for a low value item, we show the following dichotomy: under increasing returns to reputation the optimal strategy of a sufficiently patient and sufficiently high quality seller is to always advertise honestly, while under decreasing returns to reputation the seller will not always be honest. Finally, we suggest approaches for designing a reputation mechanism that maximizes the range of parameters for which it is optimal for the seller to be truthful. We show that mechanisms that use information from a larger number of past transactions tend to provide incentives for patient sellers to be more truthful, but for higher quality sellers to be less truthful.

Key words: reputation mechanisms, ratings, online markets

In online trading communities, sellers have a temptation to dishonesty, because potential buyers have to decide how much to pay for an item without being able to observe it firsthand. In particular, the buyer typically cannot know in advance whether the seller is describing the item honestly, and hence may be afraid that he might be exploited if he trusts the seller. This effect is exacerbated because a buyer will often interact with sellers with whom he has never interacted before and may only seldom interact in the future. The absence of trust created by this information asymmetry may result in market failure (Akerlof 1970).

Aggregation mechanisms can be used to encourage buyers to trust sellers. Such mechanisms provide buyers with aggregate statistics on the past behavior of sellers; as a result, dishonest advertising by the seller involves a greater immediate payoff at the expense of a lower long-term payoff. Aggregation mechanisms typically operate as follows: after a transaction, the buyer rates the seller depending on how satisfied he was with the transaction. Then all ratings are aggregated into the seller's *score*. We note that electronic marketplaces typically also allow users to search for more detailed information on the seller's past ratings; nevertheless, aggregated statistics play an important role in buyers' decisions, because of the time and cognitive cost required to go through all available information.

Empirical studies have shown that sellers with high scores enjoy a *price premium*: on average they sell at higher prices than sellers with lower scores. For example, a controlled experiment on vintage postcard auctions showed that an experienced seller enjoyed an 8% price premium (Resnick et al. 2006), and data from eBay auctions revealed that a one point increase in the percentage of negative ratings led to a 9% decline in sale price (Cabral and Hortacsu 2010). These studies suggest that aggregation mechanisms can effectively link current actions of the seller with future payoffs, and thus may incentivize the seller to act in a way that benefits buyers and promote trust.

With the goal of incentivizing truthful advertising, we study how the marketplace is affected by (i) the *aggregation mechanism*, i.e., the way ratings are aggregated into the seller's score, and (ii) the *payment function*, i.e., the way buyers interpret the seller's score. Our main contributions are the following.

1. We demonstrate shortcomings of a system where all ratings are weighted equally in the score of the seller. Such a mechanism will typically be unable to guarantee that a seller is truthful after completing a large number of transactions.

2. We define an aggregation mechanism which weights ratings on recent transactions more. We characterize conditions under which it is optimal for the seller to advertise truthfully, and relate seller truthfulness to returns to reputation.

3. We propose approaches to choose the aggregation mechanism that achieves seller truthfulness over the largest range of parameters, and show that displaying the percentage of good ratings of the seller within a fixed number of transactions is not optimal.

In Section 1 we introduce the model. We assume that in each period the seller has a high or low value item for sale. At the beginning of each period the seller observes the value of the item and decides how to advertise it. Potential buyers observe the seller's advertisement and his *score*, i.e., an aggregate of the seller's past ratings. Using this information, we postulate that buyers employ simple heuristics (e.g., Tversky and Kahneman 1974), Bayesian techniques, or some combination of the two, in order to decide how much to bid (Resnick et al. 2006). The paper, however, does not require specific models for the buyers' behavior.

The expected payment to the seller is a function of his advertisement and his score, called the *payment function*. Motivated by empirical studies (e.g., Ghose et al. 2005, Cabral and Hortacsu 2010), we assume that the payment function is increasing in the seller's score. Moreover, we assume that the payment is increasing in the seller's advertisement, that is the expected payment that the seller receives is higher when he advertises a high value item. Under these assumptions, a rational seller advertises high value items truthfully, but may exaggerate the value of a low value item in his advertisement. The seller faces the following trade-off: falsely advertising may result in a higher payment now, but also in a bad rating which implies lower scores and payments in the future.

After a transaction the buyer truthfully reports whether the seller advertised the item accurately. Then, the aggregation mechanism uses this information to compute the seller's new score. Note that here we do not concern ourselves with incentives for buyers to rate truthfully; our focus is entirely on the optimal strategy of the *seller*.

In Section 2 we study an aggregation mechanism that weights all ratings equally (the Unweighted Aggregation Mechanism). We show that with this mechanism the seller is incentivized to falsely advertise at some rating histories for a large class of payment functions, and thus demonstrate that this aggregation mechanism is not effective. In particular, the temptation for the seller to become dishonest increases as the total number of ratings increases, because each additional rating has a smaller and smaller effect on the seller's overall rating.

Our result on the Unweighted Aggregation Mechanism supports eBay's recent decision to base the Positive Feedback percentage on the past 12 months of feedback, rather than the entire lifetime of the seller[1]. We refer to the previous mechanism employed by eBay, which was in use from March 2003 until May 2008, as *EBAY03-08*. In *EBAY03-08* all ratings were weighted equally in the information shown next to the item's description, since lifetime statistics were given. Even though detailed information on recent ratings was available in that system (by clicking through to view the seller's profile), buyers did not always spend the time and effort to search for detailed information. For instance, data collected before and after the change in eBay's aggregation mechanism in March 2003 showed that buyers respond to the reputation information that is easiest to access (Cabral and Hortacsu 2010). This suggests that it is better to primarily show information that weights recent ratings more, as validated by our analysis.

[1] http://pages.ebay.com/help/feedback/scores-reputation.html

In Section 3 we study the class of Weighted Aggregation Mechanisms, which weight ratings from recent transactions more heavily. In particular, the score that is shown to the buyers is a weighted average of past ratings of the seller. We study this class of mechanisms because this is a natural way to weight recent ratings more, and because it includes as special cases several well known aggregation mechanisms. The expected payment that the seller receives is a function of his score. Depending on how buyers interpret the seller's score, the payment to the seller may take different forms.

In contrast to the Unweighted Aggregation Mechanism, under the Weighted Aggregation Mechanism it becomes possible, under some additional conditions, to ensure the seller is always truthful. Our main contribution is to identify for which payment functions it is possible or impossible to have a seller that always advertises truthfully. We show that under increasing returns to reputation (i.e., the payment function is convex) the optimal strategy of a sufficiently patient and sufficiently high quality seller is to advertise honestly. On the other hand, we show that if returns to reputation are decreasing (i.e., the payment function is concave) and there is no premium for low value items, then for any seller it is optimal to falsely advertise at some scores. The intuition behind this result is that when returns to reputation are increasing, the seller suffers a large reduction in future payment due to even a single deviation from truthful advertising. By contrast, this reduction is relatively small when returns to reputation are decreasing. We also study the effect of step payment functions on the seller's optimization, and give conditions under which it is optimal for the seller to be truthful. In this setting, our main finding is that to make a sufficiently patient and sufficiently high quality seller truthful, the mechanism must give sufficient weight to the two most recent ratings.

In Section 4 we address the design question of choosing the right aggregation mechanism. We note that online marketplaces often change their aggregation mechanisms, which indicates that it is hard to design an effective aggregation mechanism. We first study the Window Aggregation Mechanism, a widely used Weighted Aggregation Mechanism which only reveals the percentage of positive ratings within some fixed window. We define an optimal window size as one which maximizes the range of parameters for which the seller is truthful, and we study the dependence of the optimal window size on the parameters of the model. The optimal sufficient statistic (i.e., window size) neither has very short-term memory (since then current actions do not affect future scores), nor has very long term memory (since then the seller is not incentivized to be truthful). The best choice of averaging window lies between these extremes. We also note an interesting qualitative tradeoff in the choice of window size: informally, increasing the window size is more likely to make patient sellers truthful, while it is less likely to make high quality sellers truthful.

We then formulate the design problem for the general class of Weighted Aggregation Mechanisms; here the goal is to choose the optimal vector of weights applied to past ratings to compute a score. Our results here match the insights obtained in our study of the Window Aggregation Mechanism. Moreover, our insight suggests that the Window Aggregation Mechanism (which weights all ratings within some fixed window equally) is not optimal and a more robust mechanism might be one that weights more recent ratings even more heavily than older ratings. This is particularly interesting when one considers that nearly all online marketplaces that use mechanisms which only weight recent ratings, tend to use a Window Aggregation Mechanism to do so.

In Section 5 we demonstrate that the results also hold for a setting with multiple possible values for the item and the ratings, where a buyer's rating depends on the difference between the advertised and true value of the item. This setting shows our results are robust even if the seller has potentially many "levels" of dishonesty possible.

1. Model

We consider a single seller who is a long-lived player with discount factor δ. The seller interacts with short-lived potential buyers, i.e., buyers who are interested in the seller's item for exactly one round and then depart. We do not explicitly model the market mechanism at each time period, and instead abstract the aggregate behavior of all the buyers in a single time period via a single premium function, described further below.

In every period the seller has an item for sale whose value is high (v_H) with probability q_H and low (v_L) with probability $1-q_H$. We assume that $0 \le v_L < v_H \le 1$. The seller observes the value of the item at the beginning of a period and decides what advertisement to post. Potential buyers observe the seller's advertisement and his *score*, i.e., an aggregate of the seller's past ratings. The expected payment to the seller is thus a function of his advertisement and his score.

After purchase, the buyer truthfully reports whether the seller described the item correctly. The rating is good if the seller described the item truthfully and bad otherwise (i.e., a binary rating system). For simplicity, for most of this paper we assume that both the item's value and the rating can only have two possible values; however, these assumptions can be weakened. In Section 5 we demonstrate that our results can be generalized for a setting with multiple possible values for the item and ratings, where the value of the rating depends on how much the seller exaggerates the item's value in his description.

An *aggregation mechanism* specifies the rules for calculating the *score* (i.e., the information shown to potential buyers) from past ratings of the seller. Let r_i be the i-th most recent rating of the seller. In particular $r_i = 1$ or $r_i = 0$ depending on whether the corresponding rating was good or bad. Let $\vec{r} = (r_0, r_1, ...)$ be the vector of ratings that the seller has received up to now. We denote the aggregation mechanism by $s(\vec{r})$. In particular, $s(\vec{r})$ is a function that maps rating vectors to scores. Both the seller and the mechanism know $\vec{r}$, but potential buyers only observe $s(\vec{r})$. The score may either be a scalar or a vector. In the next sections we consider both cases: in the Unweighted Aggregation Mechanism (Section 2) it is a vector consisting of the total number of ratings and the number of good ratings, while in the Weighted Aggregation Mechanism (Section 3) it is a scalar.

We assume that the expected payment that the seller receives when his score is s and he chooses advertisement a, where a is either high or low is $a \cdot b(s)$; we call $b(s)$ the *premium function*. We assume that the premium function is non-negative and non-decreasing.

Two realistic assumptions are incorporated in the form of the premium function we chose. First, better ratings yield higher payments to the seller,[2] as has been shown by empirical studies (e.g., Ghose et al. 2005, Cabral and Hortacsu 2010). Second, the premium function is increasing in the advertisement a.

The seller chooses a policy that is a best response to the premium function $b(\cdot)$. In our model, we emphasize that the seller is not intrinsically honest or dishonest; he is rational and chooses the advertisement that maximizes his payoff. This is in contrast to the adverse selection approach, where the seller has an intrinsic type (see Section 6). Under the assumptions we made on the payment, it is optimal for the seller to advertise a high value item truthfully.

Let $V(\vec{r})$ be the maximum infinite horizon discounted payoff of the seller when his current vector of ratings is $\vec{r}$. The seller's optimal policy is given by solving the following dynamic program.

$$V(\vec{r}) = q_H(v_H \cdot b(s(\vec{r})) + \delta \cdot V(1,\vec{r})) + (1-q_H)\max\{v_H \cdot b(s(\vec{r})) + \delta \cdot V(0,\vec{r}), v_L \cdot b(s(\vec{r})) + \delta \cdot V(1,\vec{r})\} \tag{1}$$

In particular, with probability q_H the seller has a high value item for sale, which he advertises truthfully. The immediate payment he receives is $v_H \cdot b(s(\vec{r}))$ and his ratings "increase" to $(1,\vec{r})$. With probability $1-q_H$ the seller has a low value item for sale. If he advertises it as a high value item, his payoff is $v_H \cdot b(s(\vec{r})) + \delta \cdot V(0,\vec{r})$, since he receives $v_H \cdot b(s(\vec{r}))$ now, but his ratings "decrease" to $(0,\vec{r})$. If he advertises truthfully, he receives a low payment now ($v_L \cdot b(s(\vec{r}))$), but his ratings "increase" to $(1,\vec{r})$. The seller will choose the advertisement with the maximum payoff.

[2] In particular, we assume that for any $n, m \ge 0$, we have $b(s(\vec{x};1;\vec{y}),a) \ge b(s(\vec{x};0;\vec{y}),a)$ for all $\vec{x} \in \{0,1\}^n, \vec{y} \in \{0,1\}^m$.

We say that the seller is *truthful at* $\vec{r}$ if it is optimal for him to advertise a low value item truthfully when his vector of ratings is $\vec{r}$. By (1), it is optimal for the seller to be truthful at $\vec{r}$ if and only if

$$(v_H - v_L)b(s(\vec{r})) \leq \delta(V(1,\vec{r}) - V(0,\vec{r})). \tag{2}$$

In particular, if the seller is untruthful, his current payoff will increase by $(v_H - v_L)b(s(\vec{r}))$ but his expected payoff starting from the next period will decrease by $V(1,\vec{r}) - V(0,\vec{r})$ (relative to being truthful).

We use this model to study the seller's optimal strategy under various aggregation mechanisms. In particular, we are interested in which aggregation mechanisms and which premium functions induce truthful behavior.

2. Unweighted Aggregation Mechanism

In this section, we consider an aggregation mechanism that weights all ratings equally, and show that the seller has an incentive to falsely advertise after a sufficiently large number of ratings. This suggests that weighting recent ratings more is a necessary condition for efficiency.

In the *Unweighted Aggregation Mechanism*, the seller's reputation score consists of the total number of ratings (s_T) and the number of positive ratings (s_P). In particular, given the vector of past ratings $\vec{r}$, we have $s_T(\vec{r}) = |\vec{r}|$ and $s_P(\vec{r}) = |\{r_i : r_i = 1\}|$. (We use the notation $|\cdot|$ to denote both the number of components of a vector and the cardinality of a set.) At score (s_P, s_T), we assume that the seller receives payment $a \cdot b(s_P, s_T)$ for advertising an item of value a.

Proposition 1 shows that under some assumptions on the payment function, the seller will eventually be better off falsely advertising a low value item. The intuition is that when all ratings are weighted equally, after a large number of ratings one more positive rating does not make an appreciable difference to the seller's payoff. This was partially the case with eBay's aggregation mechanism until May 2008 (*EBAY03-08*), where the information that was shown to potential buyers on the item description page weighted all ratings equally. For clarity, all proofs are in the appendix.

PROPOSITION 1. *If* $b(s_T, s_T)$ *is bounded away from zero as* $s_T \to \infty$ *and*

$$b(s_T, s_T) - b(s_T - 1, s_T) \to 0 \text{ as } s_T \to \infty, \tag{3}$$

then there exists an s_T *at which it is optimal for the seller to falsely advertise a low value item.*

Proposition 1 relies on two assumptions. The assumption on $b(s_T, s_T)$ is not particularly restrictive: we expect the payment to a seller with maximum reputation to be bounded away from zero after a large number of ratings.

Condition (3) is a regularity condition on the slope of the payment functions; this formalizes the idea that as the total number of ratings increases, the marginal effect of each additional positive rating on the seller's expected payment eventually becomes negligible. This condition is not particularly restrictive; e.g., see Example 2. On the other hand, if Condition (3) does not hold, then it may be optimal for a sufficiently patient seller to be always truthful. In particular, this is the case if a single bad rating causes a discrete drop in premium regardless of how many positive ratings have been received, i.e., if there exist T and $\alpha > 0$ such that $b(s_T, s_T) - b(s_T - 1, s_T) > \alpha$ for all $s_T \geq T$, and $b_T(1) < q_H \alpha\delta/(1-\delta)$.

We briefly discuss related results in the literature. A similar result is shown by Fan et al. (2005) for a specific payment function and two specific unweighted reputation mechanisms. Moreover, we observe that Proposition 1 is similar to the result in Cripps et al. (2004) that reputations are not sustainable with imperfect monitoring and incomplete information. This is a different setting than ours, since buyers are optimizing with respect to the seller's strategy, while we assume that the

Figure 1 **Sample eBay auction listing. 99% of the ratings that user satmr2 received in the last twelve months were positive.**

buyers' behavior only depends on the aggregation mechanism that is being used, and the seller's score and advertisement. The result by Cripps et al. (2004) depends on imperfect monitoring: the seller chooses low effort with some small probability, because it will not significantly affect the buyers' beliefs in the near future, and eventually buyers learn the seller's type. In our model, after a large number of transactions it becomes profitable for the seller to be untruthful, even though monitoring is perfect.

Example 1. ***EBAY03-08***
The seller's score shown by the *EBAY03-08* mechanism next to the description of the item consisted of: (1) the difference between positive and negative ratings (i.e., $2s_P - s_T$), and (2) the ratio of positive ratings over the total number of ratings (i.e., s_P/s_T). Users could access more information on the seller's past ratings by clicking on the seller's pseudonym. However, if many users did not spend the time to search for more information when bidding, this aggregation mechanism would not be effective in the long run for a large class of payment functions. In particular, it is possible that a seller would exaggerate the value of the item once he had a large number of ratings. This intuition is supported by data analyzed by Resnick and Zeckhauser (2002), which show a decline in performance once the seller completed a large number of transactions. This observation may have influenced eBay's recent decision to change the initially shown information; the new system includes statistics about the ratings the seller received in the last twelve months (cf. Figure 1). □

Proposition 1 shows that the seller will eventually be better off falsely advertising a low value item. How many transactions pass before the seller is tempted to be dishonest? We can compute a bound on this time by upper bounding the increase in payment due to one more positive rating with a function of the total number of ratings. In particular, suppose there exists a function $f(s_T)$ such that

$$\max_{0\le s_P\le s_T-1}\{b(s_P+1,s_T)-b(s_P,s_T)\}\le f(s_T).$$

Let s_T^* denote the first time a seller would choose to falsely advertise a low value item; it follows from the proof of Proposition 1 that an upper bound for s_T^* is:

$$\hat{s}_T=\min\left\{s_T : b(s_T,s_T)\ge\frac{\delta}{1-\delta}f(s_T)\right\}. \tag{4}$$

By computing $\hat{s}_T$ we know that the seller will not advertise low value items truthfully for more than $\hat{s}_T$ consecutive transactions. The following example upper bounds s^*_T for a specific payment function.

EXAMPLE 2. Let $b(s_P, s_T) = \alpha/(\alpha+1-s_P/s_T)$ for some $\alpha \in (0,1)$; this model arises when buyers use a certain maximum likelihood estimate based on the rating history, (cf. Appendix 2, Aperjis and Johari 2010). Note that this function only depends on the proportion of positive ratings. Further $b(s_T, s_T) = 1$, and $b(s_P+1, s_T) - b(s_P, s_T) \leq 1/(s_T \cdot \alpha)$. Let $f(s_T) = 1/(s_T \cdot \alpha)$. By (4),

$$\hat{s}_T = \left\lceil \frac{\delta}{1-\delta} \frac{1}{\alpha} \right\rceil.$$

This upper bound on the number of consecutive truthful transactions of the seller is increasing in δ, the seller's discount factor. This is something we expect: the more the seller values the future, the greater the effectiveness of the aggregation mechanism, and thus the greater his incentive to tell the truth. □

3. Weighted Aggregation Mechanism: Characterization

In Section 2 we showed that weighting all ratings equally does not incentivize truthfulness. In this section we analyze mechanisms which put more weight on recent ratings, and show that it is possible to incentivize truthfulness and promote trust.

We first introduce the *Weighted Aggregation Mechanism*; as we will discuss, special cases of this mechanism are widely used in practice. Let r_i be the value of the i-th most recent rating. We consider an aggregation mechanism where buyers only see the following score:

$$s(\vec{r}) = \sum_{i=0}^{\infty} w_i \cdot r_i,$$

where $w_i \geq 0$, and $w_0 \geq w_1 \geq w_2 \geq \cdots$, with $\sum_i w_i = 1$. We assume that the weights w_i are non-increasing in i so that recent ratings are weighted more. We make the seller's score a scalar, because aggregation mechanisms typically rely primarily on a summary of ratings.

According to the general model introduced in Section 1, at score s, the seller receives expected payment equal to the product of his advertisement and the premium function $b_{\vec{w}}(s)$. We use the subscript $\vec{w}$ on the premium function to denote that the payment may explicitly depend on the aggregation mechanism. *We assume that $b_{\vec{w}}(s)$ is increasing in s for each vector $\vec{w}$.*

Both the seller and the mechanism have access to $\vec{r}$; the seller remembers his past actions and the mechanism keeps this information in order to update the seller's score. However, as we discuss in the following examples of Weighted Aggregation Mechanisms, it may not be necessary to keep the whole vector of ratings.

EXAMPLE 3. In the *Window Aggregation Mechanism*, $w_i = 1/T$ for $i = 0, 1, ..., T-1$, and $w_i = 0$ for $i \geq T$, for some $T \geq 1$, which we call the window size. Thus the score is the percentage of good ratings that the seller received in the last T transactions. To compute the seller's score, the aggregation mechanism needs to keep information on the T most recent ratings of the seller. This score is widely used in online marketplaces, such as eBay and Amazon, and has been studied in various settings (Dellarocas 2005, Aperjis and Johari 2010). □

EXAMPLE 4. In the *Exponential Aggregation Mechanism* $w_i = (1-\alpha)\alpha^i$, for some $\alpha \in (0,1)$. To compute the seller's score, the aggregation mechanism only needs to know the score of the seller in the previous period ($\hat{s}$) and his most recent rating (r_0). Then the new score is $\alpha\hat{s} + (1-\alpha)r_0$. We note that the Exponential Aggregation Mechanism has been previously suggested as a good design (Fan et al. 2005). Moreover, even though to the best of our knowledge exponential smoothing is not being used by any electronic marketplace to promote trust, it is a good model of how people update their impressions without an aggregation mechanism in place (e.g., Anderson 1981, Hogarth and Einhorn 1992, Kashima and Kerekes 1994). □

We study the Weighted Aggregation Mechanism because (1) it is a natural way to weight recent ratings more, and (2) it is a generalization of both the Window Aggregation Mechanism and the Exponential Aggregation Mechanism. However, there are other ways to summarize feedback by incorporating text comments (which are more descriptive than numerical ratings) and weighting recency. Such a summary score could offer richer information without increasing the buyers' search costs (Pavlou and Dimoka 2006).

In Section 3.1 we derive necessary and sufficient conditions for the truthful policy to be optimal for the seller. We use these conditions in subsequent sections to show that efficiency can be achieved under strictly increasing returns to reputation (Section 3.2), and under a step premium function (Section 3.3). In Section 3.4 we discuss empirical studies on eBay and Amazon. In Section 3.5 we show that under an additional assumption on the premium function, we can simplify the sufficient condition for the seller to be always truthful; this assumption is valuable in part because it is satisfied by all logarithmically concave premium functions. We use this assumption in Section 3.6, where we identify a dominance relation between premium functions such that if b_1 dominates b_2, then b_1 better incentivizes truthful advertisement.

3.1. Conditions for Optimality of Truthfulness

We are interested in whether it is optimal for the seller to be always truthful. It is optimal for the seller to be always truthful if and only if any one step deviation from the truthful policy (i.e., the policy of always advertising items truthfully) does not yield a higher payoff. Let $\hat{V}(\vec{r})$ be the infinite horizon discounted expected value when the seller's current ratings are $\vec{r}$ and the seller is always truthful. Let $s_i(\vec{r})$ be the seller's score after i periods if his current ratings are $\vec{r}$ and receives positive ratings in the following i periods. Note that $s_0(\vec{r}) = s(\vec{r})$. This implies that $\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r}) = (q_H v_H + (1-q_H)v_L)\sum_{i=0}^{\infty}\delta^i(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$ and $s_i(1,\vec{r}) - s_(0,\vec{r}) = w_i$. By (2), the seller is not better off deviating from the truthful policy when his ratings are $\vec{r}$ if and only if $(v_H - v_L)b_{\vec{w}}(s(\vec{r})) \le \delta(\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r}))$. We conclude that it is optimal for the seller to be always truthful if and only if the previous condition holds for all $\vec{r}$. Substituting $\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r})$ in the previous condition, the following lemma is proved.

LEMMA 1. *It is optimal for the seller to be truthful at all $\vec{r}$ if and only if*

$$b_{\vec{w}}(s(\vec{r})) \le \frac{q}{v_H - v_L}\sum_{i=0}^{\infty}\delta^{i+1}(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) \tag{5}$$

for all rating vectors $\vec{r}$.

Thus, if we know the premium function $b_{\vec{w}}(\cdot)$ and the seller's parameters δ and q_H, we can check whether it is optimal for the seller to be truthful by checking whether condition (5) is satisfied for all $\vec{r}$.

3.2. Increasing and Decreasing Returns to Reputation

Depending on how buyers interpret the information on the seller available to them, i.e., the score $\sum_i w_i \cdot r_i$, the payment to the seller may exhibit increasing or decreasing returns to reputation. Increasing returns to reputation correspond to a convex premium function, i.e., there are increasing marginal benefits to higher reputation. On the other hand, decreasing returns to reputation are associated with a premium function that is concave in the seller's score.

The following lemma indicates that it is easier to incentivize truthfulness under increasing returns to reputation, since it is more likely that (5) holds.

LEMMA 2. *(i) If the premium function $b_{\vec{w}}$ is strictly convex, $w_0 < 1$ and $b_{\vec{w}}(0) = 0$, then*

$$b_{\vec{w}}(s(\vec{r})) < \sum_{i=0}^{\infty}(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$$

for every $\vec{r}$.

(ii) If the premium function $b_{\vec{w}}$ is concave and $b_{\vec{w}}(0) \geq 0$, then

$$b_{\vec{w}}(s(\vec{r})) \geq \sum_{i=0}^{\infty}(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$$

for $\vec{r} = \vec{1}$.

We can get some intuition by considering condition (5) when the seller has maximum score, i.e., when $r_i = 1$ for all i. At maximum score, the gain for being truthful now (given that the seller is truthful in the future) is proportional to $\sum_i \delta^i(b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))$, while the gain from deviating now is $b_{\vec{w}}(1)$. When $b_{\vec{w}}$ is concave, then $b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)$ is much smaller than $b_{\vec{w}}(1)$, while for a strictly convex premium function, $b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)$ is significant relative to $b_{\vec{w}}(1)$.

Lemma 1 and Lemma 2 imply a dichotomy between convex and concave premium functions when $v_L = 0$ and $b_{\vec{w}}(0) = 0$. If the premium function exhibits increasing returns to reputation, and if δ and q_H are sufficiently large, then it is possible to achieve our goal of making the seller truthful regardless of his rating history. On the other hand, if the payment is concave in the score, then it is not optimal for the seller to be truthful at all rating histories; in particular, (5) can not be satisfied since if $v_L = 0$, then $q = q_H v_H$ and $q/(v_H - v_L) < 1$.

If $v_L = 0$, the condition $w_0 < 1$ is necessary for truthfulness under convex premia. In particular, if only the most recent rating affects the seller's reputation, then it will be optimal for the seller to exaggerate the value of a low value item in his advertisement when his reputation is high. This happens because the seller discounts future payments and prefers a high payment now to a high payment in the next period. We conclude that if the payment is convex, then the weights on the two most recent ratings must both be strictly positive to induce truthfulness.

If v_L is strictly greater than zero, then it may be possible to have a truthful seller under a concave premium. For any concave strictly increasing premium, there exists a sufficiently large v_L such that it is optimal for the seller to be always truthful for a given δ and q_H. In particular, this is the case if

$$\frac{v_L}{v_H - v_L} \geq \max_{\vec{r}} \left\{ \frac{b_{\vec{w}}(s(\vec{r}))}{\sum_{i=0}^{\infty} \delta^{i+1}(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))} \right\} - q_H.$$

We conclude that it is possible to have a truthful seller under both convex and concave premium functions. However, increasing returns to reputation tend to better incentivize truthfulness.

3.3. Step Functions

In Section 3.2 we considered premium functions that exhibit either increasing or decreasing returns to reputation throughout their domains. Another natural class of payments are step premium functions, which are functions that are constant throughout most of their domains. Our motivation for studying this class of functions is that it may provide a good approximation for the expected payment in electronic marketplaces. These premium functions capture the following intuition: a potential buyer is not willing to buy an item from a seller that has a very low score; however, there is some threshold on the seller's score, above which the buyer trusts the seller and is willing to pay up to his valuation for the item. Since different buyers may have different thresholds and different valuations for the item, the premium function may be a smoothed step function, e.g., a function

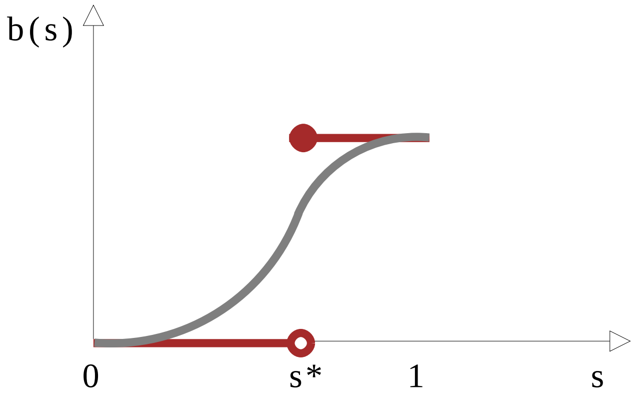

Figure 2 **Sigmoid and step payments. The grey line represents a premium function that is convex at low scores and concave at high scores. The red line represents a step function that approximates it.**

that is convex at low scores and concave at large scores (see Figure 2). In this section we study step premium functions, and note that similar results can be obtained for smooth approximations.

In this section we show that with step premium functions, optimality of the truthful policy heavily depends on the aggregation mechanism, and in particular on the magnitude of w_1 (i.e., the second largest weight) relative to the threshold of the payment. We then apply the result to the Window Aggregation Mechanism, which was introduced in Example 3.

We formally define a step premium function in the following definition.

Definition 1. The premium function $b_{\vec{w}}$ is a *step function* with threshold $s^* \in (0,1]$ if $b_{\vec{w}}(s) = b_{\vec{w}}(1) > 0$ for $s \geq s^*$, and $b_{\vec{w}}(s) = 0$ otherwise.

We are interested in understanding the conditions under which a Weighted Aggregation Mechanism can make a seller truthful when the payment is a step function. If $v_L > 0$ then it may be possible to incentivize the seller to be truthful under various vectors $\boldsymbol{w}$. In this section we focus on the case that $v_L = 0$. Our main result is the following proposition.

Proposition 2. *Suppose the premium function $b_{\vec{w}}$ is a step function with threshold s^* and $v_L = 0$.*

(i) If $s^* \leq 1 - w_1$, then there are no values of $q_H < 1$ and $\delta < 1$ for which it is optimal for the seller to be always truthful;

(ii) If $s^* > 1 - w_1$, then it is optimal for the seller to be truthful for all sufficiently large $\delta < 1$ and $q_H < 1$.

We note that w_1 is the weight on the second most recent rating and that $w_0 \geq w_1$. The result provides a useful insight: recent ratings, in particular the one and two period old ratings, must be weighted sufficiently to ensure that sellers can be made truthful.

As a specific application, we consider the Window Aggregation Mechanism, which was introduced in Example 3, and a premium function that yields a positive expected payoff for at least some reputation score(s) less than 1. These premium functions are reasonable in many applications: buyers often trust a seller with an almost maximum score. In the following example, we show that even though for any fixed window size it is not optimal for the seller to be always truthful, truthfulness may be optimal when information from multiple window sizes is aggregated. This suggests that the use of multiple window sizes, a common practice in online marketplaces, may better incentivize truthfulness.

Example 5. Assume that a Window Aggregation Mechanism with window size T is used and $v_L = 0$. The payment is a step function with threshold s^*. We consider premium functions that are positive for at least some possible scores other than 1 (the maximum possible score). This implies that $s^* \leq 1 - 1/T$. Thus, $s^* \leq 1 - w_1$. By Proposition 2 (ii), there do not exist $\delta < 1$ and $q_H < 1$ for which it is optimal for the seller to be always truthful.

Surprisingly, there is a way to aggregate information from multiple window sizes, so that it is optimal for the seller to always be truthful under step functions that are not only positive at the maximum possible score. For example, suppose the aggregation mechanism is using information on k windows with sizes $T_k \geq T_{k-1} \geq ... \geq T_1 \geq 2$. Let s_i be the score on window i and assume that the aggregate score of the seller is the average of the scores for each window, i.e., $s = (\sum_{i=1}^{k} s_i)/k$. Thus $w_1 = (1/k)\sum_{i=1}^{k} 1/T_i$, while the smallest positive weight is $w_{T_k - 1} = 1/(kT_k)$. In this case, by Proposition 2, the seller can be made truthful if $1/(kT_k) \leq 1 - s^* < (1/k)\sum_{i=1}^{k} 1/T_i$; depending on the value of s^*, these inequalities can be satisfied if more than one window is used. □

3.4. Empirical Insights

In Sections 3.2 and 3.3, we have identified conditions on the premium function that make it possible to incentivize truthfulness under the Weighted Aggregation Mechanism. A natural question arises: What is the form of the premium function in electronic marketplace? In this section we attempt to answer this question using results from empirical studies on eBay and Amazon. In particular, we wish to understand the dependence of the expected payment on the percentage of positive ratings (since this is the information shown in these markets that best matches the seller's score in a Weighted Aggregation Mechanism). We note however, that eBay's mechanism has changed multiple times within the last years, and thus different studies collect data on different versions of the mechanism.

Two empirical studies explore whether negative feedback is associated with increasing or decreasing returns, and suggest increasing returns in the percentage of positive ratings on eBay.[3]

1. Kalyanam and McIntyre (2001) collect data from Palm Pilot auctions on eBay and conjecture diminishing impact of negative feedback. In particular, they suggest that the dependence of price on percentage of negative feedback is $-f(\text{Negative }\%)$ for some concave function f, which implies increasing returns in the percentage of positive ratings.

2. A panel data set from eBay indicates that the first negative is very harmful for the seller's growth rate, while subsequent negatives have lower impact (Cabral and Hortacsu 2010). This implies increasing returns in the percentage of positive feedback for the growth, and suggests a similar effect for the expected payment to the seller.

Other studies suggest a specific dependence based on the regression that is used, without explicitly stating it. For instance, an OLS regression on data from Amazon's retail platform suggests that the price premium is convex in the number of stars, a measure similar to the percentage of positive ratings (Ghose et al. 2005). On the other hand, there are also studies which find that the percentage of positive feedback has no effect on the selling price (e.g., Resnick et al. 2006). We note that eBay was only showing the difference between the number of positive and negative ratings (and not the percentage of positive ratings) next to the description of the item when Resnick et al. (2006) conducted their experiment. Nevertheless, Kalyanam and McIntyre (2001) and others collected data under the same mechanism and found that the percentage of negative ratings had an effect on the expected price.

To interpret these contradictory results, we need to consider that most eBay sellers have an extremely high percentage of positive ratings (Resnick and Zeckhauser 2002). Low variability in the percentage of positive ratings may be one reason that, for some data sets (e.g., Resnick et al. 2006),

[3] There have also been studies that find decreasing returns in reputation; however, these studies look at different metrics than the one we are interested in here. Livingston (2005) studies the effect of the total number of positive ratings on the expected payment to the seller, and finds severely decreasing marginal returns. Obloj and Capron (2008) find decreasing marginal returns in the difference between the total number of positive and negative ratings. However, these results do not apply here, because we are interested in the effect of the percentage of positive ratings.

the price does not significantly depend on the percentage of positive ratings. Another possible interpretation is that the payment does not depend on the exact percentage of positive ratings, above a certain threshold; this behavior is consistent with a step premium function. On the other hand, the results of Kalyanam and McIntyre (2001) and Cabral and Hortacsu (2010) suggest increasing returns at least in some part of the domain. Plausible interpretations are that the premium function may either exhibit increasing returns throughout its domain, or have the form of a sigmoid function (i.e., increasing returns initially and decreasing returns for large values of the percentage, cf. Figure 2). The latter is a smooth approximation of a step function, which was studied in Section 3.3.

The fact that there are such high percentages of positive ratings on eBay may also shed light on the form of the premium function. In particular, potential explanations of this phenomenon are that (1) the reputation system works well in general, (2) users tend to not report negative ratings out of fear of retaliation (Dellarocas and Wood 2008), and (3) sellers with bad records exit and possibly re-enter under a new identity. According to our model, (1) can not be a consequence of decreasing returns, and could suggest increasing returns to reputation or a special step function. We next give evidence of (3) and argue that it implies that a low percentage of positive ratings significantly decreases the expected payment to the seller.

A seller with a non-negligible percentage of negatives may re-enter the market under a new identity, especially given the high levels of competition on eBay. Kalyanam and McIntyre (2001) provide anecdotal evidence that this happens. Moreover, data from eBay show that an increase in the percentage of negatives in a seller's record translates into an increase in exit probability (Cabral and Hortacsu 2010). Such behavior may imply that a seller without an almost perfect percentage of positive ratings is better off rejoining the system than keeping his current reputation. This suggests that a low percentage of positives has a significant effect on the expected payment to the seller, and is consistent with increasing returns to reputation or a step premium function.

3.5. A Simplified Condition for Truthfulness

In this section we present an assumption that simplifies the task of checking whether the seller is always truthful; notably, this assumption is satisfied whenever the premium function is logarithmically concave. We use this assumption in the following subsection when we study dominance relationships among premium functions, as well as throughout our analysis in Section 4.

ASSUMPTION 1. *For any parameters δ and q_H, if it is not optimal for the seller to deviate from the truthful policy at $\vec{r}=\vec{1}$, then it is optimal for him to be truthful at all rating vectors.*

Assumption 1 says that if (5) holds at $\vec{r}=\vec{1}$, then (5) holds at all $\vec{r}$, and it is optimal for the seller to always be truthful. Thus, if Assumption 1 holds, then in order to check whether truthfulness is optimal for the seller, it suffices to check whether (5) holds at $\vec{r}=\vec{1}$. Using the facts that $s(\vec{1})=1$, $s_i(\vec{1})=1$ and $s_i(0,\vec{1})=1-w_i$, we can prove the following lemma.

LEMMA 3. *If Assumption 1 holds, then it is optimal for the seller to be truthful at all $\vec{r}$ if and only if*

$$(v_H - v_L)b_{\vec{w}}(1) \le q\sum_{i=0}^{\infty}\delta^{i+1}(b_{\vec{w}}(1) - b_{\vec{w}}(1-w_i)). \tag{6}$$

The following lemma shows that a large class of premium functions satisfies Assumption 1. In particular, the assumption holds for premium functions whose logarithm is concave in the seller's score.

LEMMA 4. *A logarithmically concave premium function satisfies Assumption 1.*

We observe that if the premium function $b_{\vec{w}}$ is convex, then its logarithm may or may not be concave depending on how fast the slope of the premium function increases compared to the payment itself. We note, however, that log-concavity is not a particularly restrictive assumption. For example, the functions x^n, e^{nx} are logarithmically concave for any $n > 0$.

3.6. Dominance

In this section we define a dominance relation such that if the premium function $b_{\vec{w}}$ dominates the premium function $\tilde{b}_{\vec{w}}$, then optimality of truthfulness under $\tilde{b}_{\vec{w}}$ implies optimality of truthfulness under $b_{\vec{w}}$. The dominance relationship we consider informally captures the idea that a "steeper" premium function should lead the seller to be "more truthful."

To have a fair and interesting comparison between two premium functions $b_{\vec{w}}$ and $\tilde{b}_{\vec{w}}$, they should take the same values at minimum and maximum scores; i.e., $b_{\vec{w}}(0) = \tilde{b}_{\vec{w}}(0)$ and $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$. Under this condition and Assumption 1, the dominance relation for truthfulness is $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ for all $s \in [0,1]$, as the following proposition shows.

PROPOSITION 3. *Let $b_{\vec{w}}$, $\tilde{b}_{\vec{w}}$ be premium functions such that*
(i) $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ for all $s \in [0,1]$,
(ii) $b_{\vec{w}}(0) = \tilde{b}_{\vec{w}}(0)$, $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$,
(iii) $b_{\vec{w}}$ satisfies Assumption 1.
For any $q < 1$, $\delta < 1$, and $\vec{w}$, if it is optimal for the seller to be always truthful under $\tilde{b}_{\vec{w}}$, then it is optimal for him to be always truthful under $b_{\vec{w}}$.

For example, if always advertising truthfully is optimal for a seller with parameters δ, q_H when the premium function is $\tilde{b}_{\vec{w}}(s) = s^k$ where $k > 1$, then always advertising truthfully is also optimal under $b_{\vec{w}}(s) = s^m$ for any $m \geq k$. In particular, this observation suggests that increasing the elasticity of the premium function improves incentives for the seller to be truthful. This is consistent with our observation that it is possible to have a truthful seller with convex premium functions, but not with concave premium functions.

4. Weighted Aggregation Mechanism: Design

The results of the preceding section show that for a given vector of weights $\boldsymbol{w}$ and a given premium function $b_{\boldsymbol{w}}$, under certain conditions there exists a range of $q < 1$, $\delta < 1$, for which the seller is always truthful. Can we use this insight to guide the design of the aggregation mechanism, i.e., to maximize the range of parameters for which the seller is always truthful? In this section we consider the problem of *designing* a good aggregation mechanism; in particular, we consider a setting where the system designer is considering choosing the weight vector $\boldsymbol{w}$ in a Weighted Aggregation Mechanism.

Throughout the paper, we have considered a "good" aggregation mechanism to be one that ensures sellers are truthful. Accordingly, we assume that the mechanism designer's goal is to *maximize the range of seller parameters q and δ for which truthfulness can be guaranteed.* Throughout this section we assume that the premium function satisfies Assumption 1. We begin in Section 4.1 by formulating the design problem under a range of assumptions regarding the information available to the designer regarding q and δ. We subsequently consider specific examples, corresponding to the Window Aggregation Mechanism (Section 4.2) and Weighted Aggregation Mechanism (Section 4.3), both of which were previously introduced in Section 3. Ultimately, our analysis of this design problem lends qualitative insight into the design of aggregation mechanisms; in particular, we find that aggregation mechanisms that average over a longer past history of ratings are *more likely* to incentivize patient sellers to be truthful, but *less likely* to incentivize high quality sellers to be truthful.

4.1. The Design Problem

In this section we formulate a general design problem for the Weighted Aggregation Mechanism. We assume the mechanism designer chooses the weights $\boldsymbol{w}$ from a set W; special cases are considered in the subsequent subsections.

By Lemma 3, if Assumption 1 holds, then it is optimal for the seller to be always truthful if and only if

$$b_{\boldsymbol{w}}(1) \leq \frac{q}{v_H - v_L} \sum_{i=0}^{\infty} \delta^{i+1}(b_{\boldsymbol{w}}(1) - b_{\boldsymbol{w}}(1-w_i)).$$

We conclude that *the seller is always truthful if and only if q, δ, and $\boldsymbol{w}$ jointly satisfy the following constraint:*

$$q \cdot F(\boldsymbol{w}, \delta) \geq 1 \tag{7}$$

where

$$F(\boldsymbol{w}, \delta) = \frac{1}{v_H - v_L} \frac{\sum_{i=0}^{\infty} \delta^i (b_{\boldsymbol{w}}(1) - b_{\boldsymbol{w}}(1-w_i))}{b_{\boldsymbol{w}}(1)}. \tag{8}$$

Note that $F(\boldsymbol{w}, \delta)$ is increasing in δ for any fixed $\boldsymbol{w}$.

As discussed above, our approach is to maximize the *range* of parameter values for which the seller will be truthful, given available information regarding δ and/or q. In what follows, our analysis depends on analyzing the set $\boldsymbol{w}^*(\delta)$ defined as follows for each δ:

$$\boldsymbol{w}^*(\delta) = \arg\max_{\boldsymbol{w} \in W} F(\boldsymbol{w}, \delta). \tag{9}$$

For tractability, we assume for the duration of this section that for all δ, the set $\boldsymbol{w}^*(\delta)$ is nonempty.

We now consider the optimal choice of weights depending on the available information regarding δ and q.

1. *Both δ and q are known by the mechanism designer.* In this case, the goal is to find weights $\boldsymbol{w} \in W$ such that (7) holds. Whether or not this will be possible depends on whether q and δ are large enough, given the premium functions $b_{\boldsymbol{w}}$. In particular, we can ensure the seller is always truthful if and only if $\max_{\boldsymbol{w} \in W} F(\boldsymbol{w}, \delta) \geq 1/q$; in this case any choice of weights in $\boldsymbol{w}^*(\delta)$ is optimal.

2. *The mechanism designer knows δ, but not q.* A reasonable choice of $\boldsymbol{w}$ is one which maximizes the range of values of q for which the seller will be always truthful. From (7), this implies we should maximize $F(\boldsymbol{w}, \delta)$ subject to $\boldsymbol{w} \in W$, i.e., any $\boldsymbol{w} \in \boldsymbol{w}^*(\delta)$ is an optimal choice.

3. *The mechanism designer knows q, but not δ.* A reasonable choice of $\boldsymbol{w}$ is one which maximizes the range of values of δ for which the seller will be always truthful; thus, given q, we solve:

$$\text{minimize } \delta \tag{10}$$

$$\text{subject to } q \cdot F(\boldsymbol{w}, \delta) \geq 1; \tag{11}$$

$$\boldsymbol{w} \in W. \tag{12}$$

Let $\delta^*(q)$ denote the optimal value of the preceding problem; this is the smallest value of δ such that a seller with quality q and discount factor δ can be guaranteed to be truthful under *some* weight vector. It then follows that any $\boldsymbol{w} \in \boldsymbol{w}^*(\delta^*(q))$ is an optimal choice of weights. Observe that since the constraint is increasing in q, it follows that $\delta^*(q)$ is decreasing in q; we will use this fact in our subsequent analysis to characterize the dependence of the set $\boldsymbol{w}^*(\delta^*(q))$ on q.

Of course, there is a fourth possibility: the mechanism designer may know *neither* δ nor q. In this scenario, the mechanism designer will typically need to consider whether it is preferable to make sellers truthful over a greater range of δ, or a greater range of q. We delineate such a tradeoff in subsequent subsections.

A critical theme emerges from the preceding discussion: *regardless* of the information available to the mechanism designer, solving (9) is an important step in finding the best weight vector. We apply this insight in the following subsections to study both the Window Aggregation Mechanism and the more general Weighted Aggregation Mechanism.

4.2. Window Aggregation Mechanism

The Window Aggregation Mechanism (introduced in Example 3) is characterized by the window size, which we denote by T. As previously discussed, it is a special case of a Weighted Aggregation Mechanism, where $w_0 = w_1 = \cdots = w_{T-1} = 1/T$, and all other weights are zero; the feasible set W consists of all weight vectors of this form for window sizes $T \geq 1$. Thus the window size is the only design choice.

In this section, we apply the methodology of the Section 4.1 to determine an optimal window. Our main result is that an optimal window size is *increasing* in the discount factor δ, and *decreasing* in the quality q. Thus a longer window is *more likely* to make patient sellers always truthful, but *less likely* to make high quality sellers always truthful.

In a slight abuse of notation, we write $b_T(\cdot)$ for the premium function to indicate that the buyers' behavior may depend on window size T. Throughout the section we implicitly assume that the family of premium functions b_T is known by the mechanism designer. We further assume that $b_T(\cdot)$ is strictly convex and satisfies Assumption 1 for each T; as noted in Lemma 4, this Assumption is satisfied by logarithmically concave premium functions.

In the case of a Window Aggregation Mechanism, by a slight abuse of notation, we redefine F from (8) as:

$$F(T, \delta) = \frac{1}{v_H - v_L} \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \sum_{i=1}^{T} \delta^i.$$

Calculation of the optimal weight vector in (9) then reduces to calculating the optimal window size via the following optimization:

$$T^*(\delta) = \arg\max_{T \geq 1} F(T, \delta). \tag{13}$$

As before, we assume for the duration of this section that for all δ, the set $T^*(\delta)$ is nonempty.

The following proposition characterizes the behavior of $T^*(\delta)$.

PROPOSITION 4. *If Assumption 1 holds, then $T^*(\delta)$ is increasing in δ in the following sense: for $\delta \geq \delta'$,*
(i) $\max\{T : T \in T^*(\delta)\} \geq \max\{T : T \in T^*(\delta')\}$*; and*
(ii) $\min\{T : T \in T^*(\delta)\} \geq \min\{T : T \in T^*(\delta')\}$.

Surprisingly, note that this result holds *regardless* of the dependence of b_T on T.

Using this proposition, we now consider the optimal choice of window depending on the available information regarding δ and q, as in the preceding subsection.

1. *Both δ and q are known by the mechanism designer.* As in Section 4.1, any window size in $T^*(\delta)$ is optimal, and whether the seller can be made truthful depends on whether q and δ are sufficiently large to ensure that (7) holds. In this case any window size in $T^*(\delta)$ is optimal.

2. *The mechanism designer knows δ, but not q.* As in Section 4.1, if we wish to maximize the range of values of q for which the seller will be truthful, then any $T \in T^*(\delta)$ is an optimal choice. Note that in this case, from Proposition 4, *the set of optimal windows is increasing in δ.* This is an intuitive result, since sellers with larger δ are more patient, and thus a aggregation system with longer memory can successfully couple current behavior with distant future payoffs.

3. *The mechanism designer knows q, but not δ.* As in Section 4.1, let $\delta^*(q)$ denote an optimal solution to (10)-(12), where again the set W consists of all possible weight vectors corresponding to Window Aggregation Mechanisms. It then follows any $T \in T^*(\delta^*(q))$ is an optimal choice of window size. Recall that $\delta^*(q)$ is decreasing in q. From Proposition 4, we conclude *the set of optimal windows $T^*(\delta^*(q))$ is decreasing in q.* This is an intuitive result, since as q increases, it is possible to make less patient sellers truthful, and for such sellers a smaller window size is more appropriate.

We can now see the tradeoff discussed at the beginning of the section: informally, increasing the window size is more likely to make patient sellers (those with high δ) truthful. On the other hand, it is less likely to make high quality sellers (those with high q) truthful. When q is high and the window is large, the seller is likely to have a high score regardless of what actions he takes when he receives a low value item, because most items are high quality. This makes a smaller window more desirable, because it magnifies the impact of the seller's actions in those periods where he has a low value item for sale.

Finally, we observe that this tradeoff is also faced by a mechanism designer who knows *neither* q nor δ: a choice must be made regarding the incentives provided to patient sellers and those provided to high quality sellers.

4.3. Weighted Aggregation Mechanism

In the previous section we studied the problem of finding the best mechanism in the class of Window Aggregation Mechanisms. In this section we consider the optimization within the general class of Weighted Aggregation Mechanisms; i.e., we let W denote the set of all weight vectors such that $\sum_i w_i = 1$, and $w_i \geq 0$ for all i, with $w_0 \geq w_1 \geq w_2 \cdots$.

In this section we make the following assumption.

ASSUMPTION 2. *The premium function is $b(\cdot)$, and does not depend on the weight vector $\boldsymbol{w}$. Further, it is strictly convex and logarithmically concave.*

Note that we assume that the payment depends on the weights of the aggregation mechanism only through the score. Since Assumption 2 implies Assumption 1, we conclude (as before) by Lemma 3 that it is optimal for the seller to be always truthful if and only if condition (7) holds.

Applying the formulation of Section 4.1, it can be easily shown that finding the optimal weight vector $\boldsymbol{w}^*(\delta)$ in (9) is equivalent to the following optimization problem[4].

$$\text{minimize} \sum_{i=0}^{\infty} \delta^i \cdot b(1 - w_i) \tag{14}$$

$$\text{subject to} \sum_{i=0}^{\infty} w_i = 1; \tag{15}$$

$$w_i \geq 0, \forall i. \tag{16}$$

By Assumption 2, the premium function is strictly convex and thus (14)-(16) has a unique solution.

For the general Weighted Aggregation Mechanism it may seem plausible that all past transactions of the seller should be assigned strictly positive weights in order to best incentivize truthfulness. The following lemma shows that this is *not* the case under Assumption 2: it is optimal to only include a *finite* number of ratings in the seller's score. In particular, we note that this result implies that the Exponential Moving Average Mechanism introduced in Example 4 never arises as an optimal weight vector.

LEMMA 5. *If the premium function satisfies Assumption 2 and $\delta \in (0,1)$, then only finitely many weights are positive in $\boldsymbol{w}^*(\delta)$.*

In light of Lemma 5, let $N^*(\delta)$ be the number of strictly positive weights at the optimal solution; i.e., $w_i^*(\delta) > 0$ for $i = 0, ..., N^*(\delta) - 1$, and $w_i^*(\delta) = 0$ for $i \geq N^*(\delta)$. The following proposition characterizes $\boldsymbol{w}^*(\delta)$ and $N^*(\delta)$.

PROPOSITION 5. *If Assumption 2 holds, then*

[4] It is straightforward to show that at the optimal solution the weights will be nonincreasing, and thus we do not need to explicitly include the constraints $w_i \geq w_j$ for $i < j$.

(i) $w_i^*(\delta)$ *is strictly decreasing in* i *for* $i < N^*(\delta)$;
(ii) For sufficiently small i, $w_i^*(\delta)$ *is decreasing in* δ; *and*
(iii) $N^*(\delta)$ *is nondecreasing in* δ.

According to (i), the optimal weights are strictly decreasing in i for $i < N^*(\delta)$, and thus each strictly positive weight is different at the optimal solution. This also implies that the Window Aggregation Mechanism, which has weights $w_i = 1/T$ for $i = 0, ..., T-1$, does not arise as an optimal weight vector. The preceding observation is particularly interesting when one considers that nearly all online marketplaces that use mechanisms which only weight recent ratings, tend to use a Window Aggregation Mechanism to do so; for example, eBay's current mechanism weights all ratings equally over the past 12 months of feedback. Our insight suggests that a more robust mechanism might be one that weights more recent ratings even more heavily than older ratings.

Following Section 4.1, we now consider the optimal choice of weights depending on the available information regarding δ and q.

1. *Both* δ *and* q *are known by the mechanism designer.* As in Section 4.1, the vector of weights $\boldsymbol{w}^*(\delta)$ is optimal is this case, and whether or not it will be possible to make the seller always truthful depends on whether q and δ are large enough to ensure (7) is satisfied.

2. *The mechanism designer knows* δ, *but not* q. As in Section 4.1, the weights $\boldsymbol{w}^*(\delta)$ maximize the range of values of q for which the seller will be truthful. From Proposition 5 (ii), as δ increases, the most recent transactions are weighted less. Informally, this happens because the value to the seller of the distant future increases relative to the near future. This effect implies that the aggregation mechanism should weight distant past ratings more heavily relative to recent ratings when δ increases. Moreover, Proposition 5 (iii) states that *the optimal number of strictly positive weights is nondecreasing in* δ. As in the case of the Window Aggregation Mechanism, the intuition for this is that sellers with higher discount factors care more about the future, and so truthfulness is better incentivized when information from more transactions is included.

3. *The mechanism designer knows* q, *but not* δ. As in Section 4.1, let $\delta^*(q)$ be the solution to (10)-(12). Then the weights $\boldsymbol{w}^*(\delta^*(q))$ maximize the range of values of δ for which the seller will be truthful. Recall that $\delta^*(q)$ is decreasing in q; thus, from Proposition 5 (ii), as q increases, the most recent transactions are weighted more. Moreover, from Proposition 5 (iii), *the optimal number of strictly positive weights is decreasing in* q. As before, this is an intuitive result, since as q increases, it is possible to make less patient sellers truthful, and such sellers value the future less.

We observe that the dependence of the number of strictly positive weights on δ and q is similar to the corresponding results on the optimal Window Aggregation Mechanism, i.e., that the optimal window size is nondecreasing in δ, and nonincreasing in q. Again, this tradeoff directly affects a mechanism designer who knows neither δ nor q.

5. Non-Binary Values and Ratings

Throughout the paper we assumed that the value of the item is either high or low, and that ratings are either good or bad (binary rating system). In this section we discuss how the modeling and the results shown can be generalized when we relax this assumption. Thus, our results are robust even if the seller has potentially many "levels" of dishonesty possible.

We assume that the possible values of the item are $\Omega = \{0, 1/Q, 2/Q, \ldots, 1\}$ for some integer $Q \geq 1$, instead of only $\{\text{low}, \text{high}\}$. Let p_v be the probability that the seller has an item of value v for sale, so $\sum_{v\in\Omega} p_v = 1$. Note that we allow the possibility $p_v = 0$ for all but one value in Ω; in this case, the seller has the same item in every period, but may be tempted to exaggerate its value in his advertisement.

Let $v_a \cdot b_{\vec{w}}(s)$ be the expected payment to the seller when he advertises an item as having value v_a and his score is s. Thus the highest quality item is worth $b_{\vec{w}}(s)$, and the expected payment to the seller is scaled in proportion to the advertised quality. We assume that $b_{\vec{w}}(s)$ is an increasing function of s.

Further, we assume that when the advertised value is v_a and the true value v, then the seller receives rating $r_i = 1-(v_a - v)^+$, where $x^+ \equiv \max(x, 0)$ denotes the positive part of x. That is, the buyer "penalizes" the seller by the difference between the advertised and the true value, whenever this difference is positive.[5] Since the payment $v_a \cdot b_{\vec{w}}(s)$ is increasing in v_a and s, the seller will only ever consider exaggerating the value of his advertisement, and we can safely assume that $r_i = 1-(v_a - v)$. We consider the Weighted Aggregation Mechanism: the seller's score that the buyers see is $s(\vec{r}) = \sum_j w_j r_j$, where $w_i \geq 0$, and $w_0 \geq w_1 \geq ...$, with $\sum_i w_i = 1$.

Let $s_i(\vec{r})$ be the reputation score of the seller in i periods if in this period the rating is f, and the seller is truthful in all other periods. Our main insight is the following lemma.

LEMMA 6. *If $b_{\vec{w}}$ is convex, then it is optimal for the seller to be always truthful if and only if*

$$b_{\vec{w}}(s(\vec{r})) \leq \left(\sum_{v' \in \Omega} p_{v'} \cdot v'\right) \sum_{i=0}^{\infty} \delta^{i+1} \cdot (b_{\vec{w}}(s_i(1, \vec{r})) - b_{\vec{w}}(s_i(0, \vec{r}))) \tag{17}$$

for all rating vectors $\vec{r}$.

In words, Lemma 6 says that if the premium function is convex, then the seller is more tempted to post an advertisement that is significantly higher than the true value. The intuition for this is that a small lie now is associated with a small gain now, but a relatively large reduction in future payments (because of convexity). If the payment is convex and the seller considers exaggerating the value of the item in his description, he is better off significantly exaggerating it. Then, the reduction in future payments is not that large relative to the current gain. When the premium function is concave, the reverse occurs: small lies are more beneficial for the seller. This is because under concavity a small lie only imposes a small reduction in future payments (relative to current gains).

Inspired by the preceding result, define:

$$q \equiv \sum_{v' \in \Omega} p_{v'} \cdot v',$$

to be the seller's *quality*. As a result, *condition (17) is equivalent to (5) with q_H replaced by q*. From this observation, we immediately obtain analogs of all the results of the paper.

If we assume that $b_{\vec{w}}$ is *strictly convex and logarithmically concave*, then we get analogs of the design results of Section 4. In particular, the optimal weights are strictly decreasing whenever positive, and the number of strictly positive weights is finite. Thus, neither the Window Aggregation Mechanism nor Exponential Moving Average Mechanism are optimal. Finally, the tradeoff discussed for the binary case still holds: informally, increasing the number of strictly positive weights is more likely to make patient sellers (those with high δ) truthful, while it is less likely to make high quality sellers (those with high q) truthful.

6. Related Work

In this section we discuss related literature on seller reputation that can be applied to an online market setting. There are two main approaches. In the adverse selection approach the seller has a hidden type which buyers are trying to learn. Alternatively, reputation can only be used to

[5] The same results hold if $r_i = 1 - |v_a - v|$ instead. In this case the buyer "penalizes" the seller by the absolute difference between the advertised and true value. However, it seems more realistic and fair to only "penalize" the seller when he exaggerates the value of the item in his advertisement.

incentivize the seller to behave well. A standard assumption is that a long-lived seller interacts with short-lived buyers.

In the adverse selection approach, the seller's reputation is studied in a Bayesian setting. The seller is assumed to have a type which buyers are trying to learn from his past ratings. In this setting, the seller's reputation is the belief that buyers have about his type after observing the available information on his past behavior.

A common assumption in the adverse selection literature on reputation is that with some (small) probability the seller's type is such that he always plays an action that promotes trust (Mailath and Samuelson 2006). If the buyers have access to all past ratings of the seller, they are eventually going to discover his type (Cripps et al. 2004), which means that reputations are not sustainable; this is similar to our insight regarding the Unweighted Aggregation Mechanism (cf. Section 2). In this setting, information censoring can result in sustainable reputation (Ekmekci 2010). Also, if buyers have to pay to discover the seller's past behavior, then equilibrium behavior is cyclical: the seller builds his reputation up only to exploit it (Liu 2006).

In the setting we study, the seller does not have a hidden type and the seller's score is only used to incentivize good behavior by the seller. The objective of the reputation mechanism is to induce sellers to behave in a way that promotes trust. This approach has also been taken by Dellarocas (2005) and Fan et al. (2005).

Dellarocas (2005) studies a setting where the seller has two possible effort levels which buyers observe imperfectly. He shows that there is no equilibrium where the seller always exerts high effort, and that eBay's simple mechanism is capable of inducing the maximum theoretical efficiency. In this paper we take a non-equilibrium approach. We consider the best response of the seller to a *fixed* Markov strategy of the buyers; that is, a fixed payment function which only depends on the information available to buyers. Our work takes the point of view that aggregation mechanisms calculate sufficient statistics of the past, and act as the "state" in the interaction between buyers and sellers. We believe that this nonequilibrium approach is reasonable in practice, because of the coordination required for "convergence" to equilibrium. The large and dynamic set of participants in the major online markets makes the rationality, knowledge, and coordination required for equilibrium improbable.

Fan et al. (2005) consider a similar setting, where the seller exerts high or low effort, but take a non-equilibrium approach and assume that the seller has a belief over the average bidder behavior. For a specific behavior of bidders it is shown that two simple aggregation mechanisms are inadequate and propose exponential smoothing instead. We consider a general class of payment functions and aggregation mechanisms. Moreover, in our model, the seller decides whether to advertise truthfully and not whether to choose a high or low cost action. Thus, in this paper the goal of an aggregation mechanism is to incentivize truthful behavior by the seller.

Although we do not consider incentives for buyers to leave honest feedback in this paper, another line of research considers how truthful feedback can be elicited. In online markets agents may undertake fake transactions in order to enhance their reputation. This can be avoided if a specific relation between the reputation premium and the transaction cost holds (Bhattacharjee and Goel 2005). Alternatively, even if fake transactions can not be undertaken, buyers may not leave honest feedback after a transaction. Miller et al. (2005) devise a scoring system that induces honest reporting of feedback. In this paper, we assume that buyers always leave truthful feedback in order to focus on the seller's decisions.

We conclude by noting that a number of papers have empirically studied the effect that the seller's score has on the average payment he receives. Resnick and Zeckhauser (2002) provide a survey. Some sample studies include data about eBay auctions for coins (Lucking-Reiley et al. 2007); Palm Pilots (Kalyanam and McIntyre 2001); Pentium III processors (Houser and Wooders 2006); collectible coins, Thinkpads and Beanie babies (Cabral and Hortacsu 2010); and postcards (Resnick et al. 2006). The last of these papers (Resnick et al. 2006) conducts a controlled experiment on eBay to study the price premium of an experienced seller. Finally, Ghose et al. (2005) and Pavlou and Dimoka (2006) study the effect of different dimensions of a seller's reputation on pricing power by considering text comments on Amazon and eBay respectively.

7. Proofs

Proof of Proposition 1: We will show that there exists a sufficiently large N such that if the seller is truthful up to that time, then he is better off advertising a low value item as a high value item at time N. We assume that the seller is truthful up to time N and so his score is $(s_P, s_T) = (N, N)$. Substituting in (2), the seller will advertise a low value item truthfully at reputation (N, N) only if

$$b(N,N) \leq \delta(V(N+1,N+1) - V(N,N+1)). \tag{18}$$

We wish to upper bound $V(N+1,N+1) - V(N,N+1)$. Having one more positive rating only affects the seller's payoff when he advertises a high value item; in those periods payoff is determined by the premium function $b(\cdot)$. Let $\alpha_M = b(M,M) - b(M-1,M)$. By (3), $\alpha_M \to 0$ as $M \to \infty$. Choose $k^*(N) \in \arg\max_{k\geq 0} \alpha_{N+k+1}$. Then,

$$V(N+1,N+1) - V(N,N+1) \leq \sum_{k=0}^{\infty} \delta^k \alpha_{N+k+1} \leq \frac{1}{1-\delta} \alpha_{N+k^*(N)+1} \to 0 \text{ as } N \to \infty.$$

Thus, for every $\varepsilon > 0$ there exists N_1 such that $V(N+1,N+1) - V(N,N+1) < \varepsilon$ for all $N \geq N_1$. Since $b_N(1)$ is assumed to be bounded away from zero as $N \to \infty$, there exist $\varepsilon > 0$ and N_2 such that $b_N(1) > \varepsilon$ for all $N \geq N_2$. Thus, there exists N for which (18) is not satisfied, which implies that there exists an s_T at which it is optimal for the seller to falsely advertise a low value item.

■

Proof of Lemma 2:

We first prove (i). Fix a vector of ratings $\vec{r}$. Note that if $b_{\vec{w}}(s(\vec{r})) = 0$, then the condition trivially holds. We thus assume that $b_{\vec{w}}(s(\vec{r})) > 0$ in what follows.

We first observe that if $w_i > 0$, then $s_i(1,\vec{r}) > w_i$. In particular, we could only have $s_i(1,\vec{r}) = w_i > 0$ if $i = 0$ and $w_0 = 1$, which contradicts the assumption that $w_0 < 1$.

If $w_i > 0$, then

$$b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})) = b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(1,\vec{r}) - w_i) > w_i \cdot \frac{b_{\vec{w}}(s_i(1,\vec{r}))}{s_i(1,\vec{r})} \geq w_i \cdot b_{\vec{w}}(s(\vec{r})).$$

The equality holds because $s_i(0,\vec{r}) = s_i(1,\vec{r}) - w_i$. The first inequality follows because $b_{\vec{w}}$ is strictly convex and $b_{\vec{w}}(0) = 0$, and the second because $s(\vec{r}) \leq s_i(1,\vec{r}) \leq 1$. We conclude that $b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})) > w_i \cdot b_{\vec{w}}(s(\vec{r}))$ for all i with $w_i > 0$. Summing over all i and using the fact that $\sum_i w_i = 1$, we show (i).

We now show (ii). By concavity of $b_{\vec{w}}$,

$$\sum_{i=0}^{\infty} (b_{\vec{w}}(1) - b_{\vec{w}}(1-w_i)) \leq \sum_{i=0}^{\infty} w_i b_{\vec{w}}(1) = b_{\vec{w}}(1),$$

since $b_{\vec{w}}(0) \geq 0$, and $\sum_i w_i = 1$. ■

Proof of Proposition 2: We note that if $s(\vec{r}) < s^*$, then the seller has no reason to deviate from the truthful policy. We thus consider what happens for $s(\vec{r}) \geq s^*$.

We first show that if $s^* \leq 1 - w_1$, then the seller is better off deviating from the truthful policy at $s(\vec{r}) = 1$. If he receives a low value item and does not deviate, his infinite horizon expected payoff is

$v_H q_H \sum_{i=1}^{\infty} \delta^i b_{\vec{w}}(1)$, while if he does deviate his expected payoff is at least $v_H(1 + q_H \sum_{i=2}^{\infty} \delta^i) b_{\vec{w}}(1)$, which is strictly greater for any $\delta < 1$. This shows (i).

We now show (ii). If $s^* > 1 - w_1$, then $w_0 \geq w_1 > 1 - s^*$. We will show that (5) is satisfied for any vector of ratings $\vec{r}$ with $s(\vec{r}) \geq s^*$. We note that $b_{\vec{w}}(s(\vec{r})) = b_{\vec{w}}(s(1,\vec{r})) = b_{\vec{w}}(1)$. Moreover, $s(0,\vec{r}) = s_i(1,\vec{r}) - w_i$, so $s_0(0,\vec{r}), s_1(0,\vec{r}) < s^*$ and $b_{\vec{w}}(s_0(0,\vec{r})) = b_{\vec{w}}(s_1(0,\vec{r})) = 0$. Thus,

$$\delta q_H \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) \geq \delta q_H \sum_{i=0}^{1} \delta^i (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) = q_H(\delta + \delta^2) b_{\vec{w}}(1),$$

and thus if $q_H(\delta + \delta^2) \geq 1$, the seller does not deviate.

■

Proof of Lemma 4:

By Lemma 1, it is optimal for the seller to always be truthful if $q \geq q^*(b_{\vec{w}}, \delta, \vec{r})$ for all $\vec{r}$, where $q^*(b_{\vec{w}}, \delta, \vec{r}) \equiv (v_H - v_L) b_{\vec{w}}(s(\vec{r})) / \sum_{i=0}^{\infty} \delta^{i+1} (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$. If it is not optimal for the seller to deviate from the truthful policy at maximum reputation, then $q \geq q^*(b_{\vec{w}}, \delta, \vec{1})$.

Let $b_{\vec{w}}$ be logarithmically concave. To show the statement of the proposition, it suffices to show that $q^*(b_{\vec{w}}, \delta, \vec{1}) \geq q^*(b_{\vec{w}}, \delta, \vec{r})$ for all $\vec{r}$, i.e., if for all $\vec{r}$, $b_{\vec{w}}(1) \cdot \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) \geq b_{\vec{w}}(s(\vec{r})) \cdot \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))$.

It suffices to show that

$$b_{\vec{w}}(1) \cdot (b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^0(\vec{r}))) \geq b_{\vec{w}}(s(\vec{r})) \cdot (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)). \tag{19}$$

To conclude the proof we show that (19) is satisfied by a logarithmically concave function.

Let $w > 0$ and $x \in [w, 1]$. Since $b_{\vec{w}}$ is logarithmically concave, then $\log(b_{\vec{w}}(x)) - \log(b_{\vec{w}}(x - w))$ is nonincreasing in x, which implies that $b_{\vec{w}}(x - w)/b_{\vec{w}}(x)$ is nondecreasing in x. Thus, $b_{\vec{w}}(1 - w_i)/b_{\vec{w}}(1) \geq b_{\vec{w}}(s_i(0,\vec{r}))/b_{\vec{w}}(s_i(1,\vec{r}))$ and $(b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))/b_{\vec{w}}(1) \leq (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))/b_{\vec{w}}(s_i(1,\vec{r}))$. Using the fact that $b_{\vec{w}}(s_i(1,\vec{r})) \geq b_{\vec{w}}(s_i(0,\vec{r}))$, this implies that (19) holds and concludes the proof. ■

Proof of Proposition 3: If the seller is always truthful under $\tilde{b}_{\vec{w}}$, it must be that $(v_H - v_L)\tilde{b}_{\vec{w}}(1) \leq \delta q \sum_{i=0}^{\infty} \delta^i (\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i))$, since in particular he is truthful when his score is equal to 1. Since $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ for all s, and $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$, we have $\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i) \leq b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)$ for all w_i. Thus,

$$(v_H - v_L) b_{\vec{w}}(1) = (v_H - v_L)\tilde{b}_{\vec{w}}(1) \leq \delta q \sum_{i=0}^{\infty} \delta^i (\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i)) \leq \delta q \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)).$$

Since $b_{\vec{w}}$ satisfies Assumption 1 and condition (6), the optimal policy of the seller is to always advertise truthfully under $b_{\vec{w}}$ (by Lemma 3). ■

Proof of Proposition 4: We will show that $\log(F(T,\delta))$ satisfies increasing differences. Let $T' \geq T$.

$$\log(F(T',\delta)) - \log(F(T,\delta)) = \log\left(\frac{1 - \delta^{T'}}{1 - \delta^T}\right) + t(T,T'),$$

where $t(T,T')$ does not depend on δ. Thus, to show that $\log(F(T,\delta))$ has increasing differences in (T,δ) it suffices to show that $(1 - \delta^{T'})/(1 - \delta^T)$ is increasing in δ. The first derivative with respect to δ is positive if and only if

$$\frac{T \cdot \delta^{T-1}}{1 - \delta^T} \geq \frac{T' \cdot \delta^{T'-1}}{1 - \delta^{T'}}.$$

Since $T' \geq T$ it suffices to show that $r(x) \equiv (x \cdot \delta^{x-1})/(1 - \delta^x)$ is decreasing. We proceed by differentiating r:

$$r'(x) = \frac{\delta^{x-1}}{(1 - \delta^x)^2}(1 - \delta^x - x \ln(1/\delta)).$$

To complete the proof, we show that $\delta^T + T\ln(1/\delta) > 1$ holds for $T \geq 1$, $\delta \in (0,1)$. First note that $\delta^T + T\ln(1/\delta)$ is increasing in T, since

$$\frac{\partial(\delta^T + T\ln(1/\delta))}{\partial T} = \ln(1/\delta)\cdot(1-\delta^T) > 0.$$

So it suffices to show that $\hat{g}(\delta) \equiv \delta + \ln(1/\delta) > 1$. g is strictly decreasing in $(0,1)$, because

$$\hat{g}'(\delta) = 1 + \frac{-1/\delta^2}{1/\delta} = \frac{\delta - 1}{\delta} < 0,$$

and $\hat{g}(1) = 1$. So, $\hat{g}(\delta) > 1$ for $\delta \in (0,1)$.

This proves that $\log(F(T,\delta))$ has increasing differences in (T,δ); the result follows by applying Topkis' Theorem (Topkis 1998). ■

Proof of Lemma 5:

We first show that if b is logarithmically concave, then $b'(1) < \infty$. Suppose not, and let $f(x) = \log(b(x))$. Then f is concave and $f'(x) \to \infty$ as $x \to 1$. However, this is not possible since if f is concave, then f' is decreasing on $[0,1]$.

We next show that if $b'(1) < \infty$, then at the optimal solution only a finite number of weights are positive. Consider the vector $\vec{u}^*(\delta)$, which is the solution of (14)-(16). The optimality conditions are:

$$b'(1-u_i^*) = \frac{\lambda - \mu_i}{\delta^i};\; u_i^* \cdot \mu_i = 0, \text{ for all } i,$$

where λ and μ_i are the Lagrange multipliers of constraint (15) and (16) respectively. Suppose $u_i^* > 0$ for all i. Then $b'(1-u_i^*) = \lambda/\delta^i$ for some $\lambda > 0$. But $b'(1-u_i^*) < b'(1) < \infty$ for all i, while $\lambda/\delta^i \to \infty$ as $i \to \infty$, which is a contradiction. Thus it must be that $u_i^* > 0$ for only a finite number of i's, i.e., $N^*(\delta) < \infty$.

■

Proof of Proposition 5: We first consider $\vec{u}^*$, i.e., the solution of (14)-(16). Let λ be the Lagrange multiplier of (15). The optimality conditions imply that for $i = 0,...,N^*(\delta)-1$, $u_i^*(\delta) = 1 - b'^{-1}(\lambda/\delta^i)$. Since b is strictly convex, both b' and b'^{-1} are strictly increasing. Thus the optimal weights are strictly decreasing for $i = 0,...,N^*(\delta)-1$. For $\vec{v}^*$, we observe that the optimal weights are strictly decreasing for $i = 0,...,N^*(\delta^*(q))-1$ (according to the proof of Lemma 5). This proves (i).

Substituting the formula for $u_i^*(\delta)$ in (15),

$$\sum_{i=0}^{N^*(\delta)-1} b'^{-1}(\lambda/\delta^i) = N^*(\delta) - 1. \tag{20}$$

In order to find the optimal weights, it suffices to find $\lambda, N^*(\delta)$ such that equation (20) is satisfied and

$$\frac{\lambda}{\delta^{N^*(\delta)-1}} < b'(1) \leq \frac{\lambda}{\delta^{N^*(\delta)}}, \tag{21}$$

so that $u_{N^*-1}^*(\delta) > 0$ and $u_{N^*}^*(\delta) = 0$.

Suppose $\delta_1 < \delta_2$. Let N_1 and N_2 be the corresponding optimal numbers of strictly positive weights, and λ_1 and λ_2 be the corresponding Lagrange multipliers. By (20),

$$\sum_{i=0}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) = N_1 - 1;$$

$$\sum_{i=0}^{N_2-1} b'^{-1}(\lambda_2/\delta_2^i) = N_2 - 1.$$

We will show that $N_1 \le N_2$. Suppose not. Then, $N_1 > N_2$, and subtracting the latter from the former equation,

$$\sum_{i=N_2}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) + \sum_{i=0}^{N_2-1} \left(b'^{-1}(\lambda_1/\delta_1^i) - b'^{-1}(\lambda_2/\delta_2^i)\right) = N_1 - N_2.$$

Since $b'^{-1}(\lambda_1/\delta_1^i) < 1$ for $i \in \{N_2, ..., N_1 - 1\}$, it follows that $\sum_{i=N_2}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) < N_1 - N_2$. Thus the previous equality can only hold if $b'^{-1}(\lambda_1/\delta_1^i) > b'^{-1}(\lambda_2/\delta_2^i)$ for some $i \in \{0, 1, ..., N_2 - 1\}$. Since b is strictly convex, this implies that $\lambda_1/\delta_1^i > \lambda_2/\delta_2^i$ for some $i \in \{0, 1, ..., N_2 - 1\}$. We observe that if this is the case for i, it is also the case for $i+1$ (since $1/\delta_1 > 1/\delta_2$). We conclude that it must hold for $i = N_2 - 1$, i.e., $\lambda_1/\delta_1^{N_2-1} > \lambda_2/\delta_2^{N_2-1}$. Thus $b'^{-1}(\lambda_1/\delta_1^{N_2}) > b'^{-1}(\lambda_2/\delta_2^{N_2})$, which is a contradiction, since $N_1 > N_2$ implies $b'^{-1}(\lambda_1/\delta_1^{N_2}) < 1$ and $b'^{-1}(\lambda_2/\delta_2^{N_2}) \ge 1$ (by (21)). Thus, if $\delta_1 < \delta_2$, it must be $N_1 \le N_2$. This proves that $N^*(\delta)$ is increasing in δ.

To show that $M^*(q)$ is decreasing in q, we observe that $\delta^*(q)$ (defined in the proof of Lemma 5) is decreasing in q, and that $M^*(q) \equiv N^*(\delta^*(q))$. This concludes the proof of (ii).

■

Proof of Lemma 6: When the seller's rating vector is $\vec{r}$ and his rating in the current period is f, we denote the seller's rating vector in the next period by $u^f(\vec{r})$. In particular, $u^f(r_0, r_1, ...) = (f, r_0, r_1, ...)$. The Bellman equation for the discounted infinite horizon expected payoff of the seller is then given by:

$$V(\vec{r}) = \sum_{v \in \Omega} p_v \max_{v_a \ge v, v_a \in \Omega} \left\{ v_a b_{\vec{w}}(s(\vec{r})) + \delta \cdot V\left(u^{1-(v_a - v)}(\vec{r})\right) \right\}.$$

Similar to the analysis in Section 3.1, it is optimal for the seller to always be truthful if and only if

$$(v_a - v) \cdot b_{\vec{w}}(s(\vec{r})) \le \left(\sum_{v' \in \Omega} p_{v'} \cdot v' \right) \sum_{i=0}^{\infty} \delta^{i+1} \cdot \left(b_{\vec{w}}(s_i^1(\vec{r})) - b\left(s_i^{1-(v_a - v)}(\vec{r})\right) \right) \tag{22}$$

for all rating vectors $\vec{r}$ and all $v_a, v \in \Omega$ with $v_a > v$.

We observe that (22) depends only on the difference $v_a - v$ and not on the specific values v_a or v. Moreover, if $b_{\vec{w}}$ is convex, since $(b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^{1-d}(\vec{r})))/d = (b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^1(\vec{r}) - w_i d))/d$ is decreasing in d, we conclude that if (22) is satisfied for v_a, v, then it is also satisfied for v'_a, v' with $v'_a - v' \le v_a - v$. This establishes the lemma. ■

Acknowledgments. This work was partially supported by the National Science Foundation under grant 0428868. We also gratefully acknowledge helpful conversations with Ashish Goel, Richard Zeckhauser, and Benjamin Van Roy.

References

Akerlof, George A. 1970. The market for lemons: Quality uncertainty and the market mechanism. *Quart. J. of Econom.* **84**(3) 488–500.

Anderson, Norman H. 1981. *Foundations of Information Integration Theory*. New York: Academic Press.

Aperjis, Christina, Ramesh Johari. 2010. Optimal windows for aggregating ratings in electronic marketplaces. *Management Science (to appear)* .

Bhattacharjee, Rajat, Ashish Goel. 2005. Avoiding ballot stuffing in eBay-like reputation systems. *P2PECON*. 133–137.

Cabral, Luis, Ali Hortacsu. 2010. Dynamics of seller reputation: Theory and evidence from eBay. *J. of Industr. Econom. (to appear)* .

Cripps, Martin, George J. Mailath, Larry Samuelson. 2004. Imperfect monitoring and impermanent reputations. *Econometrica* **72** 407–432.

Dellarocas, Chrysanthos. 2005. Reputation mechanism design in online trading environments with pure moral hazard. *Inform. Systems Res.* **16**(2) 209–230.

Dellarocas, Chrysanthos N., Charles A. Wood. 2008. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Sci.* **54**(3) 460–476.

Ekmekci, Mehmet. 2010. Sustainable reputations with rating systems. *Journal of Economic Theory (to appear)* .

Fan, Ming, Yong Tan, Andrew B. Whinston. 2005. Evaluation and design of online cooperative feedback mechanisms for reputation management. *IEEE Trans. on Knowl. and Data Eng.* **17**(2) 244–254.

Ghose, Anindya, Panagiotis G. Ipeirotis, Arun Sundararajan. 2005. Reputation premiums in electronic peer-to-peer markets: Analyzing textual feedback and network structure. *P2PECON*. 150–154.

Hogarth, Robin M., Hillel J. Einhorn. 1992. Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology* **24**(1) 1–55.

Houser, Daniel, John Wooders. 2006. Reputation in auctions: Theory, and evidence from eBay. *J. Econom. & Management Str.* **15**(2) 353–369.

Kalyanam, Kirthi, Shelby McIntyre. 2001. Return on reputation in online auction markets. Working Paper, Santa Clara University.

Kashima, Yoshihisa, Andrew R. Z. Kerekes. 1994. A distributed memory model of averaging phenomena in person impression formation. *Journal of Experimental Social Psychology* **30**(5) 407 – 455.

Liu, Qignmin. 2006. Information acquisition and reputation dynamics. SIEPR Discussion Paper 06-30, Stanford University.

Livingston, Jeffrey A. 2005. How valuable is a good reputation? A sample selection model of internet auctions. *The Review of Economics and Statistics* **87**(3) 453–465.

Lucking-Reiley, David, Doug Bryan, Naghi Prasad, Daniel Reeves. 2007. Pennies from eBay: The determinants of price in online auctions. *J. Industrial Econom.* **55**(2) 223–233.

Mailath, George J., Larry Samuelson. 2006. *Repeated Games and Reputations*. Oxford University Press.

Miller, Nolan, Paul Resnick, Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Sci.* **51**(9) 1359–1373.

Obloj, Tomasz, Laurence Capron. 2008. Whether and when sellers can extract value from reputation advantage: The case of internet auctions. *Atlanta Competitive Advantage Conference Paper*.

Pavlou, Paul A., Angelika Dimoka. 2006. The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation. *Inform. Systems Res.* **17**(4) 392–414.

Resnick, Paul, Richard Zeckhauser. 2002. *The Economics of the Internet and E-Commerce*, vol. 11, chap. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. Elsevier Science Ltd., 127–157.

Resnick, Paul, Richard Zeckhauser, John Swanson, Kate Lockwood. 2006. The value of reputation on eBay: A controlled experiment. *Experimental Econom.* **9**(2) 79–101.

Topkis, Donald M. 1998. *Supermodularily and complementarity*. Princeton University Press.

Tversky, A., D. Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science* **185** 1124–1131.